

Novell Cluster Services™ 2.0 for Linux Administration Guide

Open Enterprise Server 11

August 6, 2012

Novell.

Legal Notices

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

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

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classification to export, re-export or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in the U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical biological weaponry end uses. 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 © 2007–2012 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.
1800 South Novell Place
Provo, UT 84606
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	11
1 Overview of Novell Cluster Services	13
1.1 Why Should I Use Clusters?	13
1.2 Benefits of Novell Cluster Services	14
1.3 Product Features	14
1.4 Clustering for High-Availability	15
1.5 Shared Disk Scenarios	16
1.5.1 Using Fibre Channel Storage Systems	17
1.5.2 Using iSCSI Storage Systems	18
1.5.3 Using Shared SCSI Storage Systems	19
1.6 Terminology	19
1.6.1 The Cluster	19
1.6.2 Cluster Resources	20
1.6.3 Failover Planning	22
2 What's New or Changed in Novell Cluster Services 2.0	25
2.1 What's New (April 2012 Patches)	25
2.2 What's New (OES 11)	26
3 Planning for a Cluster	29
3.1 Determining Your Design Criteria	29
3.2 Using Cluster Best Practices	30
3.3 Planning the LAN Connectivity	30
3.3.1 VLAN	30
3.3.2 Channel Bonding	31
3.3.3 Spanning Tree Protocol	31
3.3.4 IP Addresses	31
3.3.5 Name Resolution	31
3.4 Planning the Shared Storage Connectivity	32
3.5 Planning the Shared Storage Solution	32
3.6 Planning the eDirectory Deployment	32
3.6.1 Object Location	32
3.6.2 Cluster OU Context	33
3.6.3 Cluster OU Partitioning and Replication	33
3.7 Planning for Shared Storage as Cluster Resources	33
3.8 Planning for OES Services as Cluster Resources	34
4 Planning for Novell Cluster Services	37
4.1 Cluster Administration Requirements	37
4.1.1 Cluster Installation Administrator	37
4.1.2 NCS Proxy User	38
4.1.3 Cluster Administrator or Administrator-Equivalent User	39
4.2 IP Address Requirements	39
4.3 Volume ID Requirements	40
4.4 Hardware Requirements	40

4.5	Xen Virtualization Environments	41
4.6	Software Requirements for Cluster Services	41
4.6.1	Novell Open Enterprise Server 11	41
4.6.2	Novell Cluster Services	42
4.6.3	Novell eDirectory 8.8.6	42
4.6.4	SLP	45
4.6.5	Novell iManager 2.7.4	46
4.6.6	Clusters Plug-in for iManager	46
4.6.7	Storage-Related Plug-Ins for iManager	47
4.6.8	SFCB and CIMOM	48
4.6.9	CASA	49
4.6.10	Web Browser	49
4.7	Software Requirements for Cluster Resources	50
4.7.1	NCP Server for Linux	50
4.7.2	Novell Storage Services File System for Linux	51
4.7.3	LVM Volume Groups and Linux POSIX File Systems	51
4.7.4	NCP Volumes on Linux POSIX File Systems	52
4.7.5	Dynamic Storage Technology Shadow Volume Pairs	52
4.7.6	NCP File Access	52
4.7.7	Novell AFP	52
4.7.8	Novell CIFS	52
4.7.9	Novell Samba	53
4.7.10	Novell Domain Services for Windows	53
4.8	Shared Disk Configuration Requirements	53
4.8.1	Shared Devices	53
4.8.2	SBD Partitions	54
4.8.3	Shared iSCSI Devices	55
4.8.4	Shared RAID Devices	55
4.9	SAN Rules for LUN Masking	55
4.10	Multipath I/O Configuration Requirements	56
4.10.1	Path Failover Settings for Device Mapper Multipath	56
4.10.2	Modifying the Port Down Retry Setting in the modprobe.conf.local File	56
4.10.3	Modifying the Polling Interval, No Path Retry, and Failback Settings in the multipath.conf File	57
4.10.4	Modifying the Port Down Retry and Link Down Retry Settings for an HBA BIOS	58

5 Installing and Configuring Novell Cluster Services on OES 11 59

5.1	Novell Cluster Services Licensing	59
5.2	Extending the eDirectory Schema to Add Cluster Objects	60
5.2.1	Prerequisites for Extending the Schema	60
5.2.2	Extending the Schema	61
5.3	Assigning Install Rights for Container Administrators (or Non-Administrator Users)	62
5.4	Installing Novell Cluster Services	63
5.4.1	Before You Install Novell Cluster Services	63
5.4.2	Installing Novell Cluster Services during a OES 11 Installation	64
5.4.3	Installing Novell Cluster Services on an Existing OES 11 Server	65
5.5	Configuring Novell Cluster Services	66
5.5.1	Initializing and Sharing a Device to Use for the SBD Partition	67
5.5.2	Opening the Novell Open Enterprise Server Configuration Page	69
5.5.3	Using Different LDAP Credentials for the Cluster Configuration	70
5.5.4	Accessing the Novell Cluster Services Configuration Page in YaST	70
5.5.5	Configuring a New Cluster	71
5.5.6	Adding a Node to an Existing Cluster	77
5.6	Configuring Additional Administrators	82
5.7	Installing or Updating the Clusters Plug-in for iManager	82
5.7.1	Prerequisites for Installing or Updating the Clusters Plug-In	82
5.7.2	Installing the Clusters Plug-In	82

5.7.3	Uninstalling and Reinstalling the Clusters Plug-In	83
5.8	Patching Novell Cluster Services	83
5.9	What's Next	83
6	Upgrading Clusters from OES 2 to OES 11	85
6.1	What's New for OES 11	85
6.1.1	NSS Pools	85
6.1.2	Linux POSIX File Systems on an EVMS CSM (Compatibility Only)	86
6.1.3	Linux POSIX File Systems on an LVM Volume Group	86
6.2	Supported Upgrade Paths from OES 2 to OES 11	87
6.3	Requirements and Guidelines for Upgrading Clusters from OES 2	87
6.4	Adding OES 11 Nodes to an OES 2 Cluster (Rolling Cluster Upgrade)	88
7	Configuring Cluster Policies and Priorities	91
7.1	Understanding Cluster Settings	91
7.1.1	Cluster Policies	91
7.1.2	Cluster Protocols	92
7.2	Configuring Quorum Membership and Timeout Properties	93
7.2.1	Quorum Triggers (Number of Nodes)	94
7.2.2	Quorum Triggers (Timeout)	94
7.3	Configuring Cluster Protocol Properties	94
7.3.1	Heartbeat	95
7.3.2	Tolerance	95
7.3.3	Master Watchdog	95
7.3.4	Slave Watchdog	95
7.3.5	Maximum Retransmits	95
7.4	Configuring Cluster Event Email Notification	95
7.5	Viewing the Cluster Node Properties	97
7.6	Modifying the Cluster IP Address and Port Properties	97
7.7	What's Next	97
8	Managing Clusters	99
8.1	Starting and Stopping Novell Cluster Services	100
8.1.1	Starting Novell Cluster Services	100
8.1.2	Stopping Novell Cluster Services	100
8.1.3	Enabling and Disabling the Automatic Start of Novell Cluster Services	100
8.2	Monitoring Cluster and Resource States	101
8.3	Generating a Cluster Configuration Report	103
8.4	Cluster Migrating Resources to Different Nodes	104
8.5	Onlining and Offlining (Loading and Unloading) Cluster Resources from a Cluster Node	104
8.5.1	Guidelines for Onlining and Offlining Resources	104
8.5.2	Using iManager to Online and Offline Resources	105
8.5.3	Using the Cluster Online Command	105
8.5.4	Using the Cluster Offline Command	105
8.5.5	Using the Cluster Migrate Command	105
8.6	Removing (Leaving) a Node from the Cluster	106
8.6.1	Removing One Node at a Time	106
8.6.2	Removing All Nodes at the Same Time	106
8.7	Joining a Node to the Cluster	106
8.8	Shutting Down Linux Cluster Servers When Servicing Shared Storage	106
8.9	Enabling or Disabling Cluster Maintenance Mode	107
8.10	Preventing a Cluster Node Reboot after a Node Shutdown	107

8.11	Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster	108
8.11.1	Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster	108
8.11.2	Moving a Cluster or Changing IP Addresses of Cluster Nodes and Resources	109
8.12	Changing the NCS Proxy User Assignment in the NCS_Management_Group	112
8.13	Adding a Node That Was Previously in the Cluster	113
8.14	Deleting a Cluster Node from a Cluster	113
8.15	Creating or Deleting Cluster SBD Partitions	114
8.15.1	Prerequisites for Creating an SBD Partition	114
8.15.2	Before You Create a Cluster SBD Partition	115
8.15.3	Creating a Non-Mirrored Cluster SBD Partition	115
8.15.4	Creating a Mirrored Cluster SBD Partition	117
8.15.5	Deleting a Non-Mirrored Cluster SBD Partition	120
8.15.6	Deleting a Mirrored Cluster SBD Partition	121
8.15.7	Removing a Segment from a Mirrored Cluster SBD Partition	122
8.16	Configuring STONITH (Feature Preview)	122
8.16.1	Sample Script for HP iLO	123
8.16.2	Sample Script for Web-Based Power Switches	124
8.17	Customizing Cluster Services Management	124

9 Configuring and Managing Cluster Resources 127

9.1	Planning for Cluster Resources	127
9.1.1	Naming Conventions for Cluster Resources	128
9.1.2	Using Parameters with Space Characters in a Cluster Script	128
9.1.3	Script Length Limits	128
9.1.4	Planning Cluster Maintenance	128
9.1.5	Number of Resources	128
9.1.6	Linux POSIX File System Types	129
9.2	Using Cluster Resource Templates	129
9.2.1	Default Resource Templates	129
9.2.2	Creating a Resource Template	131
9.2.3	Synchronizing Locally Modified Resource Templates with eDirectory	132
9.3	Creating Cluster Resources	133
9.4	Configuring a Load Script for a Cluster Resource	133
9.5	Configuring an Unload Script for a Cluster Resource	134
9.6	Enabling Monitoring and Configuring the Monitor Script	135
9.6.1	Understanding Resource Monitoring	135
9.6.2	Configuring Resource Monitoring	137
9.6.3	Example Monitoring Scripts	138
9.6.4	Monitoring Services that Are Critical to Clustering	138
9.7	Configuring the Start, Failover, and Failback Modes for Cluster Resources	139
9.7.1	Understanding Cluster Resource Modes	139
9.7.2	Setting the Start, Failover, and Failback Modes for a Resource	140
9.8	Configuring Preferred Nodes for a Resource	140
9.9	Configuring Resource Priorities for Load Order	142
9.10	Configuring Resource Mutual Exclusion Groups	143
9.10.1	Understanding Resource Mutual Exclusion	143
9.10.2	Setting Up RME Groups	146
9.10.3	Viewing RME Group Settings	146
9.11	Controlling Resource Monitoring	147
9.12	Changing the IP Address of a Cluster Resource	147
9.13	Renaming a Cluster Resource	147
9.14	Deleting Cluster Resources	148
9.14.1	Deleting a Cluster Resource on a Master Node	148
9.14.2	Deleting a Cluster Resource on a Non-Master Node	149
9.15	Additional Information for Creating Cluster Resources	150

9.15.1	Creating Storage Cluster Resources	150
9.15.2	Creating Service Cluster Resources	150
9.15.3	Creating Virtual Machine Cluster Resources	151
10	Creating and Managing Service Cluster Resources	153
11	Configuring Cluster Resources for Shared NSS Pools and Volumes	155
11.1	Requirements for Creating Pool Cluster Resources	155
11.1.1	Novell Cluster Services	156
11.1.2	Resource IP Address	156
11.1.3	Novell eDirectory	156
11.1.4	Novell Storage Services	156
11.1.5	Shared Storage	157
11.1.6	NCP Server for Linux	158
11.1.7	Novell CIFS for Linux	158
11.1.8	Novell AFP for Linux	159
11.1.9	Novell Samba	159
11.1.10	Domain Services for Windows	159
11.2	Guidelines for Using Pool Cluster Resources	159
11.2.1	Guidelines for Shared Pools	159
11.2.2	Guidelines for Shared Volumes	160
11.3	Creating Cluster-Enabled Pools and Volumes	161
11.3.1	Initializing and Sharing a Device	161
11.3.2	Creating a Cluster-Enabled Pool and Volume with iManager	163
11.3.3	Creating a Cluster-Enabled Pool and Volume with NSSMU	165
11.4	Cluster-Enabling an Existing NSS Pool and Its Volumes	167
11.5	Configuring a Load Script for the Shared NSS Pool	171
11.6	Configuring an Unload Script for the Shared NSS Pool	173
11.7	Configuring a Monitor Script for the Shared NSS Pool	174
11.8	Adding Advertising Protocols for NSS Pool Cluster Resources	174
11.9	Adding NFS Export for a Clustered Pool Resource	175
11.9.1	Sample NFS Load Script	175
11.9.2	Sample NFS Unload Script	176
11.10	Mirroring and Cluster-Enabling Shared NSS Pools and Volumes	176
11.10.1	Understanding NSS Mirroring	176
11.10.2	Requirements for NSS Mirroring	178
11.10.3	Creating and Mirroring NSS Pools on Shared Storage	178
11.10.4	Verifying the NSS Mirror Status in the Cluster	180
11.11	Renaming a Pool for a Pool Cluster Resource	180
11.12	Adding a Volume to a Clustered Pool	181
11.13	Changing an Assigned Volume ID	181
11.14	Expanding the Size of a Clustered Pool	182
11.14.1	Planning to Expand a Clustered Pool	182
11.14.2	Increasing the Size of a Clustered Pool by Extending the Size of Its LUN	184
11.14.3	Increasing the Size of a Clustered Pool by Adding Space from a Different Shared LUN	189
11.14.4	Verifying the Expanded Pool Cluster Resource on All Nodes	193
11.15	Deleting NSS Pool Cluster Resources	194
12	Configuring and Managing Cluster Resources for LVM Volume Groups	195
12.1	Requirements for Creating LVM Cluster Resources	195
12.1.1	Novell Cluster Services	196
12.1.2	Linux Logical Volume Manager 2 (LVM2)	196
12.1.3	Clustered Logical Volume Manager Daemon (CLVMD)	196
12.1.4	Resource IP Address	196

12.1.5	Shared Storage Devices	196
12.1.6	All Nodes Must Be Present	196
12.1.7	Mixed Mode OES Clusters (Not Supported)	197
12.1.8	NCP File Access with Novell NCP Server	197
12.1.9	CIFS/Samba File Access with Novell Samba	197
12.2	Initializing a SAN Device	197
12.3	Configuring an LVM Cluster Resource with NSSMU	198
12.3.1	Creating an LVM Volume Group Cluster Resource	200
12.3.2	Configuring the LVM Cluster Resource Settings	202
12.3.3	Viewing or Modifying the LVM Resource Scripts	205
12.3.4	Sample LVM Resource Scripts	206
12.4	Configuring an LVM Cluster Resource with LVM Commands and the Generic File System Template	208
12.4.1	Creating a Shared LVM Volume with LVM Commands	209
12.4.2	Creating a Generic File System Cluster Resource for an LVM Volume Group	213
12.4.3	Sample LVM Resource Load Scripts	218
12.5	Creating a Virtual Server Object for an LVM Volume Group Cluster Resource	220
12.5.1	Creating the Virtual Server Object	221
12.5.2	Modifying the Load Script	221
12.5.3	Modifying the Unload Script	222
12.5.4	Activating the Script Changes	222
12.5.5	Verifying the Virtual Server	222
12.6	Renaming the Mount Point Path for a Clustered LVM Volume	223
12.7	Disabling Clustering for an LVM Volume	225
12.8	Deleting a Clustered LVM Volume Group and Logical Volume	226
12.9	LVM Management Tools	227
12.9.1	Using the YaST Expert Partitioner	227
12.9.2	Using LVM Commands	227

13 Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers 229

13.1	Requirements for Using CSM Cluster Resources on OES 11	229
13.2	Modifying the Scripts for a CSM Resource	230
13.2.1	Offlining the CSM Cluster Resources	231
13.2.2	Configuring the Scripts for the CSM Cluster Resource	231
13.2.3	Sample Load Script for a CSM Resource	232
13.2.4	Sample Unload Script for CSM Resource	233
13.2.5	Sample Monitor Script for a CSM Resource	233
13.3	Modifying the Scripts for a CSM Resource with a Segment Manager On It	234
13.3.1	Offlining the CSM Cluster Resources	235
13.3.2	Converting Linux POSIX Volumes to Compatibility Volumes	236
13.3.3	Deporting the CSM Container	236
13.3.4	Offlining the CSM Cluster Resource	237
13.3.5	Configuring the OES 11 NLVM Configuration File	237
13.3.6	Configuring the Scripts for the CSM Cluster Resource	237
13.3.7	Sample Load Script for a CSM Resource with a Segment Manager on It	240
13.3.8	Sample Unload Script for a CSM Resource with a Segment Manager on It	242
13.3.9	Sample Monitor Script for a CSM Resource with a Segment Manager on It	243
13.4	Configuring the Preferred Nodes and Cluster Settings for a CSM Cluster Resource	243
13.5	Verifying the Load and Unload Scripts	245
13.6	Creating a Virtual Server Object for a CSM Cluster Resource	245
13.6.1	Creating the Virtual Server Object	246
13.6.2	Modifying the Load Script	246
13.6.3	Modifying the Unload Script	247
13.6.4	Activating the Script Changes	247
13.6.5	Verifying the Virtual Server	247
13.7	Deleting a Cluster Resource	248

14	Configuring Novell Cluster Services in a Xen Virtualization Environment	249
14.1	Prerequisites for Xen Host Server Environments	249
14.2	Virtual Machines as Cluster Resources	250
14.2.1	Creating a Xen Virtual Machine Cluster Resource	251
14.2.2	Configuring Virtual Machine Load, Unload, and Monitor Scripts	252
14.2.3	Setting Up Live Migration.	257
14.3	Virtual Machines as Cluster Nodes	258
14.4	Virtual Cluster Nodes in Separate Clusters	259
14.5	Mixed Physical and Virtual Node Clusters	259
14.6	Additional Information	261
15	Troubleshooting Novell Cluster Services	263
15.1	File Location Information	264
15.2	Version Issues	264
15.3	Diagnosing a Poison Pill Condition	264
15.4	Diagnosing Cluster Problems.	264
15.5	Why didn't my resource go to the node I specified?	265
15.6	Cluster Search Times Out (Bad XML Error).	266
15.7	iManager Error 500 Occurs while Connecting to a Cluster	266
15.8	A Device Name Is Required to Create a Cluster Partition	266
15.9	Cluster Resource Goes Comatose Immediately After Migration or Failover	266
15.10	smdr.novell Is Not Registered with SLP for a New Cluster Resource	267
15.11	NSS Takes Up to 10 Minutes to Load When the Server Is Rebooted (Linux)	267
15.12	Could Not Delete This Resource Data_Server (Error 499)	267
15.13	Problem Authenticating to Remote Servers during Cluster Configuration	267
15.14	Problem Connecting to an iSCSI Target	268
15.15	Problem Deleting a Cluster Resource or Clustered Pool	268
15.16	Can't Find the Prevent Cascading Failover Option	268
15.17	Is there a way to uninstall Novell Cluster Services from a server?	268
15.18	Running supportconfig Concurrently on Multiple Nodes Causes Cluster Problems	268
16	Security Considerations	269
16.1	Cluster Administration Rights.	269
16.2	Ports.	269
16.3	Email Alerts	270
16.4	Log Files.	270
A	Console Commands for Novell Cluster Services	271
A.1	Cluster Management Commands	271
A.2	Business Continuity Clustering Commands	276
A.3	extend_schema Command	276
A.3.1	Syntax	276
A.3.2	Options	277
A.3.3	Example	277
A.4	SBD Utility	277
A.4.1	Syntax	278
A.4.2	Options	278
A.4.3	Return Value	279
A.4.4	Examples	279
A.5	Cluster Segment Manager Import/Export (csmport) Utility.	280
A.5.1	Syntax	280

A.5.2	Options	280
A.5.3	Examples	282
A.6	Virtual NCS NCP Server Object (ncs_ncpserv.py) Script	286
A.6.1	Syntax	286
A.6.2	Examples	287
A.7	Configuration Update (ncs-configd.py) Script	288
B	Files for Novell Cluster Services	289
C	Comparing Novell Cluster Services for Linux and NetWare	293
D	Comparing Resources Support for Linux and NetWare	299
E	Documentation Updates	305
E.1	August 6, 2012	305
E.1.1	Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers	305
E.2	July 20, 2012	306
E.2.1	Configuring Cluster Policies and Protocols	306
E.3	June 27, 2012	306
E.3.1	Configuring and Managing Cluster Resources	306
E.3.2	Configuring Cluster Resources for Shared NSS Pools and Volumes	306
E.4	May 31, 2012	307
E.4.1	Configuring Cluster Resources for Shared NSS Pools and Volumes	307
E.5	April 30, 2012	307
E.5.1	Configuring Cluster Policies and Priorities	307
E.5.2	Files for Novell Cluster Services	307
E.5.3	Managing Clusters	308
E.5.4	Planning for Novell Cluster Services	308
E.5.5	Troubleshooting Novell Cluster Services	308
E.6	January 31, 2012	308
E.6.1	Configuring Cluster Resources for Shared NSS Pools and Volumes	309
E.6.2	Planning for Novell Cluster Services	309
E.6.3	Upgrading Clusters from OES 2 to OES 11	309

About This Guide

This guide describes how to install, upgrade, configure, and manage Novell Cluster Services 2.0 for Novell Open Enterprise Server (OES) 11. It is divided into the following sections:

- ♦ Chapter 1, “Overview of Novell Cluster Services,” on page 13
- ♦ Chapter 2, “What’s New or Changed in Novell Cluster Services 2.0,” on page 25
- ♦ Chapter 3, “Planning for a Cluster,” on page 29
- ♦ Chapter 4, “Planning for Novell Cluster Services,” on page 37
- ♦ Chapter 5, “Installing and Configuring Novell Cluster Services on OES 11,” on page 59
- ♦ Chapter 6, “Upgrading Clusters from OES 2 to OES 11,” on page 85
- ♦ Chapter 7, “Configuring Cluster Policies and Priorities,” on page 91
- ♦ Chapter 8, “Managing Clusters,” on page 99
- ♦ Chapter 9, “Configuring and Managing Cluster Resources,” on page 127
- ♦ Chapter 10, “Creating and Managing Service Cluster Resources,” on page 153
- ♦ Chapter 11, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 155
- ♦ Chapter 12, “Configuring and Managing Cluster Resources for LVM Volume Groups,” on page 195
- ♦ Chapter 13, “Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers,” on page 229
- ♦ Chapter 14, “Configuring Novell Cluster Services in a Xen Virtualization Environment,” on page 249
- ♦ Chapter 15, “Troubleshooting Novell Cluster Services,” on page 263
- ♦ Chapter 16, “Security Considerations,” on page 269
- ♦ Appendix A, “Console Commands for Novell Cluster Services,” on page 271
- ♦ Appendix B, “Files for Novell Cluster Services,” on page 289
- ♦ Appendix C, “Comparing Novell Cluster Services for Linux and NetWare,” on page 293
- ♦ Appendix D, “Comparing Resources Support for Linux and NetWare,” on page 299
- ♦ Appendix E, “Documentation Updates,” on page 305

Audience

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

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

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comments feature at the bottom of each page of the online documentation.

Documentation Updates

The latest version of this *Novell Cluster Services for Linux Administration Guide* is available on the [OES 11 Clustering \(High Availability\) documentation Web site \(http://www.novell.com/documentation/oes11/cluster-services.html\)](http://www.novell.com/documentation/oes11/cluster-services.html).

Additional Documentation

For information about converting clusters from NetWare to Linux, see the [OES 11: Novell Cluster Services NetWare to Linux Conversion Guide](#).

For information about creating cluster resources for various Linux services on your OES 11 server, refer to the clustering sections in the individual guides. See the [“Clustering Linux Services” list on the OES 11 Clustering \(High Availability\) documentation Web site \(http://www.novell.com/documentation/oes11/cluster-services.html#clust-config-resources\)](http://www.novell.com/documentation/oes11/cluster-services.html#clust-config-resources).

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

1 Overview of Novell Cluster Services

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

- ♦ [Section 1.1, “Why Should I Use Clusters?,” on page 13](#)
- ♦ [Section 1.2, “Benefits of Novell Cluster Services,” on page 14](#)
- ♦ [Section 1.3, “Product Features,” on page 14](#)
- ♦ [Section 1.4, “Clustering for High-Availability,” on page 15](#)
- ♦ [Section 1.5, “Shared Disk Scenarios,” on page 16](#)
- ♦ [Section 1.6, “Terminology,” on page 19](#)

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 for data and services running on OES 11 servers. You can configure up to 32 OES 11 servers in a high-availability cluster, where resources can be dynamically relocated to any server in the cluster. Each cluster resource can be configured to automatically fail over to a preferred server if there is a failure on the server where it is currently running. In addition, costs are lowered through the consolidation of applications and operations onto a cluster.

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, which allows you to remotely manage the cluster. You can use the Clusters plug-in to manage and monitor clusters and cluster resources, and to configure the settings and scripts for cluster resources.
- ♦ The ability to tailor a cluster to the specific applications and hardware infrastructure that fit your organization.
- ♦ Dynamic assignment and reassignment of server storage as needed.
- ♦ The ability to use email to automatically notify administrators of cluster events and cluster state changes.

1.4 Clustering for High-Availability

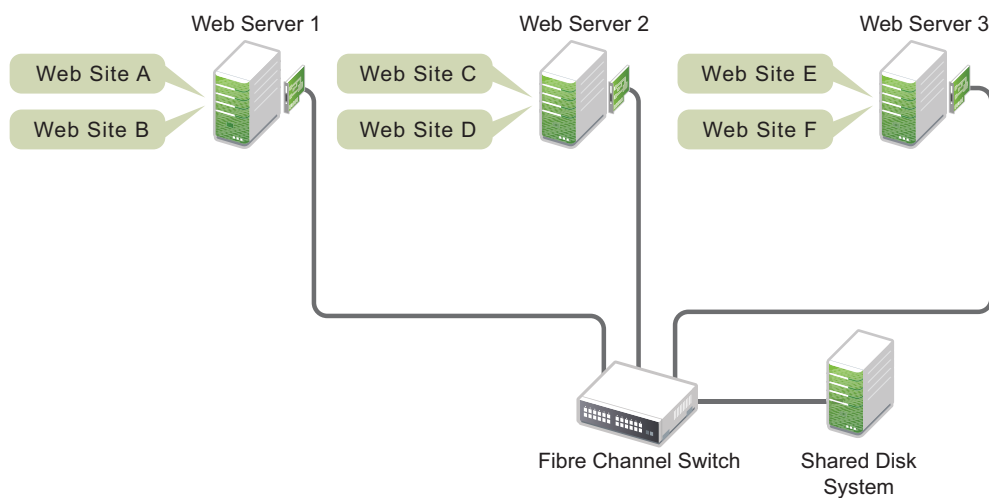
A Novell Cluster Services for Linux cluster consists of the following components:

- ♦ 2 to 32 OES 11 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.

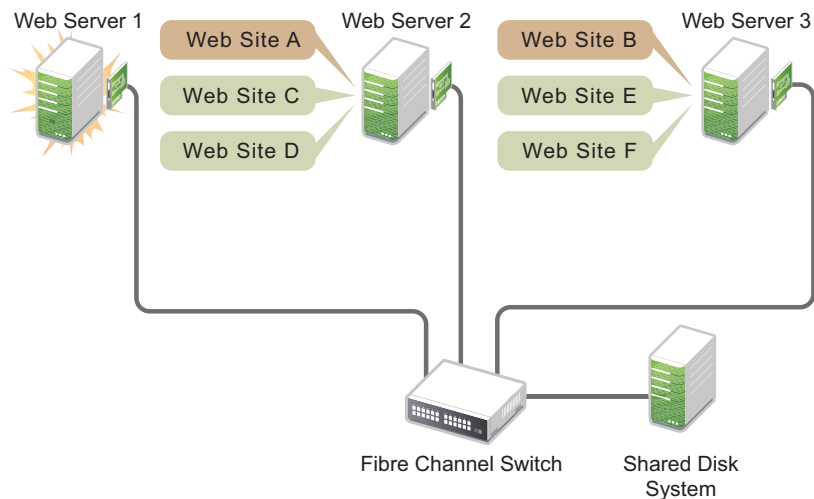
Figure 1-1 *Three-Server Cluster*



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, email, and information lose their connections. [Figure 1-2](#) shows how resources are moved when Web Server 1 fails.

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



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

When you configured the cluster, you decided where the 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, and so on 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

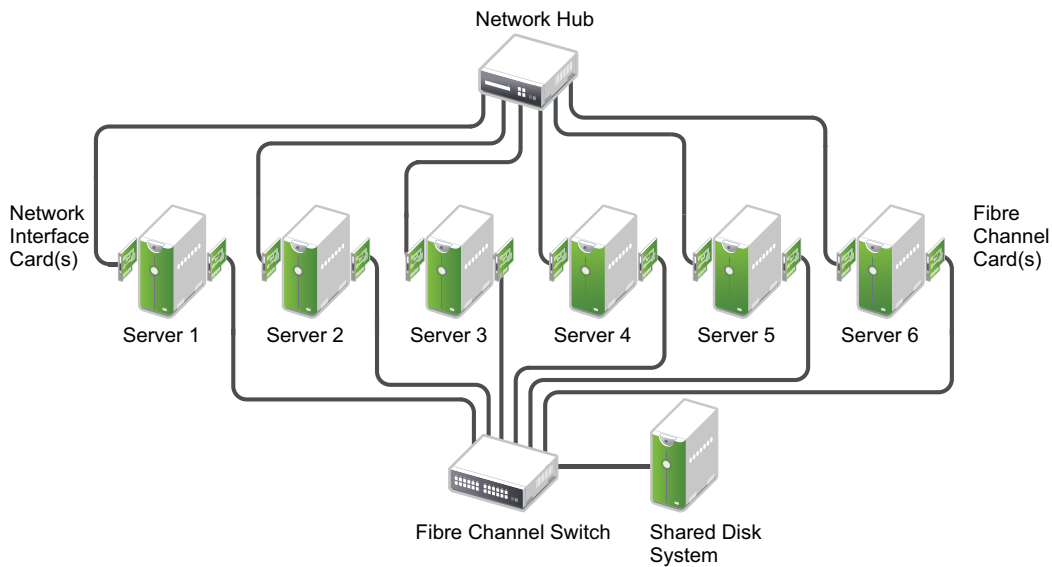
server in the cluster automatically mounts the shared disk directories previously mounted on the failed server. This gives network users continuous access to the directories on the shared disk subsystem.

- ♦ [Section 1.5.1, “Using Fibre Channel Storage Systems,” on page 17](#)
- ♦ [Section 1.5.2, “Using iSCSI Storage Systems,” on page 18](#)
- ♦ [Section 1.5.3, “Using Shared SCSI Storage Systems,” on page 19](#)

1.5.1 Using Fibre Channel Storage Systems

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

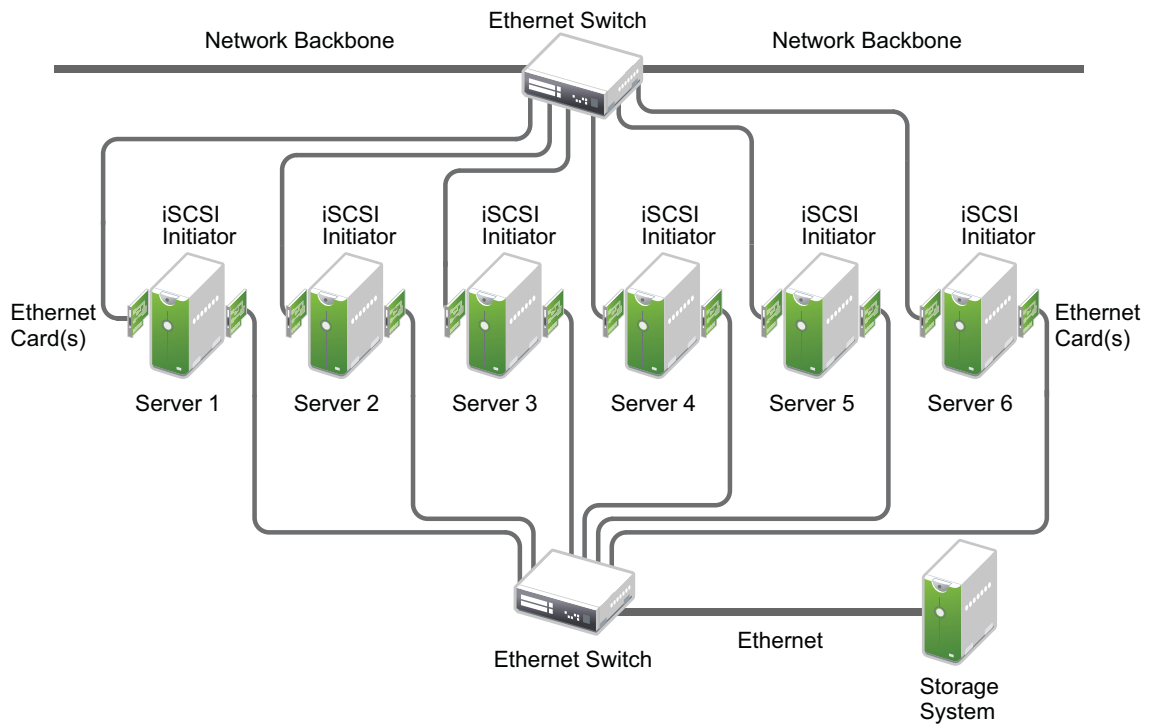
Figure 1-3 Typical Fibre Channel Cluster Configuration



1.5.2 Using iSCSI Storage Systems

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

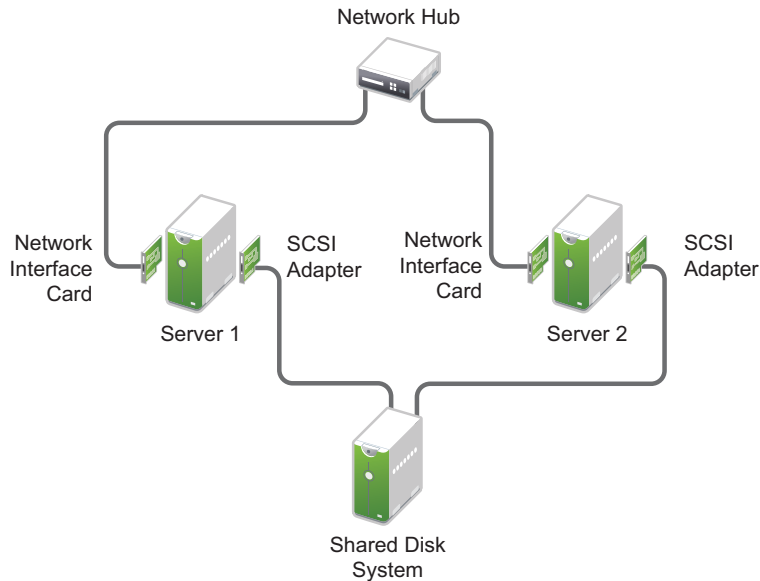
Figure 1-4 Typical iSCSI Cluster Configuration



1.5.3 Using Shared SCSI Storage Systems

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

Figure 1-5 Typical Shared SCSI Cluster Configuration



1.6 Terminology

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

- ♦ [Section 1.6.1, “The Cluster,” on page 19](#)
- ♦ [Section 1.6.2, “Cluster Resources,” on page 20](#)
- ♦ [Section 1.6.3, “Failover Planning,” on page 22](#)

1.6.1 The Cluster

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

- ♦ [“Cluster IP Address” on page 20](#)
- ♦ [“Server IP Address” on page 20](#)
- ♦ [“Master Node” on page 20](#)
- ♦ [“Slave Node” on page 20](#)
- ♦ [“Split-Brain Detector \(SBD\)” on page 20](#)
- ♦ [“Shared Storage” on page 20](#)

Cluster IP Address

The unique static IP address for the cluster.

Server IP Address

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

Master Node

The first server that comes up in an cluster is assigned the cluster IP address and becomes the master node. The master node monitors the health of the cluster nodes. It also synchronizes updates about the cluster to eDirectory. If the master node fails, Cluster Services migrates the cluster IP address to another server in the cluster, and that server becomes the master node.

Slave Node

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

Split-Brain Detector (SBD)

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

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

Shared Storage

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

1.6.2 Cluster Resources

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

- ♦ [“Resource IP Address” on page 21](#)
- ♦ [“NCS Virtual Server” on page 21](#)
- ♦ [“Resource Templates” on page 21](#)
- ♦ [“Service Cluster Resource” on page 21](#)
- ♦ [“Pool Cluster Resource” on page 22](#)
- ♦ [“Linux POSIX Volume Cluster Resource” on page 22](#)

- ♦ [“NCP Volume Cluster Resource” on page 22](#)
- ♦ [“DST Volume Cluster Resource” on page 22](#)
- ♦ [“Cluster Resource Scripts” on page 22](#)

Resource IP Address

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

NCS Virtual Server

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

Resource Templates

A resource template contains the default load, unload, and monitoring scripts and default settings for service or file system cluster resources. Resource templates are available for the following OES 11 Services and file systems.

- ♦ Novell Archive and Versioning
- ♦ Novell DHCP
- ♦ Novell DNS
- ♦ Generic file system (for LVM-based Linux POSIX volumes)
- ♦ NSS file system (for NSS pool resources)
- ♦ Generic IP service
- ♦ Novell iFolder 3.x
- ♦ Novell iPrint
- ♦ MySQL
- ♦ Novell Samba

Personalized templates can also be created. For information, see [Section 9.2, “Using Cluster Resource Templates,” on page 129](#).

Service Cluster Resource

An application or OES 11 service that has been cluster-enabled. The application or service is installed on all nodes in the cluster where the resource can be failed over. The cluster resource includes scripts for loading, unloading, and monitoring. The resource can also contain the configuration information for the application or service.

Pool Cluster Resource

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

Linux POSIX Volume Cluster Resource

A cluster-enabled Linux POSIX volume. The volume is bound to the resource IP address. This provides location independent access to data on the volume via native Linux protocols such as Samba or FTP. You can optionally create an NCS Virtual Server object (`NCS:NCP Server`) for the resource as described in [Section 12.5, “Creating a Virtual Server Object for an LVM Volume Group Cluster Resource,”](#) on page 220.

NCP Volume Cluster Resource

An NCP volume (or share) that has been created on top of a cluster-enabled Linux POSIX volume. The NCP volume is re-created by a command in the resource load script whenever the resource is brought online. The NCP volume is bound to an NCS Virtual Server object (`NCS:NCP Server`) and to the resource IP address. This provides location independent access to the data on the volume for NCP clients in addition to the native Linux protocols such as Samba or FTP. You must create an NCS Virtual Server object (`NCS:NCP Server`) for the resource as described in [Section 12.5, “Creating a Virtual Server Object for an LVM Volume Group Cluster Resource,”](#) on page 220.

DST Volume Cluster Resource

A cluster-enabled Novell Dynamic Storage Technology volume made up of two shared NSS volumes. Both shared volumes are managed in the same cluster resource. The primary volume is bound to an NCS Virtual Server object (`NCS:NCP Server`) and to the resource IP address. This provides location independent access to data on the DST volume for NCP and Novell CIFS clients. (Novell AFP does not support DST volumes.)

If Novell Samba is used instead of Novell CIFS, the cluster resource also manages FUSE and ShadowFS. You point the Samba configuration file to `/media/shadowfs/dst_primary_volume_name` to provide users a merged view of the data.

Cluster Resource Scripts

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

1.6.3 Failover Planning

- ♦ [“Heartbeat” on page 23](#)
- ♦ [“Quorum” on page 23](#)
- ♦ [“Preferred Nodes” on page 23](#)
- ♦ [“Resource Priority” on page 23](#)
- ♦ [“Resource Mutual Exclusion Groups” on page 23](#)

- ♦ [“Failover” on page 23](#)
- ♦ [“Fan-Out Failover” on page 23](#)
- ♦ [“Failback” on page 23](#)
- ♦ [“Cluster Migrate” on page 24](#)
- ♦ [“Leave a Cluster” on page 24](#)
- ♦ [“Join a Cluster” on page 24](#)

Heartbeat

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

Quorum

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

Preferred Nodes

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

Resource Priority

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

Resource Mutual Exclusion Groups

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

Failover

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

Fan-Out Failover

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

Failback

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

Cluster Migrate

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

Leave a Cluster

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

Join a Cluster

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

2 What's New or Changed in Novell Cluster Services 2.0

This section describes enhancements and changes in Novell Cluster Services 2.0 for Novell Open Enterprise Server 11.

- ♦ [Section 2.1, “What’s New \(April 2012 Patches\),” on page 25](#)
- ♦ [Section 2.2, “What’s New \(OES 11\),” on page 26](#)

2.1 What’s New (April 2012 Patches)

In addition to bug fixes, Novell Cluster Services provides the following enhancements and behavior changes in the OES 11 April 2012 Scheduled Maintenance patch:

Cluster Enabling an Existing Pool

In the Clusters plug-in to Novell iManager, when you cluster-enable an existing Novell Storage Services (NSS) pool, the default setting has been modified to automatically deactivate the pool on the master node and online the resource on a preferred node immediately after the pool cluster resource is created. The *Online resource after create* option has been renamed as:

`Deactivate the pool on the master node, and online resource after create`

If you deselect the option, the resource is in an offline state after it is created. At your convenience, you can deactivate the pool on the master node, and bring the resource online on a preferred node.

Previously, the *Online resource after create* option was by default deselected and dimmed, and the resource was automatically placed in an offline state.

For information, see “Cluster-Enabling an Existing NSS Pool and Its Volumes” (http://www.novell.com/documentation/oes11/clus_admin_lx/data/hrt0ekvg.html) in the *OES 11: Novell Cluster Services 2.0 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

STONITH (Feature Preview)

The STONITH (shoot-the-other-node-in-the-head) capability allows Novell Cluster Services to remotely kill a suspect node by using remote power control. Unlike a poison pill, it does not require a response from the suspect node. STONITH is used after a poison pill is issued; it does not replace the poison pill. For information, see “Configuring STONITH” (http://www.novell.com/documentation/oes11/clus_admin_lx/data/stonith.html) in the *OES 11 SP1: Novell Cluster Services 2.0 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

2.2 What's New (OES 11)

Novell Cluster Services 2.0 supports OES 11 services and file systems running on 64-bit SUSE Linux Enterprise Server (SLES) 11 SP1. In addition to bug fixes and performance improvements, it includes the following changes and enhancements:

EVMS Is Deprecated

The Enterprise Volume Management System (EVMS) has been deprecated in SLES 11, and is also deprecated in OES 11. Novell Linux Volume Manager (NLVM) replaces EVMS for managing NetWare partitions under Novell Storage Services (NSS) pools.

NSS Pool Cluster Resources

Novell Cluster Services for OES 11 supports NSS pools that are created on OES 11, OES 2 SP3 and earlier, and NetWare 6.5 SP8.

A new NSS capability supports the GPT partitioning scheme. This allows you to create NSS pools up to 8 TB (terabytes) in size on a single device. Pools created with GPT-partitioned devices are not backwards compatible with prior releases of OES and NetWare. The DOS partitioning scheme is also available and supports devices up to 2 TB in size.

The NSS management tools use the Novell Linux Volume Manager instead of the Enterprise Volume Management System that is used in previous OES releases. The Storage plug-in to iManager and NSSMU can be used to create pool cluster resources. You can also use NLVM commands to create shared pools and volumes at a command prompt or in scripts.

During a rolling cluster upgrade, the existing NSS pool cluster resources can be cluster migrated to any node in the mixed-mode cluster. However, you must not create new NSS pools on OES 11 nodes while you are upgrading the cluster from OES 2 to OES 11. For information, see “NSS Pools” (http://www.novell.com/documentation/oes11/clus_admin_lx/data/btvuil3.html#btvv2vn) in the *OES 11: Novell Cluster Services 2.0 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

LVM Volume Group Cluster Resources

Novell Cluster Services uses a Linux Logical Volume Manager (LVM) volume group and logical volume to create the cluster resource for shared Linux POSIX file systems (such as Ext2/3, ReiserFS, and XFS) on OES 11.

The NSS Management Utility (NSSMU) and Novell Linux Volume Manager (NLVM) commands support creating Linux POSIX file systems and Linux Logical Volume Manager (LVM) volume groups and logical volumes. The tools support both the DOS and the GPT partitioning schemes. The DOS partitioning scheme supports devices up to 2 TB in size. The GPT partitioning scheme supports devices up to 8 zettabytes (ZB, or one billion terabytes). Your actual device size is limited by your storage hardware and the size recognized by your target file system. For information about maximum file system sizes on Linux, see the *SUSE Linux Enterprise Server Technical Information: File System Support* (<http://www.suse.com/products/server/technical-information/#FileSystem>).

You can create a Linux volume group cluster resource by using NSSMU and Novell Linux Volume Manager commands. You can also use native Linux LVM2 commands to create a shared LVM volume group, and then create a resource by using the generic file system (Generic_FS) resource template in Novell iManager, or by using other application resource templates that need shared Linux POSIX file

systems. For information, see “Configuring and Managing Cluster Resources for an LVM volume groups” (http://www.novell.com/documentation/oes11/clus_admin_lx/data/ncsshvollxlv.html) in the *OES 11: Novell Cluster Services 2.0 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

Linux POSIX Cluster Resources with CSM Containers

On OES 2, a Linux POSIX volume cluster resources uses a Cluster Segment Manager (CSM) container on devices that are managed by EVMS. Because EVMS has been deprecated in OES 11, you must modify their scripts and cluster settings so they can run on OES 11 clusters. You cannot create new cluster resources with CSM containers on OES 11 clusters. For information, see “Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers” (http://www.novell.com/documentation/oes11/clus_admin_lx/data/ncsshvollx.html) in the *OES 11: Novell Cluster Services 2.0 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

CSMPORT Utility

The Cluster Segment Manager Import/Export (CSMPORT, `/opt/novell/ncs/bin/csmport`) utility allows you to import and use Linux POSIX volume cluster resources that use CSM containers in OES 11 clusters. After it is configured to run on OES 11, the resource should fail over only to OES 11 nodes. For information about CSMPORT, see “Cluster Segment Manager Import/Export (csmport) Utility” (http://www.novell.com/documentation/oes11/clus_admin_lx/data/csmport.html) in the *OES 11: Novell Cluster Services 2.0 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

Cluster Resource Templates Use LVM Volume Groups for Shared Data Locations

The following cluster resource templates have been modified to use LVM volume groups and logical volumes for cluster resources that share data on Linux POSIX file systems. Previously, the templates used the EVMS Cluster Segment Manager container and Linux POSIX volumes.

OES 11 Application	Cluster Resource Template
Archive and Version Services	AV_Template
DHCP	DHCP_Template (for an NSS pool or for an LVM volume group)
Linux POSIX file system	Generic_FS_Template
iFolder	iFolder_Template (for an NSS pool or for an LVM volume group)
iPrint	IPrint_Template (for an NSS pool or for an LVM volume group)
MySQL 5.x	MySQL_Template
Samba	Samba_Template
Xen virtual machine	Xen_Template

The `DNS_Template` uses an NSS file system. The `Generic_IP_Service` and `XenLive_Template` templates do not use a shared data location.

The monitor scripts for resources that use a Linux volume group were modified to check the status of the LVM logical volume in addition to the file system and the IP address.

Ext3 is the default file system type used in the scripts. The Ext2, Ext3, ReiserFS, and XFS file systems have been tested and are fully supported.

Virtual Server Name for Cluster Resources

The default virtual server name for cluster resources now uses hyphens instead of underscores, such as `MYCLUS-MYPOOL-SERVER`. The suggested name is compliant with the Internet Engineering Task Force (IETF) RFC 1123 standard that allows hostnames to contain only letters, digits, and hyphens. Underscores can still be used in the virtual server name if your network environment supports them.

CIFS Monitor Command in the NSS Monitor Script

Novell CIFS provides a `monitor` command option in OES 11 that provides a restart capability if the `cifs` daemon goes down. If you create a new pool cluster resource with CIFS enabled as an advertising protocol, the following line is added to the resource's monitor script:

```
exit_on_error rcnovell-cifs monitor
```

Previously, the `CIFS status` command was used. You can replace it with the `monitor` command for existing pool cluster resources to take advantage of the CIFS restart capability. For information, see “Configuring a Monitor Script for the Shared NSS Pool” (http://www.novell.com/documentation/oes11/clus_admin_lx/data/bffz0j5.html) in the *OES 11: Novell Cluster Services 2.0 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

Assigned Nodes List

If you attempt to online or migrate a cluster resource to a node that is not in the resource's Assigned Nodes list, the resource stays offline or is not migrated. This change makes the command behavior consistent with the online and migrate options in the Cluster plug-in in iManager. The node that you specify must be running in the cluster and must also be in the resource's Assigned Nodes list.

Previously, if the specified node was not a preferred node, the `cluster online` and `cluster migrate` commands brought the resource online on a node in its Assigned Nodes list.

Order of Servers in the LDAP Server List

When you configure the cluster node in YaST, the LDAP server list is created. The default order is to list the local LDAP server first and others second. In previous OES releases, the default order was based on the IP address.

You cannot change the order of LDAP servers in the list during the cluster node configuration in YaST, but you can modify it later by running the `/opt/novell/ncs/install/ncs_install.py` script. For information, see “Changing the Administrator Credentials or LDAP Server IP Address for a Cluster” (http://www.novell.com/documentation/oes11/clus_admin_lx/data/be4p892.html#bgjnbnv) in the *OES 11: Novell Cluster Services 2.0 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

3 Planning for a Cluster

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

IMPORTANT: For information about the system requirements for installing and using Novell Cluster Services, see [Chapter 4, “Planning for Novell Cluster Services,” on page 37](#).

- ♦ [Section 3.1, “Determining Your Design Criteria,” on page 29](#)
- ♦ [Section 3.2, “Using Cluster Best Practices,” on page 30](#)
- ♦ [Section 3.3, “Planning the LAN Connectivity,” on page 30](#)
- ♦ [Section 3.4, “Planning the Shared Storage Connectivity,” on page 32](#)
- ♦ [Section 3.5, “Planning the Shared Storage Solution,” on page 32](#)
- ♦ [Section 3.6, “Planning the eDirectory Deployment,” on page 32](#)
- ♦ [Section 3.7, “Planning for Shared Storage as Cluster Resources,” on page 33](#)
- ♦ [Section 3.8, “Planning for OES Services as Cluster Resources,” on page 34](#)

3.1 Determining Your Design Criteria

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

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

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

3.2 Using Cluster Best Practices

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

- ♦ Ensure that eDirectory is stable before implementing a cluster.
- ♦ Ensure that you have full Read/Write replicas of the entire eDirectory tree co-located in the data center where you are setting up the cluster.
- ♦ Ensure that IP addresses are unique.
- ♦ IP address assignments should be consistently applied for each cluster and its cluster resources.
- ♦ IP address changes for the cluster and cluster resources should only be made by using the procedure described in [Section 8.11.2, “Moving a Cluster or Changing IP Addresses of Cluster Nodes and Resources,”](#) on page 109.

IP address changes for cluster resources should always be made on the Protocols page of the iManager Clusters plug-in, not directly in load, unload, and monitor scripts. This is the only way to change the IP address on the virtual NCS:NCP Server object in eDirectory.

- ♦ Volume IDs used for a cluster resource must be unique across all nodes.

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

The Novell Client uses the volume ID to access a volume.

- ♦ Each node’s configuration must consider the configuration requirements for each of the services it is intended to host.
- ♦ Create failover matrixes for each cluster resource so that you know what service is supported and which nodes are the preferred nodes for failover.

3.3 Planning the LAN Connectivity

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

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

3.3.1 VLAN

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

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

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

3.3.2 Channel Bonding

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

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

3.3.3 Spanning Tree Protocol

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

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

3.3.4 IP Addresses

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

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

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

3.3.5 Name Resolution

Ensure that SLP is properly configured for name resolution. For information, see [Section 4.6.4, “SLP,” on page 45](#).

3.4 Planning the Shared Storage Connectivity

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

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

- ♦ Use host-based multipath I/O management. For information, see the following resources:
 - ♦ [Section 4.10, “Multipath I/O Configuration Requirements,” on page 56](#)
 - ♦ [“Managing Multipath I/O for Devices” \(http://www.suse.com/documentation/sles11/stor_admin/data/multipathing.html\)](http://www.suse.com/documentation/sles11/stor_admin/data/multipathing.html) in the *SLES 11 SP1: Storage Administration Guide* (http://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html)
- ♦ Connect each node via two fabrics to the storage area network (SAN).
- ♦ Use redundant SAN connections to provide fault-tolerant connectivity between the cluster nodes and the shared storage devices.
- ♦ Use LUN masking to exclusively assign each LUN to one or more host connections. For information, see [Section 4.9, “SAN Rules for LUN Masking,” on page 55](#).

3.5 Planning the Shared Storage Solution

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

- ♦ For maximum flexibility, we recommend that you create only one cluster resource per LUN.

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

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

3.6 Planning the eDirectory Deployment

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

- ♦ [Section 3.6.1, “Object Location,” on page 32](#)
- ♦ [Section 3.6.2, “Cluster OU Context,” on page 33](#)
- ♦ [Section 3.6.3, “Cluster OU Partitioning and Replication,” on page 33](#)

3.6.1 Object Location

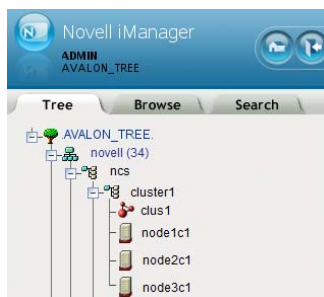
Cluster nodes and Cluster objects can exist in any container in the eDirectory tree. The Virtual Server object for the cluster and the objects for cluster resources are automatically created in the eDirectory context of the server where the cluster resource is created and cluster-enabled.

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

3.6.2 Cluster OU Context

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

Figure 3-1 Example: Cluster1 Container and Its Objects



3.6.3 Cluster OU Partitioning and Replication

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

If you do not want to put a replica of eDirectory on the node, you must configure one or multiple LDAP servers for the node to use. The LDAP servers must have a master replica or a Read/Write replica of eDirectory. For information about how to modify the LDAP server list that is used by a cluster, see [Section 8.11.1, "Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster,"](#) on page 108.

3.7 Planning for Shared Storage as Cluster Resources

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

File System or Storage Solution	For information, see
Novell Storage Services (NSS) pools	Chapter 11, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 155
Linux POSIX volumes	Chapter 13, "Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers," on page 229
NCP (NetWare Control Protocol) volumes (NCP shares on cluster-enabled Linux POSIX volumes)	"Configuring NCP Volumes with Novell Cluster Services" in the <i>OES 11: NCP Server for Linux Administration Guide</i>
Dynamic Storage Technology (DST) volumes (NSS volumes configured in a shadow volume pair)	"Configuring DST Shadow Volumes with Novell Cluster Services" in the <i>OES 11: Dynamic Storage Technology Administration Guide</i>

3.8 Planning for OES Services as Cluster Resources

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

Service	For information, see
Apache Web Server	"Apache Web Server" in the <i>OES 11: Novell Cluster Services NetWare to Linux Conversion Guide</i>
Apple Filing Protocol (Novell AFP)	"Configuring AFP with Novell Cluster Services for an NSS File System" in the <i>OES 11: Novell AFP Administration Guide</i>
Archive and Version Services	"Configuring Archive and Version Service for Novell Cluster Services" in the <i>OES 11: Novell Archive and Version Services 2.1 Administration Guide</i>
Certificate Server (eDirectory Server Certificates)	"eDirectory Server Certificates" in the <i>OES 11: Novell Cluster Services NetWare to Linux Conversion Guide</i>
CIFS (Novell CIFS)	"Configuring CIFS with Novell Cluster Services for an NSS File System" in the <i>OES 11: Novell CIFS for Linux Administration Guide</i>
DFS VLDB (Novell Distributed File Services Volume Location Database)	"Clustering Novell Distributed File Services" in the <i>OES 11 SP1: Novell Distributed File Services Administration Guide for Linux</i>
DHCP Server on a Linux POSIX volume	"Configuring DHCP with Novell Cluster Services for the Linux File System" in the <i>OES 11: Novell DNS/DHCP Services for Linux Administration Guide</i>
DHCP Server on an NSS volume	"Configuring DHCP with Novell Cluster Services for the NSS File System" in the <i>OES 11: Novell DNS/DHCP Services for Linux Administration Guide</i>
DNS Server	"Configuring DNS with Novell Cluster Services" in the <i>OES 11: Novell DNS/DHCP Services for Linux Administration Guide</i>
iFolder 3.x	"Clustering iFolder Servers with Novell Cluster Services for Linux" in the <i>Novell iFolder 3.9 Administration Guide</i>
iPrint	"Configuring iPrint with Novell Cluster Services" in the <i>OES 11: iPrint Linux Administration Guide</i>
MySQL	"High Availability and Scalability" in the <i>MySQL 5.x Reference Manual</i> (http://dev.mysql.com/doc/refman/5.5/en/ha-overview.html) "MySQL" in the <i>OES 11: Novell Cluster Services NetWare to Linux Conversion Guide</i>
NetStorage	"Configuring NetStorage with Novell Cluster Services" in the <i>OES 11: NetStorage Administration Guide for Linux</i> .
PureFTP	"Cluster Enabling Pure-FTPd in an OES 11 Environment" in the <i>OES 11: Planning and Implementation Guide</i>

Service	For information, see
QuickFinder 5.0.x (Server Synchronization Feature)	“QuickFinder Server” in the OES 11: Novell Cluster Services NetWare to Linux Conversion Guide
Samba (Novell Samba)	“Configuring Samba for LVM Volume Groups and Novell Cluster Services” in the OES 11: Novell Samba Administration Guide

4 Planning for Novell Cluster Services

This section describes the requirements for installing and using Novell Cluster Services on Novell Open Enterprise Server 11 servers.

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

- ♦ [Section 4.1, “Cluster Administration Requirements,”](#) on page 37
- ♦ [Section 4.2, “IP Address Requirements,”](#) on page 39
- ♦ [Section 4.3, “Volume ID Requirements,”](#) on page 40
- ♦ [Section 4.4, “Hardware Requirements,”](#) on page 40
- ♦ [Section 4.5, “Xen Virtualization Environments,”](#) on page 41
- ♦ [Section 4.6, “Software Requirements for Cluster Services,”](#) on page 41
- ♦ [Section 4.7, “Software Requirements for Cluster Resources,”](#) on page 50
- ♦ [Section 4.8, “Shared Disk Configuration Requirements,”](#) on page 53
- ♦ [Section 4.9, “SAN Rules for LUN Masking,”](#) on page 55
- ♦ [Section 4.10, “Multipath I/O Configuration Requirements,”](#) on page 56

4.1 Cluster Administration Requirements

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

- ♦ [Section 4.1.1, “Cluster Installation Administrator,”](#) on page 37
- ♦ [Section 4.1.2, “NCS Proxy User,”](#) on page 38
- ♦ [Section 4.1.3, “Cluster Administrator or Administrator-Equivalent User,”](#) on page 39

4.1.1 Cluster Installation Administrator

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

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

- ♦ [“eDirectory Schema Administrator” on page 38](#)
- ♦ [“Container Administrator” on page 38](#)

eDirectory Schema Administrator

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

For instructions, see [Section 5.2, “Extending the eDirectory Schema to Add Cluster Objects,” on page 60](#).

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

Container Administrator

After the schema has been extended, the container administrator (or non-administrator user) needs the following eDirectory rights to install 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 instructions, see [Section 5.3, “Assigning Install Rights for Container Administrators \(or Non-Administrator Users\),” on page 62](#)

4.1.2 NCS Proxy User

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

Novell Cluster Services supports the OES Common Proxy User enablement feature of Novell eDirectory 8.8.6 or later. If the OES Common Proxy user is enabled in eDirectory when you configure the cluster, you can specify whether to use the Common Proxy user, the LDAP Admin user, or another administrator user. The specified user is automatically assigned to the NCS_Management group that resides in the Cluster object container. This accommodates the server-specific common user for each of the nodes. As a group member, the assigned user has the necessary rights for configuring the cluster and cluster resources and for exchanging information with eDirectory.

You can modify this default administrator user name or password for the user name assigned as the NCS Proxy User after the install by following the procedure in [Section 8.11, “Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster,” on page 108](#).

Consider the following caveats for the three proxy user options:

- ♦ [“OES Common Proxy User” on page 39](#)
- ♦ [“LDAP Admin User” on page 39](#)
- ♦ [“Another Administrator User” on page 39](#)

OES Common Proxy User

If you specify the OES Common Proxy user for a cluster and later disable the Common Proxy user in eDirectory, the LDAP Admin user is automatically assigned to the NCS_Management group and the Common Proxy user is automatically removed from the group.

If a proxy user is renamed, moved, or deleted in eDirectory, eDirectory takes care of the changes needed to modify the user information in the NCS_Management group.

If a cluster node is removed from the tree, the proxy user for that server is among the group of cluster objects that needs to be deleted from the eDirectory tree.

For information about enabling or disabling the OES Common Proxy User, see the [OES 11: Installation Guide](#). For caveats and troubleshooting information for the OES Common Proxy user, see the [OES 11: Planning and Implementation Guide](#).

LDAP Admin User

If you specify the LDAP Admin user as the NCS Proxy User, you typically continue using this identity while you set up the cluster and cluster resources. After the cluster configuration is completed, you create another user identity to use for this purpose, and grant that user sufficient administrator rights as specified in [“Cluster Administrator or Administrator-Equivalent User” on page 39](#).

Another Administrator User

You can specify an existing user name and password to use for the NCS Proxy user. Novell Cluster Services adds this user name to the NCS_Management group.

4.1.3 Cluster Administrator or Administrator-Equivalent User

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

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

4.2 IP Address Requirements

- ☐ Each server in the cluster must be configured with a unique static IP address.

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

4.3 Volume ID Requirements

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

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

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

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

4.4 Hardware Requirements

The following hardware requirements for installing 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 additional memory on each server in the cluster
- ☐ One non-shared device on each server to be used for the operating system
- ☐ At least one network card per server in the same IP subnet

In addition, each server must meet the requirements for Novell Open Enterprise Server 11. For information, see [“Meeting All Server Software and Hardware Requirements”](#) in the *OES 11: Installation Guide*.

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

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

4.5 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 14](#), “Configuring Novell Cluster Services in a Xen Virtualization Environment,” on page 249.

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 14.2](#), “Virtual Machines as Cluster Resources,” on page 250.

4.6 Software Requirements for Cluster Services

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

- ♦ [Section 4.6.1](#), “Novell Open Enterprise Server 11,” on page 41
- ♦ [Section 4.6.2](#), “Novell Cluster Services,” on page 42
- ♦ [Section 4.6.3](#), “Novell eDirectory 8.8.6,” on page 42
- ♦ [Section 4.6.4](#), “SLP,” on page 45
- ♦ [Section 4.6.5](#), “Novell iManager 2.7.4,” on page 46
- ♦ [Section 4.6.6](#), “Clusters Plug-in for iManager,” on page 46
- ♦ [Section 4.6.7](#), “Storage-Related Plug-Ins for iManager,” on page 47
- ♦ [Section 4.6.8](#), “SFCB and CIMOM,” on page 48
- ♦ [Section 4.6.9](#), “CASA,” on page 49
- ♦ [Section 4.6.10](#), “Web Browser,” on page 49

4.6.1 Novell Open Enterprise Server 11

Novell Cluster Services 2.0 for Linux supports Novell Open Enterprise Server 11 running on 64-bit SUSE Linux Enterprise Server 11 SP1. Novell Cluster Services is one of the OES Services patterns for OES 11.

IMPORTANT: Ensure that you have installed the latest patches for SUSE Linux Enterprise Server 11 SP1. Clustered LVM volume groups require Linux kernel version 2.6.32.45-0.3 or later.

We recommend having uniform nodes in the cluster. The same release of OES 11 must be installed and running on each node in the cluster.

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

Upgrading from	For information, see:
OES 2 SP2 or SP3	Chapter 6, “Upgrading Clusters from OES 2 to OES 11,” on page 85
NetWare 6.5 SP8	OES 11: Novell Cluster Services NetWare to Linux Conversion Guide

4.6.2 Novell Cluster Services

Novell Cluster Services is required for creating and managing clusters and shared resources on your OES 11 servers. It is one of the OES Services patterns on the OES 11 Add-On disk.

4.6.3 Novell eDirectory 8.8.6

Novell eDirectory 8.8.6 (or later) is required for managing the Cluster object and Cluster Node objects for Novell Cluster Services. eDirectory must be installed and running in the same tree where you create the cluster. eDirectory can be installed on any node in the cluster, on a separate server, or in a separate cluster. You can install an eDirectory master replica or replica in the cluster, but it is not required to do so for Novell Cluster Services.

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

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

- ♦ [“eDirectory Tree” on page 42](#)
- ♦ [“eDirectory Context” on page 42](#)
- ♦ [“Cluster Object Container” on page 43](#)
- ♦ [“Cluster Objects Stored in eDirectory” on page 44](#)
- ♦ [“LDAP Server List” on page 45](#)

eDirectory Tree

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

eDirectory Context

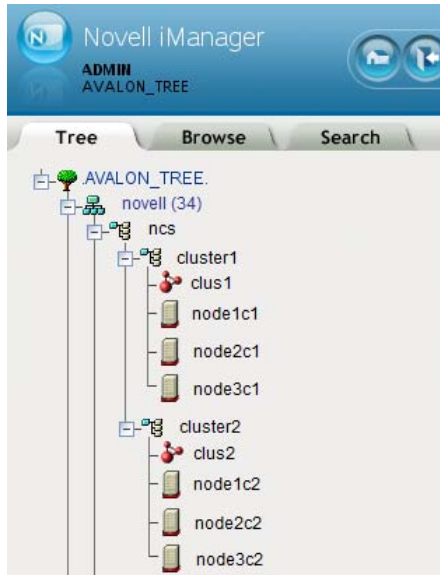
If you are creating a new cluster, the eDirectory context where the new Cluster object will reside must be an existing context. Specifying a new context during the Novell Cluster Services configuration does not create a new context.

Cluster Object Container

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

For example, [Figure 4-1](#) shows an example where all clusters are configured in the ncs organizational unit. Within the container, each cluster is in its own organizational unit, and the Server objects for the nodes are in the same container as the Cluster object:

Figure 4-1 Same Container for Cluster Object and Server Objects






If the servers in the cluster are in separate eDirectory containers, the user that administers the cluster must have rights to the cluster server containers and to the containers where any cluster-enabled pool objects are stored. You can do this by adding trustee assignments for the cluster administrator to a parent container of the containers where the cluster server objects reside. See “[eDirectory Rights](http://www.novell.com/documentation/edir88/edir88/data/fbachifb.html)” (<http://www.novell.com/documentation/edir88/edir88/data/fbachifb.html>) in the *eDirectory 8.8 Administration Guide* for more information.

Cluster Objects Stored in eDirectory

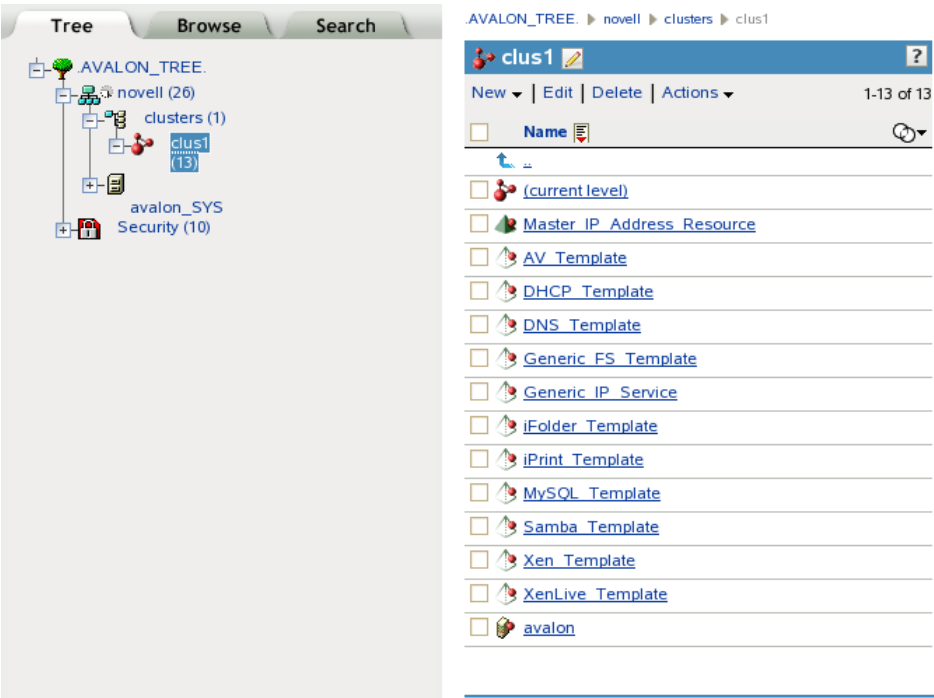
Table 4-1 shows the cluster objects that are automatically created and stored in eDirectory under the Cluster object (🍷) after you create a cluster:

Table 4-1 Cluster Objects

Icon	eDirectory Object
	Master_IP_Address_Resource
	Cluster Node object (<i>servername</i>)
	Resource Template objects. There are 11 default templates: AV_Template DHCP_Template DNS_Template Generic_FS_Template Generic_IP_Service iFolder_Template iPrint_Template MySQL_Template Samba_Template Xen_Template XenLive_Template




For example, Figure 4-2 shows the 13 default eDirectory objects that are created in the Cluster container as viewed from the Tree view in iManager:

Figure 4-2 Tree View of the Default eDirectory Objects in the Cluster






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

Table 4-2 Cluster Resource Objects

Icon	eDirectory Object
	Cluster Node object (<i>servername</i>)
	NSS Pool Resource object (<i>poolname_SERVER</i>)
	Resource object

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

Table 4-3 BCC Cluster Resource Objects

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

LDAP Server List

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

- ♦ local to the cluster
- ♦ closest physical read/write replica

For information about configuring a list of LDAP servers for the cluster, see [Section 8.11.1, “Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster,”](#) on page 108.

4.6.4 SLP

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

SLP is not automatically installed when you select Novell Cluster Services. SLP is installed as part of the Novell eDirectory configuration during the OES 11 install. You can enable and configure SLP on the eDirectory Configuration - NTP & SLP page. For information, see [“Specifying SLP Configuration Options”](#) in the *OES 11: Installation Guide*.

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

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

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

4.6.5 Novell iManager 2.7.4

Novell iManager 2.7.4 (or later) is required for configuring and managing clusters on OES 11. iManager must be installed on at least one computer in the same tree as the cluster. It can be installed in the cluster or not in the cluster. For information about using iManager, see the [iManager 2.7x documentation Web site](http://www.novell.com/documentation/imanager27/index.html) (<http://www.novell.com/documentation/imanager27/index.html>).

For SFCB (Small Footprint CIM Broker) and CIMOM requirements, see [Section 4.6.8, “SFCB and CIMOM,” on page 48](#).

For browser configuration requirements, see [“Web Browser” on page 49](#).

4.6.6 Clusters Plug-in for iManager

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

The following components must be installed in iManager:

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

For information, see [“Storage-Related Plug-Ins for iManager” on page 47](#).

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

The Clusters plug-in also provides an integrated management interface for Novell Business Continuity Clustering (BCC). The additional interface is present only if BCC is installed on the server. See the following table for information about the versions of BCC that are supported. BCC is sold separately from OES. For purchasing information, see the [BCC product page](http://www.novell.com/products/businesscontinuity/) (<http://www.novell.com/products/businesscontinuity/>).

BCC Release	OS Platform Support	iManager and Clusters Plug-In
BCC 2.0	OES 11	Planned for a future release.
BCC 1.2.2	OES 2 SP3	iManager 2.7.4 or later Requires the Clusters plug-in for OES 2 SP3 and the OES 2 SP3 April 2011 Scheduled Maintenance patch. See BCC 1.2.2: Administration Guide for OES 2 SP3 (http://www.novell.com/documentation/bcc/bcc122_admin_lx/data/bookinfo.html) .

BCC Release	OS Platform Support	iManager and Clusters Plug-In
BCC 1.2.1	OES 2 SP2	iManager 2.7.3 or later Requires the Clusters plug-in in the January 2010 Maintenance Patch for OES 2 SP2 Linux. See BCC 1.2.1: Administration Guide for OES 2 SP2 (http://www.novell.com/documentation/bcc/bcc121_admin_lx/data/bookinfo.html).
BCC 1.2.0	OES 2 SP1 Linux	iManager 2.7.3 or later Requires the Clusters plug-in patch for February 3, 2009 or later version. See BCC 1.2: Administration Guide for OES 2 SP1 (http://www.novell.com/documentation/bcc/bcc12_admin_lx/data/bookinfo.html).
BCC 1.1 SP2	NetWare 6.5 SP8	iManager 2.7.2 or later Requires the Clusters plug-in released in OES 2 SP1 Linux or NetWare 6.5 SP8, or a later version. See BCC 1.1 SP2 Administration Guide for NetWare 6.5 SP8 (http://www.novell.com/documentation/bcc/bcc11_admin_nw/data/bktitle.html).

4.6.7 Storage-Related Plug-Ins for iManager

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

Product	Plug-In	NPM File
Novell Apple Filing Protocol (AFP)	File Protocols > AFP	<code>afpmgmt.rpm</code>
Novell Archive and Version Services	Archive Versioning	<code>arkmgmt.rpm</code>
Novell CIFS	File Protocols > CIFS	<code>cifsmgmt.rpm</code>
Novell Cluster Services	Clusters	<code>ncsmgmt.rpm</code>
Novell Distributed File Services	Distributed File Services	<code>dfsmgmt.rpm</code>
Novell Storage Services	Storage	<code>nssmgmt.rpm</code>

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

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

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

The plug-in files are included on the installation disk. The latest 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 2.7.4 Administration Guide*.

For information about working with storage-related plug-ins for iManager, see “[Understanding Storage-Related Plug-Ins](#)” in the *OES 11: NSS File System Administration Guide for Linux*.

4.6.8 SFCB and CIMOM

The Small Footprint CIM Broker (SFCB) replaces OpenWBEM for CIMOM activities in OES 11. OES11 and SLES 11 offer SFCB as the default CIMOM and CIM clients. When you install any OES components that depend on WBEM, SFCB and all of its corresponding packages are installed with the components. For information, see “[Small Footprint CIM Broker \(SFCB\)](#)” in the *OES 11: Planning and Implementation Guide*.

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

Port 5989 is the default setting for Secure HTTP (HTTPS) communications. If you are using a firewall, the port must be opened for CIMOM communications.

The Clusters plug-in (and all other storage-related plug-ins) for iManager require CIMOM connections for tasks that transmit sensitive information (such as a user name and password) between iManager and the `_admin` volume on the OES 11 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.

IMPORTANT: SFCB is automatically PAM-enabled for Linux User Management (LUM) as part of the OES 11 installation. Users not enabled for LUM cannot use the CIM providers to manage OES. The user name that you use to log in to Novell iManager when you manage a cluster and the BCC cluster must be a Novell eDirectory user name that has been LUM-enabled.

For more information about the permissions and rights needed by the administrator user, see [Section 4.1, “Cluster Administration Requirements,” on page 37](#).

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 SFCB is not running.

You can use the `rscfcb` command to help resolve CIMOM and SFCB issues:

To perform this task	At a terminal console prompt, enter as the root user
To start SFCB	<code>rscfcb start</code>
To stop SFCB	<code>rscfcb stop</code>
To check SFCB status	<code>rscfcb status</code>
To restart SFCB	<code>rscfcb restart</code>

For more information, see “Web Based Enterprise Management using SFCB” (http://www.suse.com/documentation/sles11/book_sle_admin/data/cha_wbem.html) in the *SUSE Linux Enterprise Server 11 Administration Guide* (http://www.suse.com/documentation/sles11/book_sle_admin/data/book_sle_admin.html).

4.6.9 CASA

Novell Cluster Services requires CASA (Common Authentication Service Adapter) to be installed and running on each node in the cluster.

The following table contains some useful commands for resolving CASA issues:

To perform this task	At a terminal console prompt, enter as the root user
To check CASA status	<code>rcmicasad status</code>
To check that CASA is running correctly	<code>CASAccli -l</code>
To restart CASA	<code>rcmicasad restart</code>

4.6.10 Web Browser

For a information about supported Web browsers for Novell iManager, see “Using a Supported Web Browser” in the *Novell iManager 2.7.4 Administration Guide*.

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

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

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

4.7 Software Requirements for Cluster Resources

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

- ♦ [Section 4.7.1, “NCP Server for Linux,” on page 50](#)
- ♦ [Section 4.7.2, “Novell Storage Services File System for Linux,” on page 51](#)
- ♦ [Section 4.7.3, “LVM Volume Groups and Linux POSIX File Systems,” on page 51](#)
- ♦ [Section 4.7.4, “NCP Volumes on Linux POSIX File Systems,” on page 52](#)
- ♦ [Section 4.7.5, “Dynamic Storage Technology Shadow Volume Pairs,” on page 52](#)
- ♦ [Section 4.7.6, “NCP File Access,” on page 52](#)
- ♦ [Section 4.7.7, “Novell AFP,” on page 52](#)
- ♦ [Section 4.7.8, “Novell CIFS,” on page 52](#)
- ♦ [Section 4.7.9, “Novell Samba,” on page 53](#)
- ♦ [Section 4.7.10, “Novell Domain Services for Windows,” on page 53](#)

4.7.1 NCP Server for Linux

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

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

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

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

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

For information, see [“Configuring Cross-Protocol File Locks for NCP Server”](#) in the *OES 11: NCP Server for Linux Administration Guide*.

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

For information about configuring and managing NCP Server for Linux, see the [OES 11: 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 11: NCP Server for Linux Administration Guide](#).

4.7.2 Novell Storage Services File System for Linux

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

- ♦ Initializing and sharing devices used for the SBD (split-brain detector) and for shared pools. For information, see [Section 4.8.2, “SBD Partitions,” on page 54](#).
- ♦ Creating and cluster-enabling a shared pool. For information, see [Chapter 11, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 155](#).
- ♦ Creating and cluster-enabling a shared Linux Logical Volume Manager (LVM) volume group. For information, see [Chapter 12, “Configuring and Managing Cluster Resources for LVM Volume Groups,” on page 195](#).

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

- ♦ NetWare Core Protocol (NCP), which is selected by default and is mandatory for NSS. For information, see “[NCP Server for Linux](#)” on page 50.
- ♦ Novell Apple Filing Protocol (AFP). For information, see “[Novell AFP](#)” on page 52.
- ♦ Novell CIFS. For information, see “[Novell CIFS](#)” on page 52.

4.7.3 LVM Volume Groups and Linux POSIX File Systems

Novell Cluster Services supports creating shared cluster resources on Linux Logical Volume Manager (LVM) volume groups. You can configure Linux POSIX file systems on the LVM volume group, such as Ext3, XFS, and ReiserFS. LVM and Linux POSIX file systems are automatically installed as part of the OES 11 installation.

After the cluster is configured, you can create LVM volume group cluster resources as described in [Chapter 12, “Configuring and Managing Cluster Resources for LVM Volume Groups,” on page 195](#).

NCP Server is required if you want to create a virtual server name (NCS:NCP Server object) for the cluster resource. You can add an NCP volume (an NCP share) on the Linux POSIX file system to give users NCP access to the data. For information, see [Section 4.7.1, “NCP Server for Linux,” on page 50](#).

4.7.4 NCP Volumes on Linux POSIX File Systems

After you cluster-enable an LVM volume group, Novell Cluster Services supports creating NCP volumes on the volume group's Linux POSIX file systems. NCP Server is required. For information, see [Section 4.7.1, “NCP Server for Linux,” on page 50](#).

For information about creating and cluster-enabling NCP volumes, see “[Configuring NCP Volumes with Novell Cluster Services](#)” in the *OES 11: NCP Server for Linux Administration Guide*.

4.7.5 Dynamic Storage Technology Shadow Volume Pairs

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

For information about creating and cluster-enabling Dynamic Storage Technology volumes on Linux, see “[Configuring DST Shadow Volumes with Novell Cluster Services](#)” in the *OES 11: Dynamic Storage Technology Administration Guide*.

4.7.6 NCP File Access

Novell Cluster Services requires NCP file access to be enabled for cluster-enabled NSS volumes, NCP volumes, and DST volumes, even if users do not access files via NCP. This is required to support access control via the Novell Trustee model. For information, see [Section 4.7.1, “NCP Server for Linux,” on page 50](#).

4.7.7 Novell AFP

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

Novell AFP is not required to be installed when you install Novell Cluster Services, but it must be installed and running in order for the *AFP* option to be available as an advertising protocol for the NSS pool cluster resource. The AFP daemon should also be running before you bring resources online that have AFP enabled.

To install Novell AFP, select the *Novell AFP* option from the OES Services list during the install. For information about configuring and managing the Novell AFP service, see the *OES11: Novell AFP Administration Guide*.

4.7.8 Novell CIFS

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

Novell CIFS is not required to be installed when you install Novell Cluster Services, but it must be installed and running in order for the *CIFS Virtual Server Name* and *CIFS* option to be available as an advertising protocol for the NSS pool cluster resource. The CIFS daemon should also be running before you bring resources online that have CIFS enabled.

To install Novell CIFS, select the *Novell CIFS* option from the OES Services list during the install. For information about configuring and managing the Novell CIFS service, see the [OES 11: Novell CIFS for Linux Administration Guide](#).

4.7.9 Novell Samba

Novell Cluster Services supports using Novell Samba as an alternative to using Novell CIFS. It provides CIFS/Samba access for users. Users must be enabled with Linux User Management.

Samba is not integrated as an advertising protocol option for NSS pool cluster resources.

For information about configuring and managing Novell Samba, see the following sections in the [OES 11: Novell Samba Administration Guide](#):

- ♦ [“Configuring Samba for LVM Volume Groups and Novell Cluster Services”](#)
- ♦ [“Configuring Samba for NSS Pools and Novell Cluster Services”](#)

4.7.10 Novell Domain Services for Windows

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

For information using Domain Services for Windows, see the [OES 11: Domain Services for Windows Administration Guide](#).

4.8 Shared Disk Configuration Requirements

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

Ensure that your shared storage devices meet the following requirements:

- ♦ [Section 4.8.1, “Shared Devices,” on page 53](#)
- ♦ [Section 4.8.2, “SBD Partitions,” on page 54](#)
- ♦ [Section 4.8.3, “Shared iSCSI Devices,” on page 55](#)
- ♦ [Section 4.8.4, “Shared RAID Devices,” on page 55](#)

4.8.1 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 configuring 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 11 documentation or the shared disk system documentation for troubleshooting information.

Prepare the device for use in a cluster resource:

- ♦ **NSS Pool:** For new devices, you must initialize and share the device before creating the pool. For an existing pool that you want to cluster enable, use NSSMU or iManager to share the device.
- ♦ **Linux LVM volume group:** For new devices, use an unpartitioned device that has been initialized. Do not mark the device as shared because doing so creates a small partition on it. LVM uses the entire device for the volume group. For an existing volume group, do not mark the device as shared.

If this is a new cluster, connect the shared disk system to the first server so that the SBD cluster partition can be created during the Novell Cluster Services install. For information, see [Section 4.8.2, “SBD Partitions,” on page 54](#).

4.8.2 SBD Partitions

If your cluster shares storage resources, you must create an SBD (split-brain detector) partition for the cluster. An SBD must be created before you attempt to create file system cluster resources, and before you configure a second node in the cluster.

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

You can create an SBD partition in YaST as part of the first node setup. You can also create it by using the SBD Utility (`sbdutil`) before you configure a second node for the cluster and before you create file system cluster resources. Either method supports mirroring the SBD partition.

You typically carve out small LUNs in your storage array to use exclusively for the SBD partition and its mirror. The device should have at least 20 MB of free available space. The minimum amount of free available space needed is 8 MB. If you want to mirror the SBD partition to achieve greater fault tolerance, you need the same amount of free disk space on the second shared disk. When you size the LUNs, ensure that you consider 4 MB of additional space that will be needed to store its shareable state information.

After you carve the device for the SBD partition (and its mirror, if you use one), you must initialize it and mark it as *Shareable for Clustering*. You can use NSSMU, the Storage plug-in for iManager, or an NSS utility called `ncsinit` to initialize a device and set it to a shared state. The share information is added to the disk in a small partition that 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, 4 MB are used for the share information, leaving up to 1020 MB as free available space for the SBD partition.

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

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

For information about how SBD partitions work and how to create one after installing the first node, see [Section 8.15.3, “Creating a Non-Mirrored Cluster SBD Partition,” on page 115](#).

4.8.3 Shared iSCSI Devices

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

4.8.4 Shared RAID Devices

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

Consider the following when using software RAIDs:

- ♦ NSS software RAID is supported for shared disks.
- ♦ 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 Novell Cluster Services does not support concurrent activation.

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

4.9 SAN Rules for LUN Masking

When you create a 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.

4.10 Multipath I/O Configuration Requirements

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

- [Section 4.10.1, “Path Failover Settings for Device Mapper Multipath,” on page 56](#)
- [Section 4.10.2, “Modifying the Port Down Retry Setting in the modprobe.conf.local File,” on page 56](#)
- [Section 4.10.3, “Modifying the Polling Interval, No Path Retry, and Failback Settings in the multipath.conf File,” on page 57](#)
- [Section 4.10.4, “Modifying the Port Down Retry and Link Down Retry Settings for an HBA BIOS,” on page 58](#)

4.10.1 Path Failover Settings for Device Mapper Multipath

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

The default setting in DM-MP is to queue I/O if one or more HBA paths is lost. 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 automatically fail over storage resources if a disk paths go down.

Change the Retry setting in the `/etc/modprobe.conf.local` and `/etc/multipath.conf` files so that Novell Cluster Services works correctly with DM-MP. For information, see [Section 4.10.2, “Modifying the Port Down Retry Setting in the modprobe.conf.local File,” on page 56](#) and [Section 4.10.3, “Modifying the Polling Interval, No Path Retry, and Failback Settings in the multipath.conf File,” on page 57](#).

Also consider changes as needed for the retry settings in the HBA BIOS. For information, see [Section 4.10.4, “Modifying the Port Down Retry and Link Down Retry Settings for an HBA BIOS,” on page 58](#).

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

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

For example, for QLogic HBAs, ensure that you have installed the latest `qla-driver`. The following settings are suggested for the `/etc/modprobe.conf.local` file. Ensure that you verify the vendor’s preferred settings before making the changes.

```
options qla2xxx qlport_down_retry=2
```

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

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

- ♦ [“Polling Interval” on page 57](#)
- ♦ [“No Path Retry” on page 57](#)
- ♦ [“Failback” on page 57](#)
- ♦ [“Example of Multipath I/O Settings” on page 58](#)

Polling Interval

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

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

No Path Retry

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

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

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

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

Failback

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

```
failback "manual"
```

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

Example of Multipath I/O Settings

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

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

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

For information about configuring the `multipath.conf` file, see [“Managing Multipath I/O for Devices”](http://www.suse.com/documentation/sles11/stor_admin/data/multipathing.html) (http://www.suse.com/documentation/sles11/stor_admin/data/multipathing.html) in the *SLES 11 Storage Administration Guide* (http://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html).

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

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

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

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

5 Installing and Configuring Novell Cluster Services on OES 11

This section describes how to install the Novell Cluster Services software on Novell Open Enterprise Server 11 servers, how to configure the cluster on the first node, and how to configure other nodes for the cluster.

IMPORTANT: Before you install or configure Novell Cluster Services, ensure that you understand the requirements for it and have configured the environment as described in [Chapter 4, “Planning for Novell Cluster Services,”](#) on page 37.

See the following resources for information about rolling cluster upgrades or conversions to the latest version of Novell Cluster Services for OES 11:

Upgrading from	For information, see:
OES 2	Chapter 6, “Upgrading Clusters from OES 2 to OES 11,” on page 85
NetWare 6.5 SP8	OES 11: Novell Cluster Services NetWare to Linux Conversion Guide

- ♦ [Section 5.1, “Novell Cluster Services Licensing,”](#) on page 59
- ♦ [Section 5.2, “Extending the eDirectory Schema to Add Cluster Objects,”](#) on page 60
- ♦ [Section 5.3, “Assigning Install Rights for Container Administrators \(or Non-Administrator Users\),”](#) on page 62
- ♦ [Section 5.4, “Installing Novell Cluster Services,”](#) on page 63
- ♦ [Section 5.5, “Configuring Novell Cluster Services,”](#) on page 66
- ♦ [Section 5.6, “Configuring Additional Administrators,”](#) on page 82
- ♦ [Section 5.7, “Installing or Updating the Clusters Plug-in for iManager,”](#) on page 82
- ♦ [Section 5.8, “Patching Novell Cluster Services,”](#) on page 83
- ♦ [Section 5.9, “What’s Next,”](#) on page 83

5.1 Novell Cluster Services Licensing

Novell Cluster Services supports up to 32 nodes in a single cluster. OES 11 customers receive a Novell Cluster Services entitlement that covers an unlimited number of two-node clusters. Customers who want to add nodes to a two-node 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>).

5.2 Extending the eDirectory Schema to Add Cluster Objects

The first time that you install Novell Cluster Services in a tree, the Novell 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

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

IMPORTANT: For information about the eDirectory rights needed to install Novell Cluster Services in a tree after the schema has been extended, see [Section 5.3, “Assigning Install Rights for Container Administrators \(or Non-Administrator Users\),”](#) on page 62.

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

- ♦ [Section 5.2.1, “Prerequisites for Extending the Schema,”](#) on page 60
- ♦ [Section 5.2.2, “Extending the Schema,”](#) on page 61

5.2.1 Prerequisites for Extending the Schema

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

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

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

Table 5-1 Tree Information Needed for the Schema Expansion

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

5.2.2 Extending the Schema

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

IMPORTANT: It is not necessary to extend the schema separately from the 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 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: Ensure that you change the values inside the quotation marks to the actual settings for your cluster.

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

- 3 As the `root` user, enter the following command at a 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 that you created in [Step 2](#).

- 5 Delete the configuration file (*configuration_filename*) that you created.
This file contains a password in clear text. Ensure that you delete the file for security reasons.
- 6 Continue with [Section 5.3, “Assigning Install Rights for Container Administrators \(or Non-Administrator Users\),”](#) on page 62.

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

If the eDirectory schema has been extended in the tree where you want to create clusters, the container administrator (or non-administrator user) needs the following eDirectory rights to install and configure Novell Cluster Services. These rights are also required if a different user configures Novell Cluster Services after the install as described in [Section 5.5.3, “Using Different LDAP Credentials for the Cluster Configuration,”](#) on page 70.

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

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

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

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

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

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

- ♦ Object Create rights where the cluster container will be.

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

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

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

For information about eDirectory rights, see [“eDirectory Rights” \(http://www.novell.com/documentation/edir88/edir88/data/fbachifb.html\)](http://www.novell.com/documentation/edir88/edir88/data/fbachifb.html) in the *eDirectory 8.8 Administration Guide*.

5.4 Installing Novell Cluster Services

Novell Cluster Services for Linux is included on the add-on media for OES 11. It is necessary to install OES 11 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 5.1, “Novell Cluster Services Licensing,” on page 59](#).

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 install Novell Cluster Services when you install OES 11, or later on an existing OES 11 server.

- ♦ [Section 5.4.1, “Before You Install Novell Cluster Services,” on page 63](#)
- ♦ [Section 5.4.2, “Installing Novell Cluster Services during a OES 11 Installation,” on page 64](#)
- ♦ [Section 5.4.3, “Installing Novell Cluster Services on an Existing OES 11 Server,” on page 65](#)

5.4.1 Before You Install Novell Cluster Services

Ensure that your system meets the requirements and guidelines in [Chapter 4, “Planning for Novell Cluster Services,” on page 37](#). Before you install Novell Cluster Services, verify that your setup meets these requirements:

- ♦ [“Using a Local eDirectory Database” on page 63](#)
- ♦ [“Rights to Extend the Schema” on page 63](#)
- ♦ [“Shared Devices for the SBD” on page 63](#)
- ♦ [“SBD Required before You Install the Second Node in the Cluster” on page 64](#)

Using a Local eDirectory Database

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

Rights to Extend the Schema

If the eDirectory schema was not previously extended in the tree as described in [Section 5.2, “Extending the eDirectory Schema to Add Cluster Objects,” on page 60](#), the administrator user that installs Novell Cluster Services must have the rights to extend the schema, such as the tree administrator.

Shared Devices for the SBD

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

SBD Required before You Install the Second Node in the Cluster

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

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

Typically, you create the SBD when you configure the cluster on the first node as described in [Section 5.5.5, “Configuring a New Cluster,” on page 71](#). A one-node cluster that has shared disk storage can be configured without an SBD, but the SBD must be created before you add another node. For information about manually creating an SBD, see:

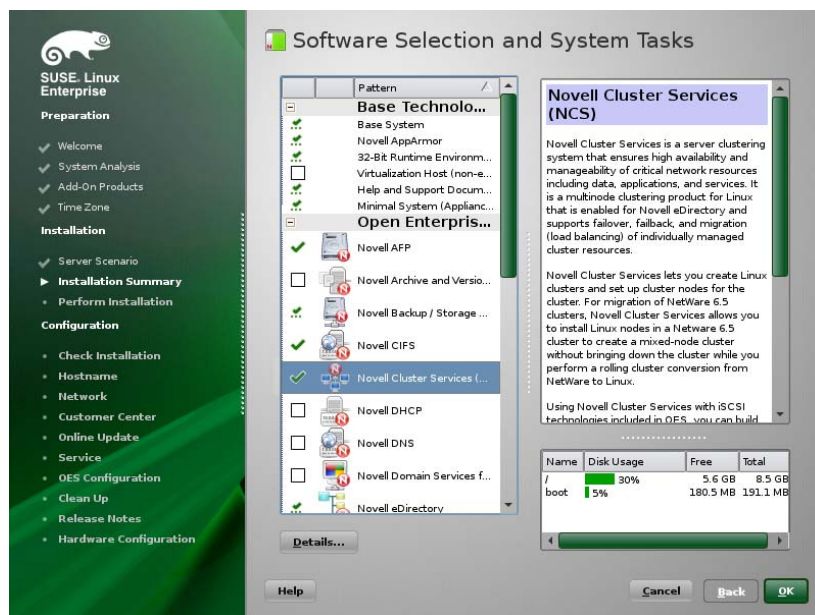
- ♦ [Section 8.15.3, “Creating a Non-Mirrored Cluster SBD Partition,” on page 115](#)
- ♦ [Section 8.15.4, “Creating a Mirrored Cluster SBD Partition,” on page 117](#)

5.4.2 Installing Novell Cluster Services during a OES 11 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 11, see the [OES 11: Installation Guide](#).

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

- 1 Start the YaST install for SUSE Linux Enterprise Server and continue to the Installation Mode page.
- 2 Select *New Installation*, select *Include Add-On Products from Separate Media*, click *Next*, then continue through the OES add-on part of the install until you get to the Installation Settings page.
- 3 On the Installation Settings page, click *Software* to open the Software Selection and System Tasks page.
- 4 Under *Open Enterprise Server*, 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 are automatically selected:

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

The following Open Enterprise Server 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 the same tree.
- ♦ 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 other protocols and services according to your planned setup. For information, see [Section 4.7, “Software Requirements for Cluster Resources,” on page 50](#).

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.

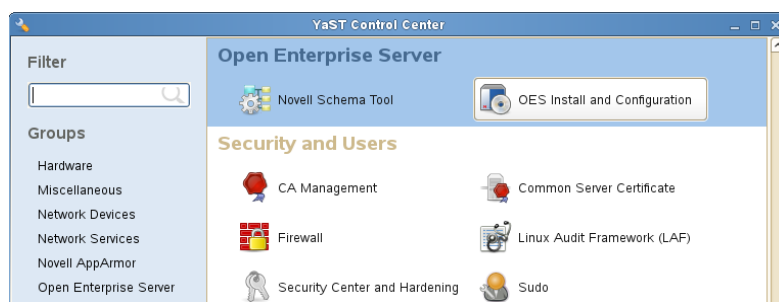
If you click *Accept*, then return to software pattern selection page, the selections that you made become your base selections and must be deselected if you want to remove them from the installation proposal.

- 5 Continue through the installation process, but do not configure Novell Cluster Services when you reach the Novell Open Enterprise Server Configuration page. You configure it later.
- 6 After the install, use the Software Updater (or other update methods) to download and apply any updates that are in the SLES and OES Updates channels.
- 7 Ensure that your SAN storage is working properly for the server.
- 8 Continue with [Section 5.5, “Configuring Novell Cluster Services,” on page 66](#).

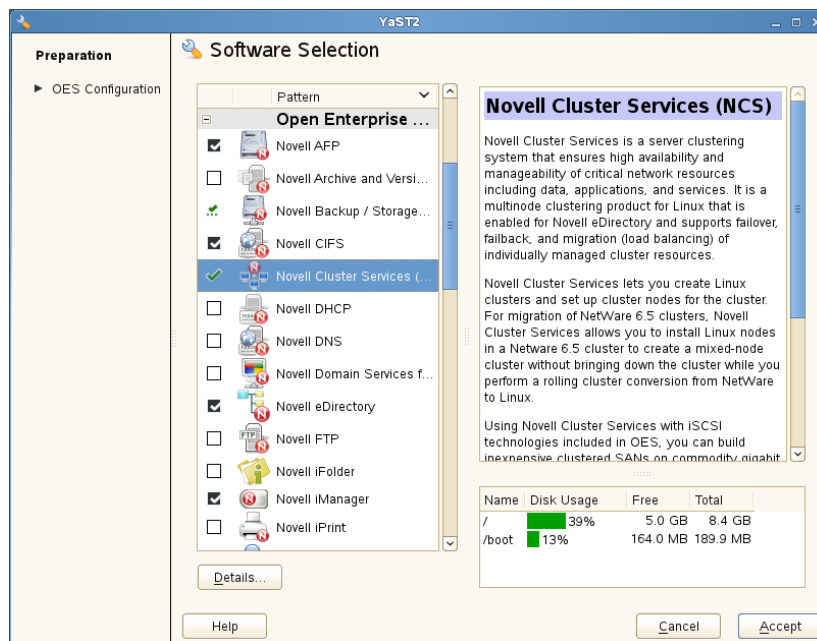
5.4.3 Installing Novell Cluster Services on an Existing OES 11 Server

If you did not install Novell Cluster Services during the OES 11 installation, you can install it later by using *YaST > Open Enterprise Server > OES Install and Configuration*.

- 1 Log in to the server as the `root` user.
- 2 In YaST, select *Open Enterprise Server > OES Install and Configuration*.



- 3 On the Software Selection page under *Open Enterprise Server*, select *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 4 in Section 5.4.2, “Installing Novell Cluster Services during a OES 11 Installation,” on page 64.](#)

- 4 Click *Accept* to begin the install, then click *Continue* to accept changed packages.
- 5 Continue through the installation process, but do not configure Novell Cluster Services when you reach the Novell Open Enterprise Server Configuration page. You configure it later.
- 6 After the install, use the Software Updater (or other update methods) to download and apply any updates that are in the SLES and OES Updates channels.
- 7 Ensure that your SAN storage is working properly for the server.
- 8 Continue with [Section 5.5, “Configuring Novell Cluster Services,” on page 66.](#)

5.5 Configuring Novell Cluster Services

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

IMPORTANT: The YaST-based configuration is not used to modify the settings for an existing cluster. For information about modifying the settings for an existing cluster, see [Section 8.11, “Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster,” on page 108.](#)

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.

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

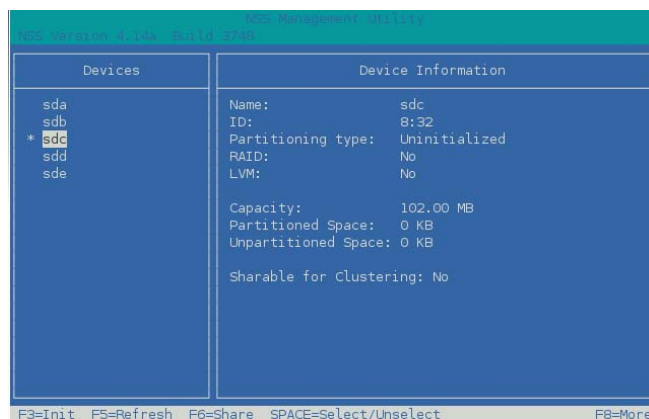
- ♦ [Section 5.5.1, “Initializing and Sharing a Device to Use for the SBD Partition,” on page 67](#)
- ♦ [Section 5.5.2, “Opening the Novell Open Enterprise Server Configuration Page,” on page 69](#)
- ♦ [Section 5.5.3, “Using Different LDAP Credentials for the Cluster Configuration,” on page 70](#)
- ♦ [Section 5.5.4, “Accessing the Novell Cluster Services Configuration Page in YaST,” on page 70](#)
- ♦ [Section 5.5.5, “Configuring a New Cluster,” on page 71](#)
- ♦ [Section 5.5.6, “Adding a Node to an Existing Cluster,” on page 77](#)

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

If you plan to use shared storage in the cluster, the cluster needs a split-brain detector (SBD). You must create the SBD before you add a second node to the cluster. You can create the SBD when you set up the cluster on the first node. In preparation, you must initialize the partition, and mark its device as shareable for clustering. If you plan to mirror the SBD, you must also initialize the partition that you want to use as the mirrored SBD, and mark its device as shareable for clustering. You need at least 20 MB for the SBD.

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

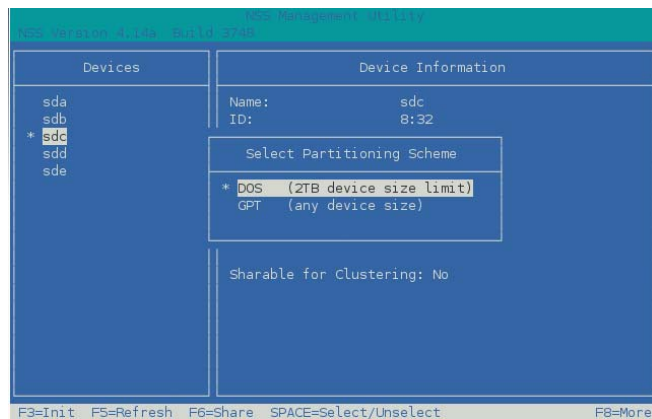
nssmu
- 3 In the NSSMU, select *Devices* and press Enter.
- 4 In the *Devices* list, select the device that you want to use for the SBD.



- 5 If the device has not been initialized or if you want to delete the current partitioning structures on the device, press F3, then press Y (Yes) to confirm and continue.

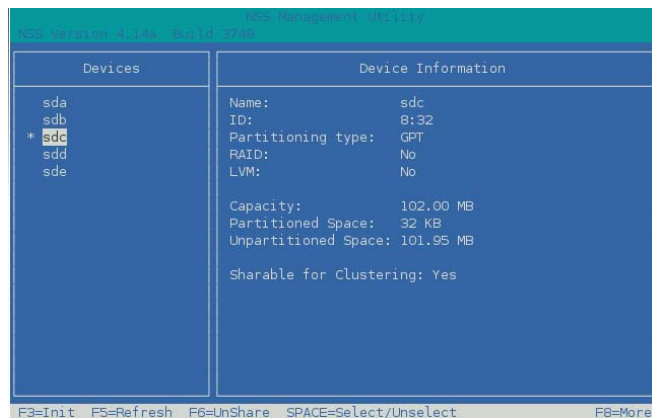
WARNING: Initializing a disk destroys all of the data on it.

- 6 Select the DOS or GPT partitioning scheme for the device, then press Enter.
DOS supports devices up to 2 TB in size. GPT supports devices of any size.



Wait for the page to refresh before continuing.

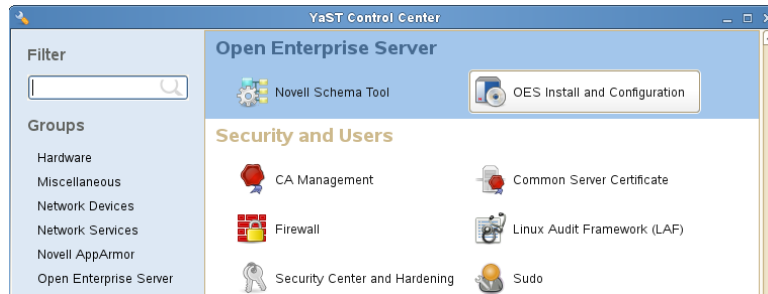
- 7 Press F6 to mark the device as shareable for clustering.
The *Shareable for Clustering* value changes from *No* to *Yes*.



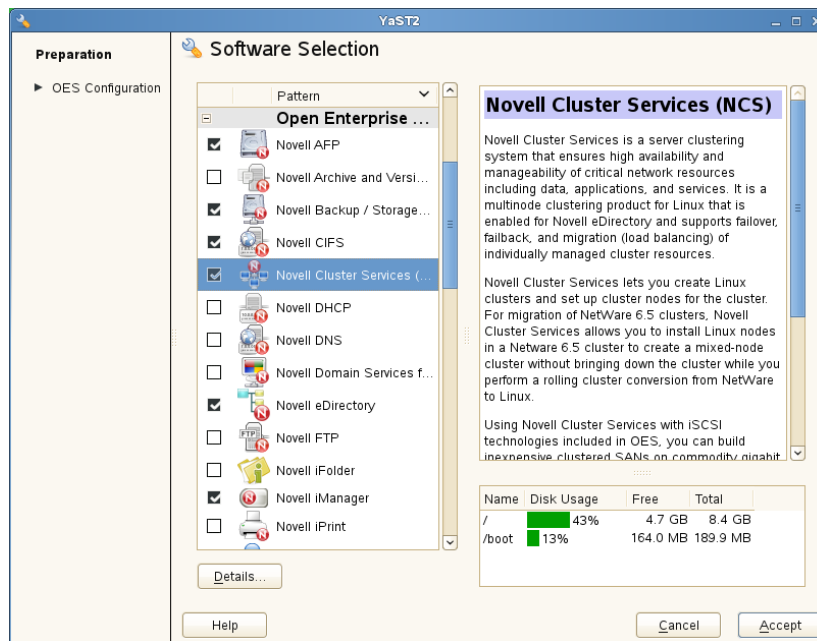
- 8 If you plan to mirror the SBD, repeat [Step 4](#) through [Step 7](#) for the second device.
- 9 Exit NMSU.
- 10 Continue with [Section 5.5.2, “Opening the Novell Open Enterprise Server Configuration Page,”](#) on page 69.

5.5.2 Opening the Novell Open Enterprise Server Configuration Page

- 1 Log in to the server as the `root` user.
- 2 Verify that the OES Add-on CD is mounted on or available to the server.
- 3 In YaST, select *Open Enterprise Server* > *OES Install and Configuration*.



- 4 On the Software Selection page under *Open Enterprise Server*, verify that the *Novell Cluster Services* option is already installed as indicated by a blue check mark.



- 5 Click *Accept* to proceed to the Novell Open Enterprise Server Configuration page.
- 6 Do one of the following:
 - ♦ **Same Administrator:** To use the same administrator credentials that were used to install Novell Cluster Services, continue with [Section 5.5.4, “Accessing the Novell Cluster Services Configuration Page in YaST,” on page 70.](#)
 - ♦ **Different Administrator:** To use different administrator credentials than those used to install Novell Cluster Services, continue with [Section 5.5.3, “Using Different LDAP Credentials for the Cluster Configuration,” on page 70.](#)

5.5.3 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 Open Enterprise Server on the server by reconfiguring the settings for the *LDAP Configuration of Open Enterprise Services* option.

For information about what rights are needed, see [Section 5.3, “Assigning Install Rights for Container Administrators \(or Non-Administrator Users\),”](#) on page 62.

- 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 user name 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 5.5.4, “Accessing the Novell Cluster Services Configuration Page in YaST,”](#) on page 70.

5.5.4 Accessing 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.

[Novell Cluster Services \(NCS\)](#)

Configure is [disabled](#)

- 2 Click the *disabled* link to enable configuration.

The sentence changes to *Configure is enabled*.

[Novell Cluster Services \(NCS\)](#)

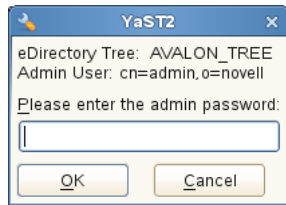
Novell Cluster Services (NCS) requires additional configuration information before continuing or disable the configuration.

Configure is [enabled](#)

- New or Existing Cluster: Existing Cluster
- LDAP Servers: 10.10.10.37
- LDAP Username: cn=admin,o=novell
- Cluster object in the directory: cn=cluster,o=novell
- NCS Proxy User: cn=OESCommonProxy_avalon,o=novell
- Node name: avalon
- Node's IP address: 10.10.10.37
- Start Cluster Services: Now

Default settings, not yet configured.

- 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**.



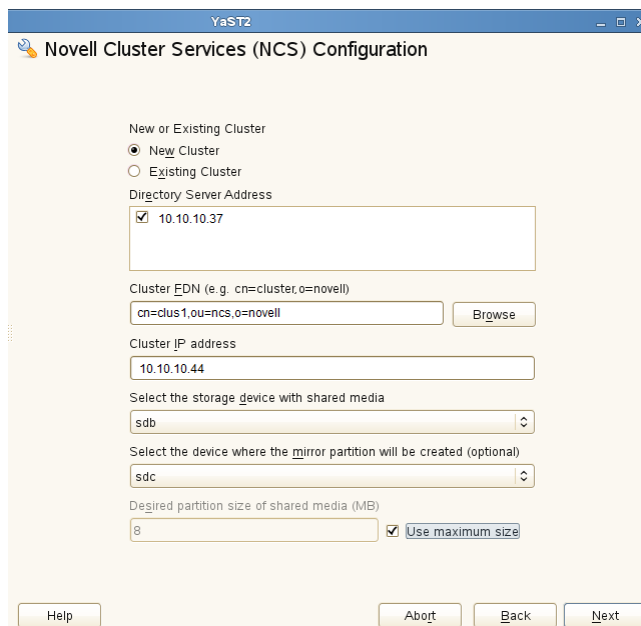
This is either the LDAP administrator identity that was used when the Open Enterprise Server were installed, or the identity configured after the install in [Section 5.5.3, “Using Different LDAP Credentials for the Cluster Configuration,”](#) on page 70.

- 5 On the Novell Cluster Services Configuration page, continue with one of the following:
 - ♦ [Section 5.5.5, “Configuring a New Cluster,”](#) on page 71
 - ♦ [Section 5.5.6, “Adding a Node to an Existing Cluster,”](#) on page 77

5.5.5 Configuring a New Cluster

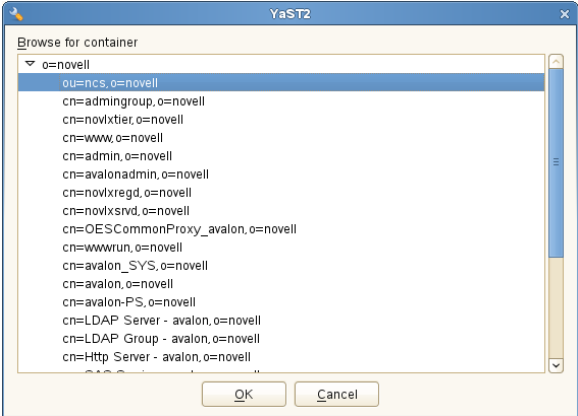
Perform the following configuration for the first node that you configure for a cluster:


- 1 Go to the Novell Cluster Services Configuration page as described in [Section 5.5.4, “Accessing the Novell Cluster Services Configuration Page in YaST,”](#) on page 70.



- 2 On the first configuration page, specify the following settings for a new cluster, then click *Next*:

Parameter	Action
New or existing cluster	Select <i>New Cluster</i> .

Parameter	Action
Directory Server Address	<p>Select the check box next to the IP address of the LDAP server you want to use as the default for this node. The local LDAP server is selected by default.</p> <p>The IP addresses shown are the LDAP servers available for this service to use. The LDAP servers must have a master replica or a Read/Write replica of eDirectory.</p> <p>You can add, remove, or change the order of available LDAP servers for the node after the setup is complete. See Section 8.11.1, "Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster," on page 108.</p>
Cluster FDN	<p>Click <i>Browse</i> and navigate the tree to select the container where you want to create the cluster, then click <i>OK</i>. The FDN is automatically added to the field with a suggested cluster name. You can specify a different cluster name.</p>  <p>You can also specify the typeful FDN (fully distinguished name) of the existing cluster. The name is case sensitive. Use the comma format illustrated in the example. Do not use dots.</p> <p>For example:</p> <pre>cn=clus1,ou=ncs,o=mycompany</pre> <p>You must specify an existing context. Specifying a new context does not create a new context.</p> <p>Cluster names must be unique. You cannot create two clusters with the same name in the same eDirectory tree.</p>

Parameter	Action
Cluster IP Address	<p>The cluster IP address represents the cluster on the network. It is different than the server IP address. For example:</p> <p>10.10.10.44</p> <p>The cluster IP address provides a single point for cluster access, configuration, and management. A Master IP Address resource is created automatically during the Cluster Services installation. The cluster IP address is bound to the master node and remains with the master node regardless of which server is the master node.</p> <p>The cluster IP address is required to be on the same IP subnet as the other servers in the same cluster, which is necessary for certain external network management programs to get cluster status alerts.</p>
Select the storage device with shared media	<p>Select the shared device that you want to use for the SBD from the drop-down list.</p> <p>For a new cluster, you can create the SBD partition now, or you can create it manually at any time before you add a second node to the cluster.</p> <p>A split-brain detector (SBD) is required if you plan to use shared disks in the cluster. For information about SBD requirements, see “SBD Partitions” on page 54.</p> <p>The drop-down menus show only devices that have been initialized and shared. See Section 5.5.1, “Initializing and Sharing a Device to Use for the SBD Partition,” on page 67.</p> 
Select the device where the mirror partition will be created	<p>If you want to mirror the SBD, select a different shared device from the drop-down list.</p>
Desired partition size of shared media	<p>Specify a size of 20 MB or greater to use for the SBD partition, or select <i>Use maximum size</i> to use the entire device.</p> <p>If you selected a device for the SBD mirror, the specified size is also used for the mirror segment.</p>

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

YaST2

Novell Cluster Services (NCS) Proxy User Configuration

NCS Proxy User Name (e.g. cn=admin,o=novell)

cn=OESCommonProxy_avalon,o=novell

Specify NCS Proxy User Password

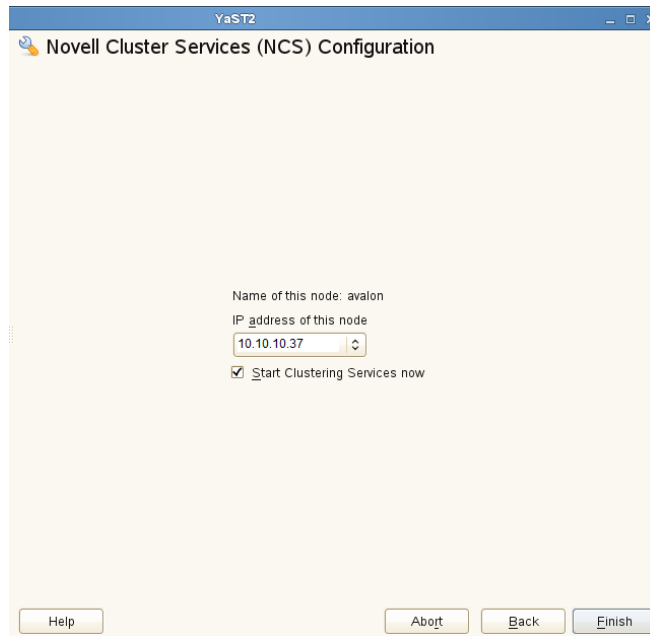
☒ Use OES Common Proxy User

Help Abort Back Next

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

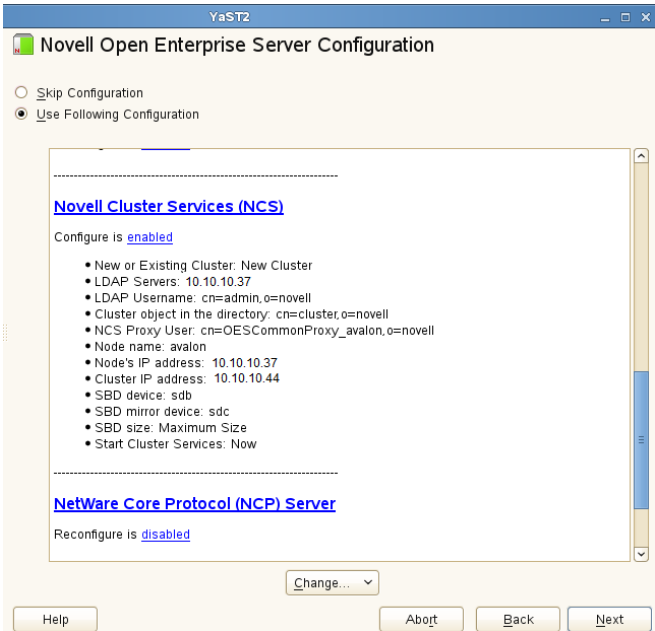
You can reset the default settings by clicking *Back* to return to the Novell Cluster Services Configuration page, then clicking *Next* to continue again to the Proxy User Configuration page.

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



Parameter	Action
IP address of this node	From the drop-down list, select the IP address that 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.
Start Clustering Services now	Select the check box to start Novell Cluster Services software on this node after configuring it.

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



- 6 Continue through the setup process, then click *Finish* to exit the OES Configuration
- 7 Start Novell Cluster Services using one of these methods:

Setup Condition	Instructions
<i>Start Cluster Services now was enabled</i>	Novell Cluster Services starts automatically after the configuration completes.
<i>Start Cluster Services now was disabled</i>	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: <code>./novell-ncs start</code>◆ At a terminal console prompt, enter the following as the <code>root</code> user: <code>rcnovell-ncs start</code>

- 8 Continue with [Section 5.6, “Configuring Additional Administrators,”](#) on page 82.

5.5.6 Adding a Node to an Existing Cluster

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

- 1 If you have not previously configured the SBD partition for the cluster, create an SBD partition for the cluster by using one of the following procedures:
 - ♦ [Section 8.15.3, “Creating a Non-Mirrored Cluster SBD Partition,” on page 115](#)
 - ♦ [Section 8.15.4, “Creating a Mirrored Cluster SBD Partition,” on page 117](#)

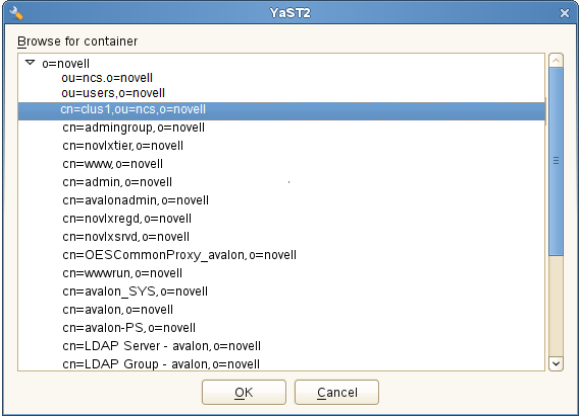
If you have a shared disk system attached to your cluster servers, an SBD partition is required and must be created before you configure the second node in the cluster.

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 as described in [Section 5.5.4, “Accessing the Novell Cluster Services Configuration Page in YaST,” on page 70](#).

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

Parameter	Action
New or existing cluster	Select <i>Existing Cluster</i> .

Parameter	Action
Directory Server Address	<p>Select the check box next to the IP address of the LDAP server you want to use as the default for this node. The local LDAP server is selected by default.</p> <p>The IP addresses shown are the LDAP servers available for this service to use. The LDAP servers must have a master replica or a Read/Write replica of eDirectory.</p> <p>You can add, remove, or change the order of available LDAP servers for the node after the setup is complete. See Section 8.11.1, “Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster,” on page 108.</p>
Cluster FDN	<p>Click <i>Browse</i> and navigate the tree to select the Cluster object of the existing cluster, then click <i>OK</i>. The FDN is automatically added to the field.</p>  <p>You can also specify the typeful FDN (fully distinguished name) of the existing cluster. The name is case sensitive. Use the comma format illustrated in the example. Do not use dots.</p> <p>For example:</p> <pre>cn=clus1,ou=ncs,o=mycompany</pre>
Select the device where the mirror partition will be created	If you want to mirror the SBD, select a different shared device from the drop-down list.

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

YaST2

Novell Cluster Services (NCS) Proxy User Configuration

NCS Proxy User Name (e.g. cn=admin,o=novell)

cn=OESCommonProxy_avalon,o=novell

Specify NCS Proxy User Password

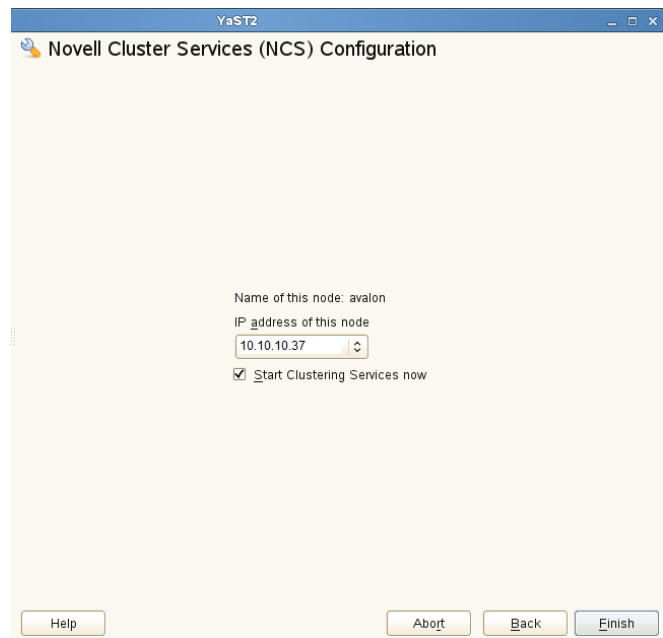
☒ Use OES Common Proxy User

Help Abort Back Next

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

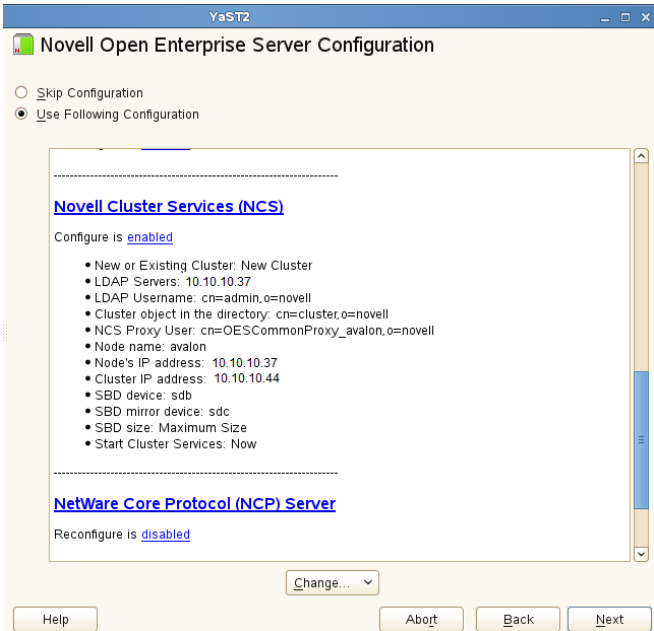
You can reset the default settings by clicking *Back* to return to the Novell Cluster Services Configuration page, then clicking *Next* to continue again to the Proxy User Configuration page.

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



Parameter	Action
IP address of this node	From the drop-down list, select the IP address that 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.
Start Clustering Services now	Select the check box to start Novell Cluster Services software on this node after configuring it.

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



- 7 Continue through the setup process, then click *Finish* to exit the OES Configuration
- 8 Start Novell Cluster Services using one of these methods:

Setup Condition	Instructions
<i>Start Cluster Services now was enabled</i>	Novell Cluster Services starts automatically after the configuration completes.
<i>Start Cluster Services now was disabled</i>	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>

- 9 Continue with [Section 5.6, “Configuring Additional Administrators,”](#) on page 82.

5.6 Configuring Additional Administrators

The Administrator user that you specify as the Novell Proxy User 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' user names must be similarly configured.

You can modify the default administrator user name or password after the install by following the procedure in [Section 8.11, "Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster,"](#) on page 108.

5.7 Installing or Updating the Clusters Plug-in for iManager

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

- ♦ [Section 5.7.1, "Prerequisites for Installing or Updating the Clusters Plug-In,"](#) on page 82
- ♦ [Section 5.7.2, "Installing the Clusters Plug-In,"](#) on page 82
- ♦ [Section 5.7.3, "Uninstalling and Reinstalling the Clusters Plug-In,"](#) on page 83

5.7.1 Prerequisites for Installing or Updating the Clusters Plug-In

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

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

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

For information, see ["Storage-Related Plug-Ins for iManager"](#) on page 47.

5.7.2 Installing the Clusters Plug-In

If you are installing a new instance of Novell iManager 2.7.4 or later and the Clusters plug-in on a computer, follow the instructions in ["Downloading and Installing Plug-in Modules"](#) in the [Novell iManager 2.7.4 Administration Guide](#). Ensure that you install the Clusters plug-in (`ncsmgmt.rpm`) as well as the common code for storage-related plug-ins (`storagemgmt.rpm`).

To configure role based services for the Clusters plug-in, see ["RBS Configuration"](#) in the [Novell iManager 2.7.4 Administration Guide](#).

5.7.3 Uninstalling and Reinstalling the Clusters Plug-In

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

- 1 In iManager, uninstall the currently installed storage-related plug-ins.
- 2 Copy the new version of the plug-in files to the iManager plug-ins location, manually overwriting the older version of the plug-in in the `packages` folder with the newer version of the plug-in.
- 3 In iManager, install all of the storage-related plug-ins, or install the plug-ins you need, plus the common code.
- 4 If you are updating the Clusters plug-in, delete the contents of the `/var/opt/novell/tomcat6/work/Catalina/localhost/nps` work folder.
You must manually delete the old compiled JSPs that are cached in the `nps` work folder. This ensures that only the latest files are cached.
- 5 Restart Tomcat by entering the following command, or reboot the server if a reboot is required by any of the patches applied.

```
rcnovell-tomcat6 restart
```

5.8 Patching Novell Cluster Services

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

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

```
rcnovell-ncs stop
```
- 4 Apply the OES 11 or Novell Cluster Services patches.
- 5 Start Novell Cluster Services by entering

```
rcnovell-ncs start
```
- 6 Use the `cluster join` command to rejoin the node to the cluster.

5.9 What's Next

After installing Novell Cluster Services on OES 11 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 91](#)
- ♦ [Chapter 8, “Managing Clusters,” on page 99.](#)
- ♦ [Chapter 9, “Configuring and Managing Cluster Resources,” on page 127](#)
- ♦ [Chapter 11, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 155](#)

- ♦ [Chapter 12, “Configuring and Managing Cluster Resources for LVM Volume Groups,”](#) on page 195
- ♦ [Chapter 13, “Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers,”](#) on page 229

If you install Novell Cluster Services at the host level of an OES 11 (Xen) server, you can create cluster resources for the virtual machines. For information, see [Section 14.2, “Virtual Machines as Cluster Resources,”](#) on page 250.

6 Upgrading Clusters from OES 2 to OES 11

This section describes how to upgrade a Novell Cluster Services cluster and the file system resources from Novell Open Enterprise Server (OES) 2 SP3 to OES 11. See the individual service guides for information about converting services to OES 11.

- ♦ [Section 6.1, “What’s New for OES 11,” on page 85](#)
- ♦ [Section 6.2, “Supported Upgrade Paths from OES 2 to OES 11,” on page 87](#)
- ♦ [Section 6.3, “Requirements and Guidelines for Upgrading Clusters from OES 2,” on page 87](#)
- ♦ [Section 6.4, “Adding OES 11 Nodes to an OES 2 Cluster \(Rolling Cluster Upgrade\),” on page 88](#)

6.1 What’s New for OES 11

For OES 11, Novell Cluster Services supports the following file system resources:

- ♦ [Section 6.1.1, “NSS Pools,” on page 85](#)
- ♦ [Section 6.1.2, “Linux POSIX File Systems on an EVMS CSM \(Compatibility Only\),” on page 86](#)
- ♦ [Section 6.1.3, “Linux POSIX File Systems on an LVM Volume Group,” on page 86](#)

6.1.1 NSS Pools

Novell Cluster Services for OES 11 supports NSS pools that are created on OES 11, OES 2 SP3 and earlier, and NetWare 6.5 SP8. The NSS and Cluster Services management tools can recognize NSS pools and volumes no matter where they were created. In a mixed-mode cluster, existing NSS pool cluster resources can be cluster migrated to any node in the cluster.

IMPORTANT: Do not create, modify, or delete NSS pool cluster resources on OES 11 nodes while the cluster is in mixed-mode. Any such action can result in data corruption. For information, see [Section 6.3, “Requirements and Guidelines for Upgrading Clusters from OES 2,” on page 87.](#)

Wait until after all nodes in the cluster are upgraded to OES 11 to create new NSS pools on the OES 11 nodes. For OES 11 servers, the NSS management tools can initialize a device to use the DOS partition table scheme or GUID partition table (GPT) scheme. DOS supports device sizes up to 2 TB. GPT supports device sizes up to 8 TB, which allows you to create NSS pools up to 8 TB (terabytes) in size on a single device.

6.1.2 Linux POSIX File Systems on an EVMS CSM (Compatibility Only)

Novell Cluster Services for OES 11 provides compatibility support for Linux POSIX file systems that were created by using the Cluster Segment Manager (CSM) in EVMS (Enterprise Volume Management System).

For OES 2, Novell Cluster Services provided a Cluster Segment Manager (CSM) for EVMS to support cluster resources for Linux POSIX file systems. EVMS was deprecated in the SUSE Linux Enterprise Server (SLES) 11 release, and is no longer available. Accordingly, Novell Cluster Services provides compatibility support for these cluster resources. It has been modified in the following ways to allow you to move clustered Linux POSIX file systems from OES 2 to OES 11:

- ♦ Recognizes and supports the EVMS-based file system structure.
- ♦ Provides a way for managing a resource in the load, unload, and monitor scripts that works with the EVMS device and segment structure.
- ♦ Changes the device path from `/dev/evms/` to `/dev/mapper/`, which is the same location where LVM devices are stored.

Before any of the new OES11 nodes join an OES2 cluster, you must offline all resources that use Linux POSIX volumes. Modify their load, unload, and monitor scripts to use the new commands for compatibility mode. Allow some OES 11 nodes to join the OES 2 cluster, then modify each resources' preferred nodes list to use only the OES 11 nodes in the cluster. Ensure that you disable *Resource Follows Master* setting for the resource until the conversion is finalized. After a resource's scripts and preferred nodes list have been modified, you can bring the resource online. As you add other OES 11 nodes, you can add those to the preferred nodes lists for failover.

After a resource has been moved from an OES 2 node to an OES 11 node, it is not supported to move it back to an OES 2 node. The resource should fail over only to other OES 11 nodes in a mixed cluster.

In the NSS management tools, the CSM layer is reported in the Details page for the shared device. This allows you to recognize the difference between the Linux POSIX resources that were created with EVMS and the new Linux LVM volume group resources that are created on OES 11.

6.1.3 Linux POSIX File Systems on an LVM Volume Group

Novell Cluster Services for OES 11 supports Linux POSIX file systems on Linux Volume Manager (LVM) volume groups.

You should not create new Linux POSIX cluster resources on OES 11 while the cluster is in mixed mode, except as part of a documented conversion procedure.

The LVM-based cluster resources can be used only on nodes running OES 11 or later. The resource must not be moved to a node running earlier versions of OES or NetWare. Ensure that you set up your preferred nodes for failover so that the LVM volume group cluster resource fails over only to OES 11 nodes in a mixed-mode cluster.

LVM tools are available in the YaST Expert Partitioner and in command line commands. For information using the, see "LVM Configuration" (http://www.suse.com/documentation/sles11/stor_admin/data/lvm.html) in the *SLES 11 SP1: Storage Administration Guide* (http://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html). For information about using LVM commands, see the man pages for the commands described in Table 12-1, "LVM Commands," on page 227.

6.2 Supported Upgrade Paths from OES 2 to OES 11

The following upgrade paths are supported from OES 2 to OES 11. Ensure that you update the OES 2 SP3 cluster with the latest patches on each node before you begin the upgrade to OES 11.

From this OES 2 platform	To this OES 11 platform
OES 2 SP3 on SUSE Linux Enterprise Server 10 SP3	OES 11 on SUSE Linux Enterprise Server 11 SP1
OES 2 SP3 on SUSE Linux Enterprise Server 10 SP4	OES 11 on SUSE Linux Enterprise Server 11 SP1

6.3 Requirements and Guidelines for Upgrading Clusters from OES 2

In addition to the [Chapter 4, “Planning for Novell Cluster Services,” on page 37](#), consider the following rules and recommendations for mixed OES 2 and OES 11 clusters:

- Mixed OES 2 and OES 11 clusters should be considered a temporary configuration that exists only during an upgrade.
- Adding a new OES 2 cluster node to a mixed-mode OES 2 and OES 11 cluster is not supported.
- Linux POSIX cluster resources on OES 2 nodes must be taken offline before you allow any of the OES 11 nodes to join the cluster.

Before bringing a resource online again, ensure that you modify its scripts for compatibility mode and configure its preferred nodes list to ensure that it is loaded on and fails over to only OES 11 nodes. You should also disable the *Resource Follows Master* settings until after the cluster conversion is finalized. For instructions, see [Chapter 13, “Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers,” on page 229](#).

- After OES 11 nodes are added to the OES 2 cluster, the Linux POSIX cluster resources must be mounted on and fail over to only OES 11 nodes.
- NSS pool cluster resources can fail over to any node in the mixed-mode OES 2 and OES 11 cluster.
- No storage management functions should be executed while a cluster is in mixed-mode unless you must do so as part of the documented conversion process. Do not attempt to create, delete, expand, or modify the properties for partitions, pools, or volumes for any shared resources in the cluster.

IMPORTANT: If it is necessary to create, modify, or delete a pool as part of a documented cluster resource upgrade process, you must shut down the OES 11 nodes, create the NSS pool cluster resource on an OES 2 node, then start the OES 11 nodes. This allows the OES 11 nodes to see the new configuration. This is critical so that pools are able to migrate, and to prevent possible corruption.

- You should not create new LVM volume group cluster resources in a mixed-mode cluster. Wait until after the cluster upgrade is completed.

6.4 Adding OES 11 Nodes to an OES 2 Cluster (Rolling Cluster Upgrade)

Performing a rolling cluster upgrade from OES 2 to OES 11 lets you keep your cluster up and running and lets your users continue to access cluster resources while the upgrade is being performed.

During a rolling cluster upgrade, one OES 11 server is added at a time. You should complete the upgrade as soon as possible. Don't leave the cluster in a mixed-mode state for an extended period.

To perform a rolling cluster upgrade:

- 1 Make a note of the OES components that are installed on the OES 2 cluster nodes.
You will probably want to install the same OES components on the OES 11 nodes that you add to the cluster.
- 2 Prepare some OES 11 nodes for the cluster, but do not allow them to join the cluster.

IMPORTANT: If you are upgrading an OES 2 cluster node to OES 11, leave the cluster before the upgrade and do not rejoin until [Step 4](#).

- 2a Install SUSE Linux Enterprise Server (SLES) 11 and the OES Services from the OES Add-on disk, including Novell Cluster Services, but do not configure the node for the cluster at this time.
- 2b Verify that the server is functioning properly.
- 2c In YaST, go to the Novell Cluster Services Configuration page as described in [Section 5.5.4, "Accessing the Novell Cluster Services Configuration Page in YaST,"](#) on page 70.
- 2d Configure the node for the existing OES 2 cluster as described in [Section 5.5.6, "Adding a Node to an Existing Cluster,"](#) on page 77.

Parameter	Action
LDAP administrator user password	<p>Specify the password of the user name shown, then click OK.</p> <p>Typically, this is the LDAP administrator identity that was used when the OES Services were installed, or the identity configured after the install in Section 5.5.3, "Using Different LDAP Credentials for the Cluster Configuration," on page 70.</p>
New or existing cluster	Select <i>Existing Cluster</i> .
Directory Server Address	<p>Select the check box next to the IP address of the LDAP server you want to use as the default for this node. The local LDAP server is selected by default.</p> <p>You can add, remove, or change the order of available LDAP servers for the node after the setup is complete. See Section 8.11.1, "Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster," on page 108.</p>
Cluster FDN	Click <i>Browse</i> , navigate the tree to select the existing Cluster object for the OES 2 cluster, click OK , then click <i>Next</i> .

Parameter	Action
NCS Proxy User Name	Specify one of the following users as the NCS Proxy user for this server, then click <i>Next</i> : <ul style="list-style-type: none"> ♦ OES Common Proxy User (default if it was configured for the server) ♦ LDAP Admin User ♦ Another Administrator User
IP address of this node	From the drop-down list, select the IP address that Novell Cluster Services will use for this node.
Start Cluster Services now	Deselect the check box. Do not start cluster services on this node or join the cluster until Step 4 , after all of the Linux POSIX cluster resources have been taken offline.

- 2e Click *Finish*, review the summary, then click *Next* to complete the configuration.
- 2f Click *Finish* to close the OES Configuration dialog box.
- 3 On each OES 2 SP3 node, prepare each of the existing Linux POSIX cluster resources for OES 11 by doing the following:
 - 3a Take the Linux POSIX cluster resource offline.
 - 3b Disable the *Resource Follows Master* option for the resource.
 - 3c Modify the resource's load, unload, and monitor scripts.
Use one of the following methods to prepare the resource, depending on how it is configured:
 - ♦ **Cluster Segment Manager:** If the Linux POSIX volume is built directly on the Cluster Segment Manager, follow the procedure in [Section 13.2, "Modifying the Scripts for a CSM Resource,"](#) on page 230. This is the typical configuration in OES 2.
 - ♦ **Cluster Segment Manager plus Another Segment Manager:** If one or multiple Linux POSIX volumes are built on a segment manager on top of the Cluster Segment Manager, follow the procedure in [Section 13.3, "Modifying the Scripts for a CSM Resource with a Segment Manager On It,"](#) on page 234. This is an optional configuration in OES 2 that is rarely used.
 - 3d Modify other resource policy settings as desired, but do not enable the *Resource Follows Master* option.
 - 3e Do not bring a Linux POSIX resources online until you are ready to do so in [Step 5b](#), after its *Preferred Nodes* list has been modified to use only OES 11 nodes.
 - 3f Repeat [Step 3a](#) to [Step 3e](#) for each Linux POSIX cluster resource.
- 4 For each of the OES 11 nodes, start Novell Cluster services manually by using one of the following methods, and allow the node to join the OES 2 cluster.
 - ♦ Reboot the cluster server.
 - ♦ At a terminal console prompt, go to the `/etc/init.d` directory, then enter the following as the root user:

```
./novell-ncs start
```

- ♦ At a terminal console prompt, enter the following as the root user:

```
rcnovell-ncs start
```

- 5 For each of the Linux POSIX cluster resources in turn, do the following:

- 5a In iManager, modify the resource's *Preferred Nodes* list to include only OES 11 nodes as described in [Section 13.4, "Configuring the Preferred Nodes and Cluster Settings for a CSM Cluster Resource,"](#) on page 243.

- 5b In iManager, bring the resource online on its first preferred node. Click *Clusters > Cluster Manager*, select the resource object, then click *Online*.

You can also use the `cluster online` command.

```
cluster online <resource_name> <node_name>
```

- 5c Verify that the resource and any related services are working properly for the resource.

- 5d Repeat [Step 5a](#) to [Step 5c](#) for each Linux POSIX cluster resource.

- 6 If you add more OES 11 nodes, ensure that you modify the *Preferred Nodes* list accordingly for each of the Linux POSIX cluster resources.

- 7 Bring down the OES 2 cluster node that you want to remove from the cluster.

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

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

- 8 After the last cluster resource has been moved to the OES 11 nodes and all OES 2 nodes have left the cluster, the conversion is complete.

7 Configuring Cluster Policies and Priorities

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 91](#)
- ♦ [Section 7.2, “Configuring Quorum Membership and Timeout Properties,” on page 93](#)
- ♦ [Section 7.3, “Configuring Cluster Protocol Properties,” on page 94](#)
- ♦ [Section 7.4, “Configuring Cluster Event Email Notification,” on page 95](#)
- ♦ [Section 7.5, “Viewing the Cluster Node Properties,” on page 97](#)
- ♦ [Section 7.6, “Modifying the Cluster IP Address and Port Properties,” on page 97](#)
- ♦ [Section 7.7, “What’s Next,” on page 97](#)

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 91](#)
- ♦ [Section 7.1.2, “Cluster Protocols,” on page 92](#)

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.

Table 7-1 Cluster Policies

Property	Description
Cluster IP address	<p>Specifies the IP address for the cluster.</p> <p>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. For information, see Section 7.6, “Modifying the Cluster IP Address and Port Properties,” on page 97.</p>

Property	Description
Port	<p>Specifies the port used for cluster communication.</p> <p>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 97.</p>
Quorum membership	<p>Specifies number of nodes that must be up and running in the cluster in order for cluster resources to begin loading.</p> <p>Specify a value between 1 and the number of nodes.</p> <p>For instructions, see Section 7.2, “Configuring Quorum Membership and Timeout Properties,” on page 93.</p>
Quorum timeout	<p>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.</p> <p>For instructions, see Section 7.2, “Configuring Quorum Membership and Timeout Properties,” on page 93.</p>
Email notification	<p>Enables or disables email notification for the cluster. If it is enabled, you can specify up to eight administrator email addresses for cluster events notification.</p> <p>For instructions, see Section 7.4, “Configuring Cluster Event Email Notification,” on page 95.</p>

7.1.2 Cluster Protocols

[Table 7-2](#) describes the configurable cluster protocols that govern inter-node communication transmission and tolerances. You can manage cluster protocols in iManager by going to the *Clusters > Cluster Options > Protocols* page. For instructions, see [Section 7.3, “Configuring Cluster Protocol Properties,” on page 94](#).

Table 7-2 Cluster Protocols

Property	Description
Heartbeat	Specifies the interval of time in seconds between signals sent by each of the non-master nodes in the cluster to the master node to indicate that it is alive.
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	<p>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.</p> <p>Modify this parameter setting only when supervised by Novell Technical Support.</p>
Slave watchdog	<p>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.</p> <p>Modify this parameter setting only when supervised by Novell Technical Support.</p>
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 94](#).
- 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 94](#).
- 7 Click *Apply* or *OK* to save your changes.
- 8 Restart Cluster Services by entering the following at a terminal console prompt:

```
rcnovell-ncs restart
```

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.

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

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.

Modify this parameter setting only when supervised by Novell Technical Support.

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.

Modify this parameter setting only when supervised by Novell Technical Support.

7.3.5 Maximum Retransmits

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

7.4 Configuring Cluster Event Email Notification

Novell Cluster Services can automatically send out email messages for certain cluster events like cluster and resource state changes or nodes joining or leaving the cluster. The subject line provides information about the cluster name, resource name, action taken, and node name. For example:

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

IMPORTANT: Novell Cluster Services uses Postfix to send cluster email alerts, using the primary IP address. Postfix uses port 25 by default. If you have a cluster resource that uses SMTP (which also defaults to port 25), the port conflict between Postfix and SMTP can prevent that resource from working in the cluster. To avoid a port conflict, you can change the Postfix port configuration, or you can bind the clustered service to its resource’s secondary IP address.

To configure Postfix to use a different port, you can edit the `/etc/postfix/main.cf` file and change the values for the `inet_interfaces`, `mydestination`, and `mynetworks_style` directives. You can change the listen port for the `smtpd` process in the `/etc/postfix/master.cf` file. For more information about configuring Postfix, see [the Postfix Web site \(http://www.postfix.org/documentation.html\)](http://www.postfix.org/documentation.html).

To bind a service to its resource's secondary IP address, refer to the cluster documentation for the service. For example, the GroupWise Internet Agent (GWIA) uses SMTP. Because both Postfix and the GWIA default to using port 25, you must configure the GWIA to bind exclusively to its resource's secondary IP address in order to avoid a port conflict between Postfix and the GWIA. For information about how to set up the secondary IP address and to bind the GWIA to it, see the following sections in the *GroupWise 2012 Interoperability Guide* (http://www.novell.com/documentation/groupwise2012/gw2012_guide_interop/data/a20gkue.html):

- ♦ “Selecting the GWIA Partition and Secondary IP Address” (http://www.novell.com/documentation/groupwise2012/gw2012_guide_interop/data/bwe3bwa.html#bwe3bwk)
 - ♦ “Forcing Use of the GWIA Secondary IP Address” (http://www.novell.com/documentation/groupwise2012/gw2012_guide_interop/data/bwe3bx1.html#bwe3byz)
-

You can enable or disable email notification for the cluster and specify up to eight administrator email 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 email notification.
- 6 If you enable email notification, add up to eight email addresses in the field provided.
You can click the buttons next to the field to add, delete, or edit email addresses. Repeat this process for each email address you want on the notification list.
- 7 If you enable email 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 email 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 are ready to configure cluster resources for it. For information, see the following:

- ♦ [Chapter 9, "Configuring and Managing Cluster Resources," on page 127](#)
- ♦ [Chapter 10, "Creating and Managing Service Cluster Resources," on page 153](#)
- ♦ [Chapter 11, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 155](#)
- ♦ [Chapter 13, "Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers," on page 229](#)

For information about managing the cluster, see [Chapter 8, "Managing Clusters," on page 99](#).

8 Managing Clusters

After you have installed, set up, and configured Novell Cluster Services on Novell Open Enterprise (OES) 11 servers and configured cluster resources, use the information in this section to help you effectively manage your cluster. This section provides instructions for migrating resources, identifying cluster and resource states, and customizing cluster management.

IMPORTANT: For information about using console commands to manage a cluster, see [Appendix A, “Console Commands for Novell Cluster Services,”](#) on page 271.

- ♦ [Section 8.1, “Starting and Stopping Novell Cluster Services,”](#) on page 100
- ♦ [Section 8.2, “Monitoring Cluster and Resource States,”](#) on page 101
- ♦ [Section 8.3, “Generating a Cluster Configuration Report,”](#) on page 103
- ♦ [Section 8.4, “Cluster Migrating Resources to Different Nodes,”](#) on page 104
- ♦ [Section 8.5, “Onlining and Offlining \(Loading and Unloading\) Cluster Resources from a Cluster Node,”](#) on page 104
- ♦ [Section 8.6, “Removing \(Leaving\) a Node from the Cluster,”](#) on page 106
- ♦ [Section 8.7, “Joining a Node to the Cluster,”](#) on page 106
- ♦ [Section 8.8, “Shutting Down Linux Cluster Servers When Servicing Shared Storage,”](#) on page 106
- ♦ [Section 8.9, “Enabling or Disabling Cluster Maintenance Mode,”](#) on page 107
- ♦ [Section 8.10, “Preventing a Cluster Node Reboot after a Node Shutdown,”](#) on page 107
- ♦ [Section 8.11, “Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster,”](#) on page 108
- ♦ [Section 8.12, “Changing the NCS Proxy User Assignment in the NCS_Management_Group,”](#) on page 112
- ♦ [Section 8.13, “Adding a Node That Was Previously in the Cluster,”](#) on page 113
- ♦ [Section 8.14, “Deleting a Cluster Node from a Cluster,”](#) on page 113
- ♦ [Section 8.15, “Creating or Deleting Cluster SBD Partitions,”](#) on page 114
- ♦ [Section 8.16, “Configuring STONITH \(Feature Preview\),”](#) on page 122
- ♦ [Section 8.17, “Customizing Cluster Services Management,”](#) on page 124

8.1 Starting and Stopping Novell Cluster Services

Novell Cluster Services automatically starts after it is installed. It also automatically starts when you reboot your OES server.

If you need to restart `adminfs`, you must stop Novell Cluster Services before you stop `adminfs`, or you can reboot the 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:

```
chkconfig open-iscsi on
```

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 stop 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, select *Cluster > 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 Alert types are: Failback Alert Failover Alert Start Alert	Alerts can result if the Start mode, Failover mode, or Failback mode for the resource has been set to Manual. The resource is waiting for the administrator to intervene to start, fail over, or fail back the resource on the specified server.	On the Cluster Manager page in Novell iManager, click the <i>Respond to Alert</i> option to resolve the alert. Depending on the resource state, you are prompted to start, fail over, or fail back the resource. If you attempt to offline a resource that is in the <i>Start Alert</i> state, nothing happens and the following message is displayed: This operation cannot be completed. It is only available when the resource is in an offline state. You must clear the <i>Start Alert</i> before you can offline the resource. After the resource is offline, you can online the resource.
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 Mutual Exclusion Groups
Resource details	Information listed (as it applies) by resource for the master IP address resource and each of the cluster resources: Policies (Follows masters, Ignore quorum, Start mode, Failover mode, Failback mode, Assigned nodes) Load script and timeout Unload script and timeout Monitoring 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*.
When you select *Run Report*, the generated report opens in a new window, instead of in the child frame. This allows you to print and save the report from any Web browser.
- 4 Use the scroll bar to view the report in the browser.
- 5 (Optional) Use the browser options to print the report.
- 6 (Optional) Save the report to a text file.
 - 6a Click the cursor in the browser window where the report results are displayed, press Ctrl+A to select all, press Ctrl+C to copy the report to the clipboard.
 - 6b In a text editor, open a new document, then press Ctrl+V to paste the report in the document.
 - 6c Save the report document.

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, select *Clusters > 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

After a cluster resource is enabled, the load and unload scripts take care of starting and stopping the services or mounting and dismounting the volumes that are configured in the resource. You start services and mount devices by onlining the resource. You stop services and unmount devices by offlining the resource.

- ♦ [Section 8.5.1, “Guidelines for Onlining and Offlining Resources,” on page 104](#)
- ♦ [Section 8.5.2, “Using iManager to Online and Offline Resources,” on page 105](#)
- ♦ [Section 8.5.3, “Using the Cluster Online Command,” on page 105](#)
- ♦ [Section 8.5.4, “Using the Cluster Offline Command,” on page 105](#)
- ♦ [Section 8.5.5, “Using the Cluster Migrate Command,” on page 105](#)

8.5.1 Guidelines for Onlining and Offlining Resources

Onlining a resource runs the load script, which loads the resource on its primary preferred node, or on an alternate preferred node if the primary node is not available.

Offlining a resource runs the unload script and unloads the resource from the server. The resource cannot be loaded on any other servers in the cluster and remains unloaded until you load it again.

If monitoring is enabled for a resource and you start, stop, or restart services outside the cluster framework while the cluster resource is online, a monitoring failure occurs, which triggers a failover of the cluster resource.

If you edit the scripts for a resource that is online, the changes you made do not take effect until the resource is taken offline and brought online again.

If a resource is in a comatose state, you must take the resource offline before it can be brought online again.

There are maintenance situations where you might need to control services by using the standard interfaces. In these rare cases, you must first offline the cluster resource, then use non-cluster methods to start and stop a service or mount and dismount a volume outside the cluster framework. When you are done, online the cluster resource to resume normal cluster handling.

8.5.2 Using iManager to Online and Offline Resources

- 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*.

8.5.3 Using the Cluster Online Command

To bring a resource online on its most preferred node that is currently active:

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

```
cluster online resource_name
```

To bring a resource online on a specific active node:

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

```
cluster online resource_name node_name
```

8.5.4 Using the Cluster Offline Command

To take a resource offline:

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

```
cluster offline resource_name
```

8.5.5 Using the Cluster Migrate Command

The `cluster migrate` command can be used to move a specified resource from the node where it is currently running to the node you specify in the command. The destination node must be running in the cluster and also be assigned in the resource's *Preferred Nodes* list.

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

```
cluster migrate resource_name node_name
```

8.6 Removing (Leaving) a Node from the Cluster

When a node leaves a cluster, the node is no longer visible to the other nodes in the cluster. You might want a node to leave the cluster to perform maintenance on the server, or to upgrade the server during a rolling cluster upgrade.

- ♦ [Section 8.6.1, “Removing One Node at a Time,” on page 106](#)
- ♦ [Section 8.6.2, “Removing All Nodes at the Same Time,” on page 106](#)

8.6.1 Removing One Node at a Time

The `cluster leave` command allows you to remove a node from a cluster so that the node is not visible to other servers in the cluster.

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

```
cluster leave
```

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

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

8.6.2 Removing All Nodes at the Same Time

The `cluster down` command removes all cluster nodes from the cluster. This command has the same effect as executing the `cluster leave` command on every server in the cluster. You are prompted to confirm the cluster down.

- 1 Log in as the `root` user on any node in the cluster, then enter the following at a terminal console prompt:

```
cluster down
```

- 2 When you are prompted to confirm the `cluster down` command, specify *Yes*, then press Enter.

8.7 Joining a Node to the Cluster

The `cluster join` command allows you to add a node back to the cluster so that the node is again visible to other servers in the cluster.

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

```
cluster join
```

8.8 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.

8.9 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
```

8.10 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.

The Novell Cluster Services reboot behavior conforms to the kernel panic setting for the Linux operating system. By default the kernel panic setting is set for no reboot after a node shutdown.

You can set the kernel panic behavior in the `/etc/sysctl.conf` file by adding a `kernel.panic` command line. Set the value to 0 for no reboot after a node shutdown. Set the value to a positive integer value to indicate that the server should be rebooted after waiting the specified number of seconds. For information about the Linux `sysctl`, see the Linux man pages on `sysctl` and `sysctl.conf`.

- 1 As the `root` user, open the `/etc/sysctl.conf` file in a text editor.
- 2 If the `kernel.panic` token is not present, add it.

```
kernel.panic = 0
```

- 3 Set the `kernel.panic` value to 0 or to a positive integer value, depending on the desired behavior.
 - ♦ **No Reboot:** To prevent an automatic cluster reboot after a node shutdown, set the `kernel.panic` token to value to 0. This allows the administrator to determine what caused the kernel panic condition before manually rebooting the server. This is the recommended setting.

```
kernel.panic = 0
```

- ♦ **Reboot:** To allow a cluster node to reboot automatically after a node shutdown, set the `kernel.panic` token to a positive integer value that represents the seconds to delay the reboot.

```
kernel.panic = <seconds>
```

For example, to wait 1 minute (60 seconds) before rebooting the server, specify the following:

```
kernel.panic = 60
```

- 4 Save your changes.

8.11 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 8.11.1, “Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster,” on page 108](#)
- ♦ [Section 8.11.2, “Moving a Cluster or Changing IP Addresses of Cluster Nodes and Resources,” on page 109](#)

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

As a best practice, you should list the LDAP servers in the following order:

- ♦ local to the cluster
- ♦ closest physical read/write replica

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 Ensure that the IP addresses and administrator user name that you plan to use meet the requirements specified in [Section 4.2, “IP Address Requirements,” on page 39](#).
- 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: Ensure that you change the values inside the quotation marks to the actual settings for your cluster.

The following lines are the content of a sample configuration file for a 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"
```

8.11.2 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 110](#)
- ♦ [“Changing the IP Addresses of Cluster Resources” on page 110](#)
- ♦ [“Changing the IP Addresses of Servers in a Cluster” on page 111](#)
- ♦ [“Modifying the Cluster Configuration Information” on page 111](#)

Prerequisites

Before you begin, ensure that the IP addresses that you plan to use meet the requirements specified in [Section 4.2, “IP Address Requirements,” on page 39](#). Ensure that the administrator user name that you will use in the new location has sufficient rights as described in [Section 4.1.3, “Cluster Administrator or Administrator-Equivalent User,” on page 39](#).

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, select *Cluster > 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 the IP address on the Protocols page, then modify the IP address information as needed in the resource load, unload, and monitor scripts.
 - 2a In iManager, select *Clusters > 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 *Protocols* tab, modify the IP address for the resource, then click *Apply*.

This is necessary to change the IP address for the NCS:NCP Server object for the resource.
 - 2e Click the *Scripts* tab, then click the *Load Script* link.
 - 2f Edit the script by replacing variables with actual values for your new configuration, such as the resource IP address 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, volume group name, file system type, and mount device.

- 2g Click *Apply* to save the changed script.
 - 2h Make 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 111](#).

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 110](#)), 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 11 Server’s IP Address”](#) in the *OES 11: 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 111](#).

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 8.11.1, “Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster,” on page 108](#).
- 2 For each node in the cluster, including that of the master IP address resource, modify the NCS: *Network Address* attribute of its Cluster Node object.
 - 2a In iManager, select *Directory Administration > Modify Object*.
 - 2b Browse to locate and select the Cluster Node object of the cluster node you want to manage.
 - 2c In the *Valued Attributes* list, select the attribute *NCS: Network Address*, click *Edit*, modify the IP address, then click *OK*.
 - 2d Repeat this process for each node in the cluster and the master IP resource.
- 3 For the cluster container of the cluster you want to manage, modify the NCS: *Network Address* and *Network Address* attributes of its Cluster object to specify the new IP address information. Both *TCP* and *UDP* addresses need to be replaced.
 - 3a In iManager, select *Directory Administration > Modify Object*.
 - 3b Browse to locate and select the Cluster object of the cluster you want to manage.
 - 3c In the *Valued Attributes* list, select the attribute *NCS: Network Address* (the attribute for the TCP address), click *Edit*, modify the IP address, then click *OK*.
 - 3d In the *Valued Attributes* list, select the attribute *Network Address* (the attribute for the UDP address), click *Edit*, modify the IP address, then click *OK*.
- 4 Ensure that LDAP server is running before restarting Novell Cluster Services.

IMPORTANT: Novell Cluster Services requires LDAP.

- 5 Ensure that NSS is running if there are NSS cluster resources that you will be onlineing.

- 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, select *Cluster > 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*.

8.12 Changing the NCS Proxy User Assignment in the NCS_Management_Group

When you configure a cluster, the NCS_Management group is created in the Cluster object container. The group contains the user that you specify for the NCS Proxy User on the Proxy User Configuration page as described in [Step 3](#) in [Section 5.5.5, “Configuring a New Cluster,” on page 71](#):

- ♦ OES Common Proxy User
- ♦ LDAP Admin User
- ♦ Another administrator user

The specified user is automatically added as a member of the NCS_Management group. If you specify the OES Common Proxy User, each nodes' OES Common Proxy User is added to the group.

If the OES Common Proxy User is disabled in eDirectory for a node, the LDAP Admin automatically replaces that user in the NCS_Management_Group.

You can modify the users assigned as the NCS Proxy User by using the `/opt/novell/ncs/install/ncs_install.py` script. You are prompted for the distinguished user name and password for the user that you want to assign. The specified user is added to the NCS_Management group, and the old user is removed from the group.

- 1 Log in to the server as the root user, then open a terminal console.
- 2 At the console prompt, enter

```
/opt/novell/ncs/install/ncs_install.py -l -u
```

In addition, you can change other configurations by providing a configuration file (such as `/root/ncs.conf`) to the above command. You must create the file and its content before issuing the command. For example:

```
/opt/novell/ncs/install/ncs_install.py -l -f /root/ncs.conf -u
```

- 3 When you are prompted, provide the comma-delimited distinguished name of the user to be assigned to the NCS_Management group, such as

```
cn=ncs_admin,o=novell
```

- 4 When you are prompted, provide the password for the user you specified.

8.13 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 11 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 8.14, “Deleting a Cluster Node from a Cluster,” on page 113](#).
- 4 Run the Novell Cluster Services installation by following the procedure outlined in [Section 5.4.3, “Installing Novell Cluster Services on an Existing OES 11 Server,” on page 65](#).
The node assumes its former identity.

8.14 Deleting a Cluster Node from a Cluster

To permanently remove a cluster node from a cluster:

- 1 Log in as the `root` user to the node in the cluster that you want to remove, then enter the following at a 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, delete the node's Cluster Node object from the cluster container:
 - 3a Select *Directory Administration > Delete Objects*.
 - 3b Browse to the Cluster container (🍎) of the cluster, locate and select the Cluster Node object (🍎) for the node in the container, then click *OK*.
 - 3c On the Delete Objects page, click *OK*, then click *OK* again to confirm the deletion of the Cluster Node object.
- 4 Select *Directory Administration > Modify Object*, select the eDirectory Server object for the node that is being removed from the cluster, remove its NCS attributes, then click *OK* to save and apply your changes.
- 5 If the deleted cluster node persists in iManager, do one of the following:
 - ♦ Run the following command as the `root` user on the master cluster node. It is safe to run the command on an active cluster:

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

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

```
rcnovell-ncs restart
```

8.15 Creating or Deleting Cluster SBD Partitions

If the cluster has a shared disk system, the cluster must have a cluster partition for the Split Brain Detector (SBD) before you attempt to install Novell Cluster Services on the second node of the cluster or to create shared file systems. You might also need to delete and re-create an SBD partition if the SBD becomes corrupted or its device fails.

For information about how the split brain detector works, see *NetWare Cluster Services: The Gory Details of Heartbeats, Split Brains, and Poison Pills* (TID 10053882) (<http://support.novell.com/docs/Tids/Solutions/10053882.html>).

IMPORTANT: For instructions about setting up the SBD partition during the Novell Cluster Services configuration, see [Section 5.5.5, “Configuring a New Cluster,” on page 71](#).

- [Section 8.15.1, “Prerequisites for Creating an SBD Partition,” on page 114](#)
- [Section 8.15.2, “Before You Create a Cluster SBD Partition,” on page 115](#)
- [Section 8.15.3, “Creating a Non-Mirrored Cluster SBD Partition,” on page 115](#)
- [Section 8.15.4, “Creating a Mirrored Cluster SBD Partition,” on page 117](#)
- [Section 8.15.5, “Deleting a Non-Mirrored Cluster SBD Partition,” on page 120](#)
- [Section 8.15.6, “Deleting a Mirrored Cluster SBD Partition,” on page 121](#)
- [Section 8.15.7, “Removing a Segment from a Mirrored Cluster SBD Partition,” on page 122](#)

8.15.1 Prerequisites for Creating an SBD Partition

You must have a shared disk system (such as a Fibre Channel SAN or an iSCSI SAN) connected to your cluster nodes before attempting to create a cluster partition. For information, see [Section 4.8, “Shared Disk Configuration Requirements,” on page 53](#) for more information.

You need one small LUN in your storage array to use exclusively for the SBD partition. You can create another LUN of the same size to use as its mirror. The device should have at least 20 MB of free available space. About 4 MB of space that is used to store its shareable state information. For information about the SBD partition requirements, see [“SBD Partitions” on page 54](#).

IMPORTANT: The SBD Utility mirrors the devices. Use only single devices. You cannot present an existing NSS software RAID 1 device to the utility.

After you carve the device for the SBD partition, you must initialize it and mark it as *Shareable for Clustering*. You must also initialize and share a second device if you plan to mirror the SBD. You can use NSSMU, the Storage plug-in for iManager, or an NSS utility called `ncsinit` to initialize a device and set it to a shared state.

IMPORTANT: Check to see if a cluster SBD partition already exists before you create a new one. See [Section 8.15.2, “Before You Create a Cluster SBD Partition,” on page 115](#).

8.15.2 Before You Create a Cluster SBD Partition

Before creating an SBD partition, you should ensure 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 for the cluster, and identifies the device on the SAN where the SBD partition is located. It returns `NotFound` if the SBD partition does not exist.

- 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 8.15.5, “Deleting a Non-Mirrored Cluster SBD Partition,” on page 120](#)
 - ♦ [Section 8.15.7, “Removing a Segment from a Mirrored Cluster SBD Partition,” on page 122](#)

8.15.3 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 Ensure that the device you want to use has been initialized and marked as *Shareable for Clustering* as described in [Section 8.15.1, “Prerequisites for Creating an SBD Partition,” on page 114](#).
- 2 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 YaST Partitioner or NSSMU to check the names of the devices if needed. Use only the base (leaf) names (such as `sdb` or `mpathd`, not the full path) with the `-d` option.

For the `-s` option, replace *size* with the size (in MB) to use for the SBD partition. If the `-s` option is not used, the default size is 8 MB. For information about size requirements, see [“SBD Partitions” on page 54](#).

You can specify the size as `-1` to use all free space on the device. This option allows Novell Cluster Services to use a whole disk/LUN (or LUNs) that you set aside for SBD.

For example, the following command creates the `/dev/evms/.nodes/mycluster1.sbd` partition:

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

For example, the following command creates the `/dev/evms/.nodes/cl1.sbd` partition:

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

You can use the `sbdutil -f -s` command to view the path and name of the SBD partition.

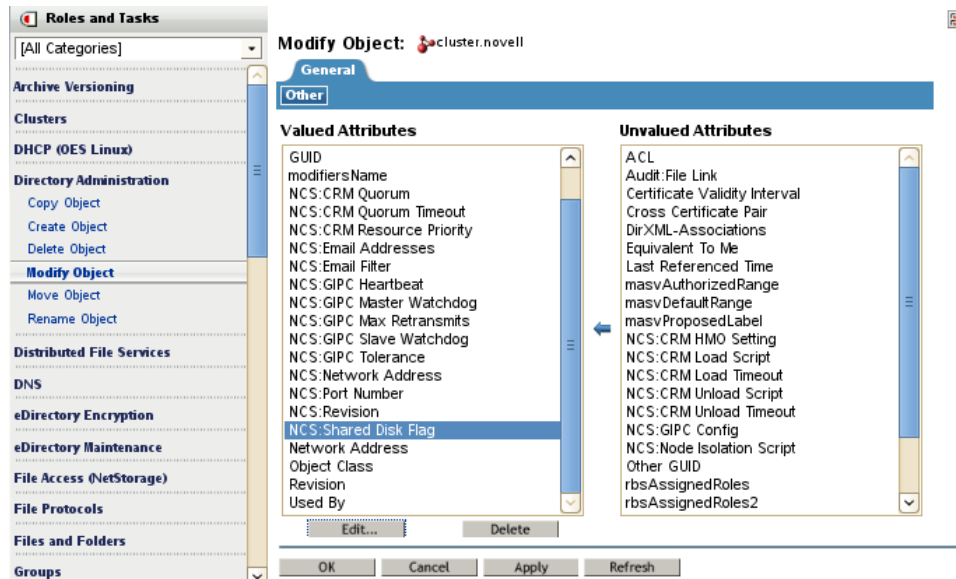
- 3 If the cluster has never had an SBD partition, modify the Cluster object in eDirectory to enable its *NCS: Shared Disk Flag* attribute.

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.

- 3a** 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.

- 3b** Select *Directory Administration*, then select *Modify Object*.
- 3c** Browse to locate and select the Cluster object of the cluster you want to manage, then click *OK*.
- 3d** Under *Valued Attributes*, select the *NCS: Shared Disk Flag*, then click *Edit*.



- 3e** Select (enable) the *NCS: Shared Disk Flag* check box, then click *OK*.



- 3f** Click *Apply* to save changes.
- 4** Reboot all cluster nodes.

8.15.4 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 NSSMU to create a mirrored cluster partition.

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

Using SBDUTIL to Create a Mirrored Cluster SBD Partition

- 1 Ensure that the devices you want to use have been initialized and marked as *Shareable for Clustering* as described in [Section 8.15.1, “Prerequisites for Creating an SBD Partition,” on page 114](#).
- 2 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 YaST Partitioner tool or NSSMU to check the names of the devices if needed. Use only 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. If the `-s` option is not used, the default size is 8 MB. For information about size requirements, see [“SBD Partitions” on page 54](#).

You can specify the size as `-1` to use all free space on the device. This option allows Novell Cluster Services to use a whole disk/LUN (or LUNs) that you set aside for SBD.

For example, the following command creates the `/dev/evms/.nodes/mycluster1.sbd` mirrored RAID device and the `mycluster1.msbd1` and `mycluster1.msbd2` partitions:

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

For example, the following command creates the `/dev/evms/.nodes/cl1.sbd` mirrored RAID device and the `cl1.msbd1` and `cl1.msbd2` partitions:

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

You can use the `sbdutil -f -s` command to view the path and name of the SBD RAID device. You can use NSSMU or the Storage plug-in to iManager to view the partitions used by the RAID device.

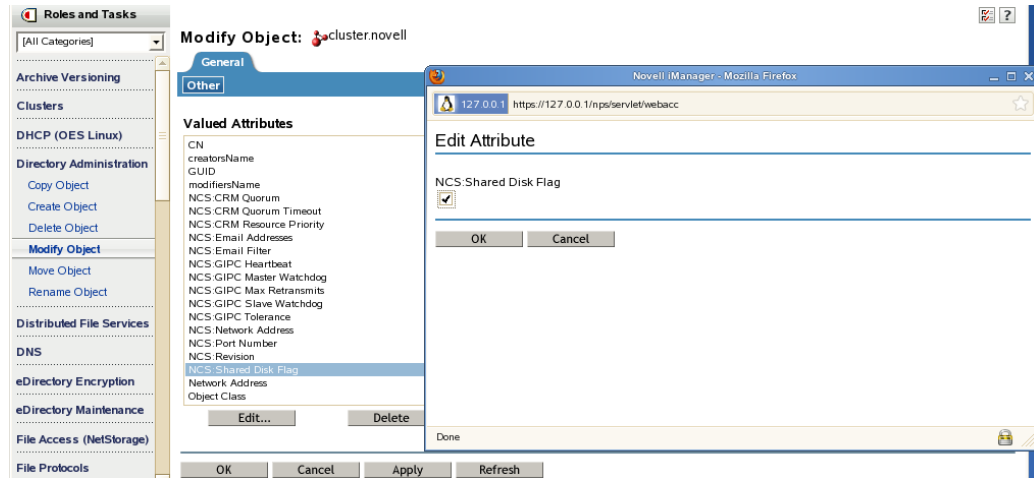
- 3 If the cluster has never had an SBD partition, modify the Cluster object in eDirectory to enable its *NCS: Shared Disk Flag* attribute.

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.

- 3a 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.

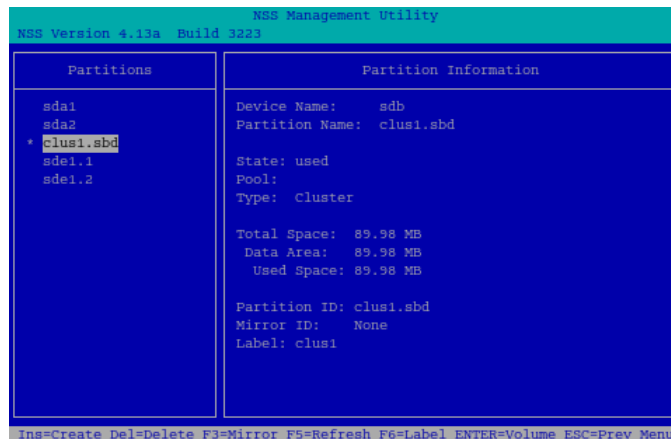
- 3b Click *Directory Administration*, then click *Modify Object*.
- 3c Browse to locate and select the Cluster object of the cluster you want to manage, then click *OK*.
- 3d Under *Valued Attributes*, select the *NCS: Shared Disk Flag*, then click *Edit*.
- 3e Select (enable) the *NCS: Shared Disk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



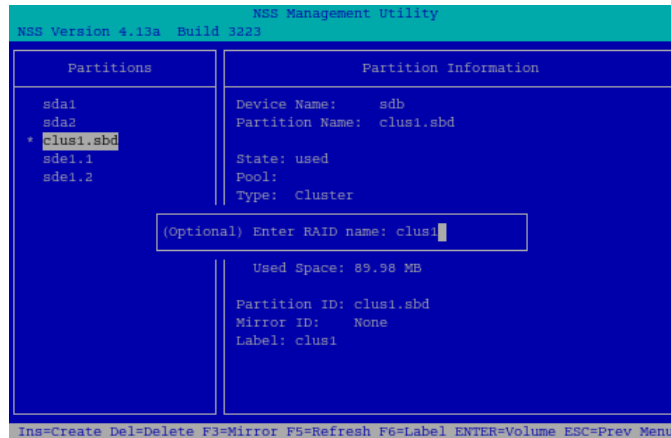
- 3f Click *Apply* to save changes.
- 4 Reboot all cluster nodes.

Using NSSMU to Mirror an Existing SBD Partition

- 1 Ensure that the devices you want to use have been initialized and marked as *Shareable for Clustering* as described in [Section 8.15.1, “Prerequisites for Creating an SBD Partition,”](#) on page 114.
- 2 At the Linux terminal console of a cluster server, log in as the root user, then enter `nssmu` to start NSSMU.
- 3 In NSSMU, mirror the SBD partition:
 - 3a Select *Partitions* from the main menu.
 - 3b From the list of partitions, select the SBD partition that you want to mirror.

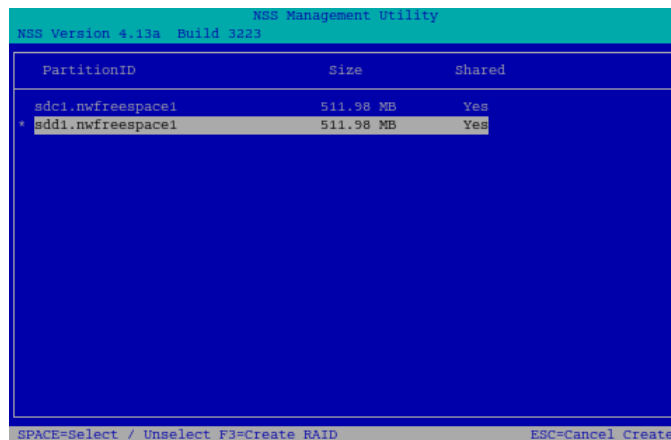


- 3c** Press F3 to open the dialog to create the RAID1 mirror.
- 3d** Type a name for the RAID with the same name as the SBD partition but without the .sbd extension, then press Enter.



Typically, the SBD partition name is the same as the cluster name plus the .sbd extension, such as `clus1.sbd`. In this case, the SBD RAID1 device should be named `clus1`.

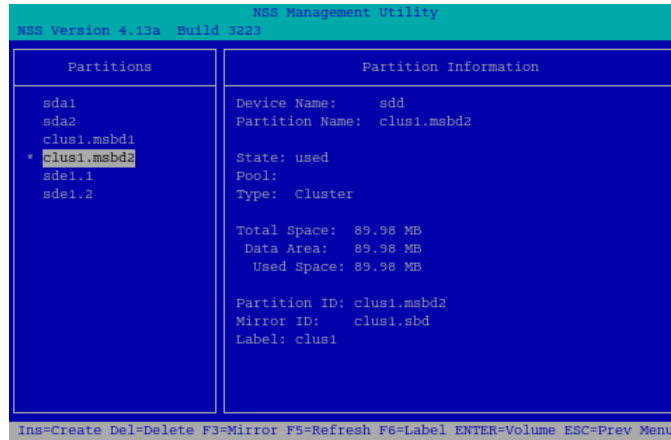
- 3e** From the list of available devices, select a device to use as the second segment of the mirror, then press the space bar to choose it.



When the device is selected, the asterisk next to the device stays there even if you move the cursor up and down in the list.

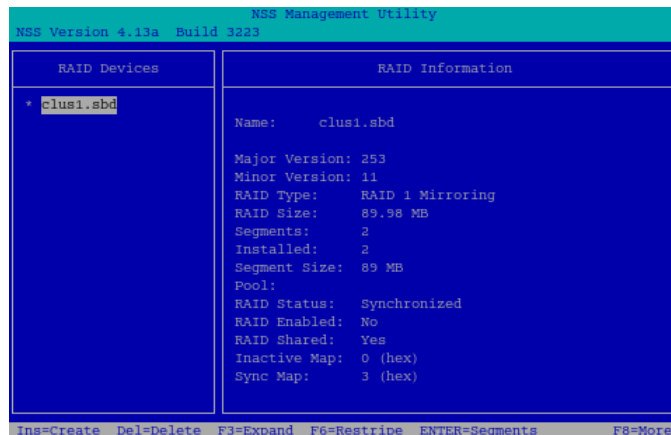
- 3f** Press F3 again to create the mirror.

- 3g** In the confirmation message, press Y (Yes) to approve the RAID1 creation.



The SBD partitions are named with the partition name, and now have the extension of .msbd1 (mirrored SBD partition 1) and .msbd2 (mirrored SBD partition 2).

- 4** In NSSMU, verify that the RAID1 device was created for the SBD:
- 4a** Select *RAID Devices* from the main menu.
 - 4b** Select the RAID1 device that you created to view details about the mirrored SBD device (such as clus1.sbd).



- 4c** View the RAID Status to ensure that synchronization has begun.
Synchronization is complete when the status is *Synchronized*.
- 5** Exit NSSMU.
- 6** Reboot all cluster nodes.

8.15.5 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 Stop Novell Cluster Services by entering the following on each node:

```
rcnovell-ncs stop
```

- 4 Delete the SBD partition.

- 4a Enter `nssmu` to open the NSS Utility, then select *Partitions*.

- 4b Select the SBD partition you want to delete.

- 4c Click *Delete* to delete the partition and its contents, then click *OK* to confirm the deletion.

- 5 If you have more than one node in the cluster, use one of the following methods to create a new SBD partition:

- ♦ [Section 8.15.3, “Creating a Non-Mirrored Cluster SBD Partition,” on page 115](#)
- ♦ [Section 8.15.4, “Creating a Mirrored Cluster SBD Partition,” on page 117](#)

Ensure that the SBD partition exists before continuing.

- 6 Start Novell Cluster Services by entering

```
rcnovell-ncs start
```

- 7 Join the nodes to the cluster by entering `cluster join` at the terminal console on each node in the cluster.

8.15.6 Deleting a Mirrored Cluster SBD Partition

Before you attempt to create (or re-create) an SBD partition, you must delete an existing SBD partition for the cluster. If the SBD partition 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 NSSMU utilities.

- 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 Stop Novell Cluster Services by entering

```
rcnovell-ncs stop
```

- 4 Delete the mirrored software RAID that you used for the SBD partition.

- 4a Enter `nssmu` to open the NSS Utility, then select *Software RAIDS*.

- 4b Select the software RAID 1 for the SBD partition you want to delete.

- 4c Click *Delete* to delete the software RAID and its member segments, then click *OK* to confirm the deletion.

- 5 If you have more than one node in the cluster, use one of the following methods to create a new SBD partition:

- ♦ [Section 8.15.3, “Creating a Non-Mirrored Cluster SBD Partition,” on page 115](#)
- ♦ [Section 8.15.4, “Creating a Mirrored Cluster SBD Partition,” on page 117](#)

Ensure that the SBD partition exists before continuing.

- 6 Start Novell Cluster Services by entering

```
rcnovell-ncs start
```

- 7 Join the nodes to the cluster by entering `cluster join` at the terminal console of each node in the cluster:

8.15.7 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 8.15.6, “Deleting a Mirrored Cluster SBD Partition,” on page 121](#), then re-create a non-mirrored SBD partition, as described in [Section 8.15.3, “Creating a Non-Mirrored Cluster SBD Partition,” on page 115](#).

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 At the terminal console of one cluster server, enter

```
cluster maintenance on
```

This causes all cluster servers to enter maintenance mode.

- 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 At the terminal console of one cluster server, enter

```
cluster maintenance off
```

This causes all cluster servers to return to normal mode.

8.16 Configuring STONITH (Feature Preview)

The STONITH (shoot-the-other-node-in-the-head) capability allows Novell Cluster Services to remotely kill a suspect node by using remote power control instead of using a poison pill. STONITH does not require any action from the node being killed, unlike poison pills, which allows it to kill non-responsive nodes.

Using STONITH requires that you have server power management technology for all nodes in the cluster. STONITH supports remote accessible cards integrated in a cluster node's hardware, such as Integrated Lights Out (iLO) from Hewlett-Packard (HP) and Dell Remote Access Card (DRAC) from Dell, and stand-alone Web-based power switches. Refer to the vendor documentation to determine how your power management system works.

To use STONITH in Novell Cluster Services, you must create an executable `/opt/novell/ncs/bin/NCS_STONITH_SCRIPT` script file to authenticate to your power controller and turn off the power, cycle the power, or reset the power for the node. The script should take the node number as the only parameter. Node numbers are assigned as 0 to 31 for nodes 1 to 32. These are the same node numbers that appear in the `/var/opt/novell/ncs/gipc.conf` file. Creating the script file automatically enables STONITH for the node; you do not need to restart anything.

IMPORTANT: STONITH does not replace poison pills in Novell Cluster Services. Novell Cluster Services issues poison pills before running the STONITH script.

Use the following sample scripts as a guide for how to create a script that works with your power management technology:

- ♦ [Section 8.16.1, “Sample Script for HP iLO,” on page 123](#)
- ♦ [Section 8.16.2, “Sample Script for Web-Based Power Switches,” on page 124](#)

8.16.1 Sample Script for HP iLO

This section provides a sample script for HP iLO power management cards. The sample code assumes the following setup for the cluster and iLO cards:

- ♦ The iLO card on each node has been pre-configured to trust any instructions sent from each of the nodes in the same cluster.
- ♦ Each cluster node's iLO card is assigned a sequential static IP address, beginning with 192.168.188.201 on node 0, 192.168.188.202 on node 2, and so on up to 192.168.188.232 on node 32.

If you alternatively use DNS names in your script, the translation must be performed by the script.

- ♦ The iLO card's command to reset power is `power reset`.

Refer to the HP documentation to determine the commands available for your iLO cards.

For each node in the cluster, create the executable `/opt/novell/ncs/bin/NCS_STONITH_SCRIPT` file and add the script below. The presence of the file automatically enables STONITH for the cluster.

IMPORTANT: Ensure that you replace the sample information with the settings for your system.

```
#!/bin/bash

if [ -n "$1" ]; then
    echo "Recycling power of node number $1 ... "
    iloIP=$(printf "192.168.188.2%02d" $((($1+1)))
    nodeIP=`grep "^nodeid .* ${1}\>" /var/opt/novell/ncs/gipc.conf | cut -d' ' -f2`
    while [ 1 ]; do
        ssh ${iloIP} power reset
        echo
        sleep 2
        if ! ping -c 1 ${nodeIP} | grep "1 received" &> /dev/null
        then
            break
        fi
        sleep 5
    done
fi
```

8.16.2 Sample Script for Web-Based Power Switches

This section provides a sample script for Web-based power switches. The sample code assumes the following setup for the cluster and power switch:

- ♦ Each cluster node is assigned a sequential static IP address, beginning with 10.10.189.100.
- ♦ Each cluster node is plugged in to a sequential outlet in the power switch, beginning with node 0 in the first outlet.
- ♦ The authentication information for the power switch management interface is:
 - ♦ **User name:** admin
 - ♦ **Password:** novell
 - ♦ **IP address:** 10.10.189.149
- ♦ The power management switch's command to cycle power is CCL.

Each vendor can have different command options available. Refer to your vendor documentation to determine the commands used by your power switch.

For each node in the cluster, create the executable `/opt/novell/ncs/bin/NCS_STONITH_SCRIPT` file and add the script below. The presence of the file automatically enables STONITH for the cluster.

IMPORTANT: Ensure that you replace the sample information with the settings for your system.

```
#!/bin/bash

until [ -z "$1" ]
do
    nodeNum=`expr $1 + 1`
    echo "Recycling power of node number $nodeNum ... "
    while [ 1 ]; do
        curl -u admin:novell http://10.10.189.149/outlet?${nodeNum}=CCL
        echo
        sleep 2
        if ! ping -c 1 10.10.189.10${nodeNum} | grep "1 received" &> /dev/null
        then
            break
        fi
        sleep 5
    done
    shift
done
```

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

[Table 8-4](#) lists the cluster-related CMD files and gives a brief description of each.

Table 8-4 Cluster-Related CMD Files

CMD File Name	Description
Node.cmd	Write-then-read command file used in conjunction with a Perl script to issue node-specific commands to the cluster and retrieve resulting node status and configuration information.
Cluster.cmd	Write-then-read command file used in conjunction with a Perl script to issue cluster-specific commands to the cluster and retrieve resulting cluster status and configuration information.
Resource.cmd	Write-then-read command file used in conjunction with a Perl script to issue resource-specific commands to the cluster and retrieve resulting resource status and configuration information.

9 Configuring and Managing Cluster Resources

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 99.

- ♦ [Section 9.1, “Planning for Cluster Resources,”](#) on page 127
- ♦ [Section 9.2, “Using Cluster Resource Templates,”](#) on page 129
- ♦ [Section 9.3, “Creating Cluster Resources,”](#) on page 133
- ♦ [Section 9.4, “Configuring a Load Script for a Cluster Resource,”](#) on page 133
- ♦ [Section 9.5, “Configuring an Unload Script for a Cluster Resource,”](#) on page 134
- ♦ [Section 9.6, “Enabling Monitoring and Configuring the Monitor Script,”](#) on page 135
- ♦ [Section 9.7, “Configuring the Start, Failover, and Failback Modes for Cluster Resources,”](#) on page 139
- ♦ [Section 9.8, “Configuring Preferred Nodes for a Resource,”](#) on page 140
- ♦ [Section 9.9, “Configuring Resource Priorities for Load Order,”](#) on page 142
- ♦ [Section 9.10, “Configuring Resource Mutual Exclusion Groups,”](#) on page 143
- ♦ [Section 9.11, “Controlling Resource Monitoring,”](#) on page 147
- ♦ [Section 9.12, “Changing the IP Address of a Cluster Resource,”](#) on page 147
- ♦ [Section 9.13, “Renaming a Cluster Resource,”](#) on page 147
- ♦ [Section 9.14, “Deleting Cluster Resources,”](#) on page 148
- ♦ [Section 9.15, “Additional Information for Creating Cluster Resources,”](#) on page 150

9.1 Planning for Cluster Resources

Consider the guidelines in this section when planning for your cluster resources.

- ♦ [Section 9.1.1, “Naming Conventions for Cluster Resources,”](#) on page 128
- ♦ [Section 9.1.2, “Using Parameters with Space Characters in a Cluster Script,”](#) on page 128
- ♦ [Section 9.1.3, “Script Length Limits,”](#) on page 128
- ♦ [Section 9.1.4, “Planning Cluster Maintenance,”](#) on page 128
- ♦ [Section 9.1.5, “Number of Resources,”](#) on page 128
- ♦ [Section 9.1.6, “Linux POSIX File System Types,”](#) on page 129

9.1.1 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 and Linux volume names are automatically used in the cluster resource name, do not use special characters !@#\$%&() in names of NSS pools or Linux POSIX volumes that you plan to cluster enable.

9.1.2 Using Parameters with Space Characters in a Cluster Script

In cluster scripts, if a parameter contains space characters, you must quote the parameter twice. For example, you can use one of the following techniques:

```
exit_on_error echo "'Errors will be reported here.'"
```

```
exit_on_error echo '"Errors will be reported here.'"'
```

```
ignore_error echo "\"Errors will be ignored here.\""
```

9.1.3 Script Length Limits

Each cluster load, unload, or monitor script can be up to 3200 bytes in length. This limit includes commands, comments, and spaces. For example, non-special ASCII characters (including the space character) are 1 byte per character, so you could have up to 3200 of these characters in a script.

IMPORTANT: Creating a script that exceeds the maximum script length can prevent a resource from loading. If your script commands and comments require more memory than 3200 bytes, you can spawn external scripts from a script.

9.1.4 Planning Cluster Maintenance

When performing cluster maintenance tasks, two cluster best practices should be observed:

- ♦ Perform maintenance tasks during non-peak hours so that users are minimally affected.
- ♦ Before performing maintenance on a node, cluster migrate its cluster resources to another node if you want the related users to be undisturbed.

IMPORTANT: Changes for a resource's properties are not applied while the resource is loaded or running on a server. You must offline the resource, then online it again in order to apply any changes that you make to its resource properties, policies, or scripts.

9.1.5 Number of Resources

Novell Cluster Services supports up to 254 resources in a cluster, regardless of the size of the cluster.

9.1.6 Linux POSIX File System Types

Ext3 is the default file system type used in the template scripts. The Ext2, Ext3, ReiserFS, and XFS file systems have been tested and are fully supported.

9.2 Using Cluster Resource Templates

Templates simplify the process of creating similar or identical cluster resources. For example, templates are helpful when you want to create multiple instances of the same resource on different servers. Several templates are provided for you. You can also create templates for any server application or resource you want to add to your cluster.

- ♦ [Section 9.2.1, “Default Resource Templates,” on page 129](#)
- ♦ [Section 9.2.2, “Creating a Resource Template,” on page 131](#)
- ♦ [Section 9.2.3, “Synchronizing Locally Modified Resource Templates with eDirectory,” on page 132](#)

9.2.1 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 9.2.2, “Creating a Resource Template,” on page 131](#). 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 11 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)	LVM volume groups
Generic IP Service	This template can be modified to create cluster resources for certain server applications that run on your cluster.
iFolder	Novell iFolder 3.9 The template scripts contain examples for NSS pools and LVM volume groups.
iPrint	Novell iPrint
MySQL	MySQL 5.0x for Linux
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. This template is available only if Xen is installed on the host server. For information, see Section 14.2.1, “Creating a Xen Virtual Machine Cluster Resource,” on page 251.
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. This template is available only if Xen is installed on the host server. For information, see Section 14.2.3, “Setting Up Live Migration,” on page 257.

To view the default templates in iManager:

- 1 In iManager *Roles and Tasks*, select *Clusters > Cluster Options*.
- 2 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.

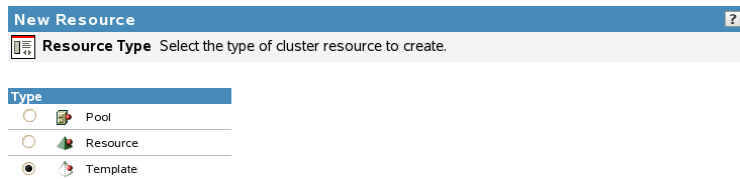
You can also view the templates outside of iManager. They are cached as `/var/opt/novell/ncs/*_Template.*` files on the master node of the cluster.

9.2.2 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 In iManager *Roles and Tasks*, select *Clusters > Cluster Options*.
- 2 Browse to locate and select the Cluster object of the cluster you want to manage.
- 3 Click the *New* link.
- 4 Specify *Template* as the resource type you want to create by clicking the *Template* radio button, then click *Next*.

Clusters > Cluster Options



- 5 In *Cluster Resource Name*, specify the name of the template you want to create.
- 6 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.
- 7 Ensure that the *Define Additional Properties* check box is selected, then click *Next* to continue to the Load Script page.
- 8 On the Load Script page, configure the load script for the cluster resource template.
 - 8a Edit or add variables with example values for your template configuration, such as the mount point, IP address, volume group name, file system type, and mount device.
 - 8b Edit or add any lines to the load script that are required to load dependent services such as Web servers or file access protocols.
 - 8c 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.
 - 8d 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.
- 9 On the Unload Script page, configure the unload script for the cluster resource template.
 - 9a Edit or add variables with example values for your template configuration, such as the mount point, IP address, volume group name, file system type, and mount device.
 - 9b Edit or add 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.

- 9c** Edit or add any lines to the unload script that are required to unload the dependent services that are loaded by this cluster resource.

- 9d** 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.

- 10** On the Monitor Script page, configure the monitor script for the cluster resource template.

- 10a** Edit or add the variables with example values for your template configuration, such as the mount point, IP address, volume group name, file system type, and mount device.

- 10b** 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.

- 10c** Specify the default *Monitor Script Timeout* value, then click *Next*.

The timeout value determines how much time the script is given to complete. If the script does not complete within the specified time, the failure action the administrator chooses for monitoring (comatose, migrate, or reboot) initiates. Cluster Services marks the process as failed right after the defined timeout expires, but it must wait for the process to conclude before it can start other resource operations.

- 11** On the Resource Policies page, specify the default *Start*, *Failover*, and *Failback* modes, then click *Next*.

- 12** 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.

9.2.3 Synchronizing Locally Modified Resource Templates with eDirectory

Typically, you update the templates by using the Clusters plug-in to iManager. However, if you add new template files or modify the files locally on a cluster node, you must synchronize those changes with the resource templates and scripts that are held in eDirectory .

To synchronize the locally modified resource templates with those held in the Cluster container in eDirectory,

- 1** Log in to the master node of the cluster, then open a terminal console.
- 2** Enter the following commands:

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

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

9.3 Creating Cluster Resources

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

- 1 In iManager *Roles and Tasks*, select *Clusters > Cluster Options*.
iManager displays four links under Clusters that you can use to configure and manage your cluster.
- 2 Browse to locate and select the Cluster object of the cluster you want to manage.
- 3 Click the *New* link.
- 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.

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.

- 6 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 9.2, “Using Cluster Resource Templates,” on page 129](#).
- 7 Select the *Define Additional Properties* check box, then click *Next*.
- 8 If you are creating a new cluster resource, continue with [“Configuring a Load Script for a Cluster Resource” on page 133](#).

9.4 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 11.5, “Configuring a Load Script for the Shared NSS Pool,” on page 171](#)
- ♦ [Section 13.2.3, “Sample Load Script for a CSM Resource,” on page 232](#)
- ♦ [“Sample LVM Resource Load Script” on page 219](#)

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, select *Clusters > 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, and so on.

You also need to personalize the script by replacing variables with actual values for your specific configuration, such as the mount point, IP address, volume group name, file system type, and mount device.

IMPORTANT: Do not comment out commands that are automatically generated for parameters that define the cluster resource, such as the mount point, IP address, volume group name, file system type, and device. If you need to modify the IP address, administrator credentials, or other attributes of an existing resource, follow the procedure in [Section 8.11, “Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster,”](#) on page 108.

- 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.

The timeout value is applied only when the resource is migrated to another node. It is not used during resource online/offline procedures.

- 7 Do one of the following:

- ♦ If you are configuring a new resource, click *Next*, then continue with [Section 9.5, “Configuring an Unload Script for a Cluster Resource,”](#) on page 134.
- ♦ Click *Apply* to save your changes.

Changes for a resource’s properties are not applied while the resource is loaded or running on a server. You must offline, then online the resource to activate the changes for the resource.

9.5 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 11.6, “Configuring an Unload Script for the Shared NSS Pool,”](#) on page 173
- ♦ [Section 13.2.4, “Sample Unload Script for CSM Resource,”](#) on page 233
- ♦ [“Sample LVM Resource Unload Script”](#) on page 219

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, volume group name, file system type, and mount 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.

The timeout value is applied only when the resource is migrated to another node. It is not used during resource online/offline procedures.

7 Do one of the following:

- ♦ If you are configuring a new resource, click *Next*, then continue with [Section 9.6.2, “Configuring Resource Monitoring,” on page 137](#).
- ♦ Click *Apply* to save your changes.

Changes for a resource’s properties are not applied while the resource is loaded or running on a server. You must offline, then online the resource to activate the changes for the resource.

9.6 Enabling Monitoring and Configuring the Monitor Script

Resource monitoring allows Novell Cluster Services to detect when an individual resource on a node has failed independently of its ability to detect node failures. Monitoring is disabled by default. It is enabled separately for each cluster resource.

- ♦ [Section 9.6.1, “Understanding Resource Monitoring,” on page 135](#)
- ♦ [Section 9.6.2, “Configuring Resource Monitoring,” on page 137](#)
- ♦ [Section 9.6.3, “Example Monitoring Scripts,” on page 138](#)
- ♦ [Section 9.6.4, “Monitoring Services that Are Critical to Clustering,” on page 138](#)

9.6.1 Understanding Resource Monitoring

When you enable resource monitoring, you must specify a polling interval, a failure rate, a failure action, and a timeout value. These settings control how error conditions are resolved for the resource.

- ♦ [“Polling Interval” on page 136](#)
- ♦ [“Failure Rate” on page 136](#)
- ♦ [“Failure Action” on page 136](#)
- ♦ [“Timeout Value” on page 136](#)
- ♦ [“How Resource Monitoring Works” on page 137](#)

Polling Interval

The monitoring script runs at a frequency specified by the polling interval. By default, it runs every minute when the resource is online. You can specify the polling interval in minutes or seconds. The polling interval applies only to a given resource.

Failure Rate

The failure rate is the maximum number of failures (*Maximum Local Failures*) detected by the monitoring script during a specified amount of time (*Time Interval*).

A failure action is initiated when the resource monitor detects that the resource fails more times than the maximum number of local failures allowed to occur during the specified time interval. For failures that occur before it exceeds the maximum, Cluster Services automatically attempts to unload and load the resource. The progress and output of executing a monitor script are appended to `/var/opt/novell/log/ncs/resource_name.monitor.out` file.

For example, if you set the failure rate to 3 failures in 10 minutes, the failure action is initiated if it fails 4 times in a 10 minute period. For the first 3 failures, Cluster Services automatically attempts to unload and load the resource.

Failure Action

The *Failover Action* indicates whether you want the resource to be set to a comatose state, to migrate to another server, or to reboot the hosting node (without synchronizing or unmounting the disks) if a failure action initiates. The reboot option is normally used only for a mission-critical cluster resource that must remain available.

If the failure action initiates and you chose the option to migrate the resource to another server, the resource migrates to the next server in its *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.

If the failure action initiates and you chose the option to reboot the hosting node without synchronizing or unmounting the disks, each of the resources on the hosting node will fail over to the next server in its *Assigned Nodes* list because of the reboot. This is a hard reboot, not a graceful one.

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.

Timeout Value

The timeout value determines how much time the script is given to complete. If the script does not complete within the specified time, the configured failure action is initiated. Cluster Services marks the process as failed right after the defined timeout expires, but it must wait for the process to conclude before it can start other resource operations.

The timeout value is applied only when the resource is migrated to another node. It is not used during resource online/offline procedures.

How Resource Monitoring Works

- 1 The monitoring script runs at the frequency you specify as the polling interval.
- 2 There are two conditions that trigger a response by Novell Cluster Services:
 - ♦ An error is returned. Go to [Step 3](#).
 - ♦ The script times out, and the process fails. Go to [Step 4](#).
- 3 Novell Cluster Services tallies the error occurrence, compares it to the configured failure rate, then does one of the following:
 - ♦ **Total errors in the interval are less than or equal to the Maximum Local Failures:** Novell Cluster Services tries to resolve the error by offlining the resource, then onlineing the resource.

If this problem resolution effort fails, Novell Cluster Services goes to [Step 4](#) immediately regardless of the failure rate condition at that time.
 - ♦ **Total errors in the interval are more than the Maximum Local Failures:** Go to [Step 4](#).
- 4 Novell Cluster Services initiates the configured failure action. Possible actions are:
 - ♦ Puts the resource in a comatose state
 - ♦ Migrates the resource to another server
 - ♦ Reboots the hosting node (without synchronizing or unmounting the disks)

9.6.2 Configuring Resource Monitoring

The resource monitoring function allows you to monitor the health of a specified resource by using a script that you create or customize. If you want 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, a failure action, and a timeout value.

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

- 1 In iManager, select *Clusters > 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 specify the value in minutes or seconds.
- 7 Specify the number of failures (*Maximum Local Failures*) for the specified amount of time (*Time Interval*).

For information, see [“Failure Rate” on page 136](#).

- 8 Specify the *Failover Action* by indicating whether you want the resource to be set to a comatose state, to migrate to another server, or to reboot the hosting node (without synchronizing or unmounting the disks) if a failure action initiates. The reboot option is normally used only for a mission-critical cluster resource that must remain available.

For information, see [“Failure Action” on page 136](#).

- 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.

The resource templates included with Novell Cluster Services for Linux include resource monitoring scripts that you can customize.

You also need to personalize the script by replacing variables with actual values for your specific configuration, such as the mount point, IP address, volume group name, file system type, and mount device.

You can use the same commands that would be used at the Linux terminal console. For example, see [Section 9.6.4, “Monitoring Services that Are Critical to Clustering,” on page 138](#).

- 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.

- 12 Do one of the following:

- ♦ If you are configuring a new resource, click *Next*, then continue with [Section 9.7.2, “Setting the Start, Failover, and Failback Modes for a Resource,” on page 140](#).
- ♦ Click *Apply* to save your changes.

Changes for a resource’s properties are not applied while the resource is loaded or running on a server. You must offline, then online the resource to activate the changes for the resource.

9.6.3 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 11.7, “Configuring a Monitor Script for the Shared NSS Pool,” on page 174](#)
- ♦ [Section 13.2.5, “Sample Monitor Script for a CSM Resource,” on page 233](#)
- ♦ [“Sample LVM Resource Monitor Script” on page 220](#)

9.6.4 Monitoring Services that Are 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 (ndscd). 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 rcnamcd status
```

You can use the `namcd status` command instead of `rcnamcd status` in the Monitor script if you want to automatically restart `namcd` if its daemon is not loaded and running. However, `namcd` creates messages in `/var/log/messages` with each check.

9.7 Configuring the Start, Failover, and Failback Modes for Cluster Resources

You can configure the start, failover, and failback of cluster resources to happen manually or automatically.

IMPORTANT: Cluster Services works with NCP user connections so that the user data sessions are resumed after failover. However, non-NCP users might experience service interruption and need to reconnect to the server after the failover. Applications using server based storage must be restarted on the client even with NCP unless they are NCP reconnect aware.

- [Section 9.7.1, “Understanding Cluster Resource Modes,” on page 139](#)
- [Section 9.7.2, “Setting the Start, Failover, and Failback Modes for a Resource,” on page 140](#)

9.7.1 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 if there is a hardware or software failure. If the resource *Failover* mode is set to `MANUAL`, you can intervene after a failure occurs and before the resource is moved to another node.

With the resource *Failback* mode set to `DISABLE`, the resource does not fail back to its most preferred node when the most preferred node rejoins the cluster. If the resource *Failback* mode is set to `AUTO`, the resource automatically fails back to its most preferred node when the most preferred node rejoins the cluster. Set the resource *Failback* mode to `MANUAL` to prevent the resource from moving back to its preferred node when that node is brought back online, until you are ready to allow it to happen.

The preferred node is the first server in the *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.

9.7.2 Setting the Start, Failover, and Failback Modes for a Resource

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

- 1 In iManager, select *Clusters > 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 (Optional) Select the *Resource Follows Master* check box if you want to ensure that the resource runs only on the master node in the cluster.

If the master node in the cluster fails, the resource fails over to whichever node becomes the master.

- 6 (Optional) Select the *Ignore Quorum* check box if you don't want the cluster-wide timeout period and node number limit enforced.

The quorum default values were set when you installed 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 Do one of the following:

- ♦ If you are configuring a new resource, click *Next*, then continue with [Section 9.8, "Configuring Preferred Nodes for a Resource,"](#) on page 140.
- ♦ Click *Apply* to save your changes.

Changes for a resource's properties are not applied while the resource is loaded or running on a server. You must offline, then online the resource to activate the changes for the resource.

9.8 Configuring Preferred Nodes for a Resource

The Preferred Nodes page allows you to control the order in which a resource attempts to fail over to one of its assigned nodes when the resource is brought online or during a failover. For example, if a node fails, the resource attempts to fail over to the one of the assigned nodes in its Preferred Nodes

list. The preferred failover order determines which node is tried first. This is useful for ensuring that the resource runs on its most preferred node of the nodes available when it fails over or is brought online. Changes are not allowed for the Preferred Nodes list for the `Master_IP_Address_Resource`.

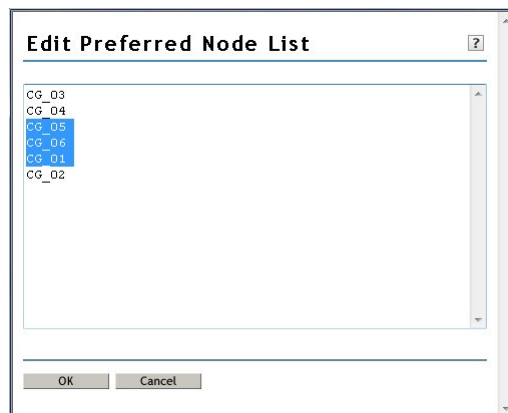
If you are creating a new cluster resource, the Preferred Nodes page should already be displayed. If you are assigning nodes for an existing resource, the Preferred Nodes page is displayed as part of the Resource Policies page. You can start with [Step 5](#).

- 1 In iManager, select *Clusters > 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 Use either of the following methods to change the preferred failover order of the nodes assigned to a resource:
 - ♦ **Arrows:** Select one of the assigned nodes by clicking on it, then click the up-arrow and down-arrow buttons to move the node up or down in the list. The page refreshes between each click on an arrow.
 - ♦ **Edit:** Open the Preferred Nodes Edit function by clicking the *Edit* pen icon, list the assigned nodes in the preferred failover order with one node per line, then click *OK* to accept the revised order.

Be careful to not alter the names of the nodes while editing the list.



- 8 Click *Apply* to save node assignment changes.
- 9 Offline the resource, then online it again to apply the changes to the resource.

Changes for a resource's properties are not applied while the resource is loaded or running on a server.

9.9 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.

Clusters > Cluster Options

Cluster Properties: cluster.novell

Policies **Priorities** Protocols RME Groups Business Continuity

View or change cluster resource priorities.

Resource Priorities:

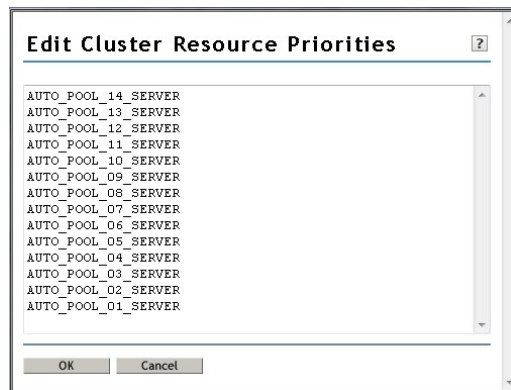
```
AUTO_POOL_14_SERVER
AUTO_POOL_13_SERVER
AUTO_POOL_12_SERVER
AUTO_POOL_11_SERVER
AUTO_POOL_10_SERVER
AUTO_POOL_09_SERVER
AUTO_POOL_08_SERVER
AUTO_POOL_07_SERVER
AUTO_POOL_06_SERVER
AUTO_POOL_05_SERVER
AUTO_POOL_04_SERVER
AUTO_POOL_03_SERVER
AUTO_POOL_02_SERVER
AUTO_POOL_01_SERVER
```



- 5 Use either of the following methods to change the load order (from highest priority to lowest priority) of a resource relative to other cluster resources on the same node:
 - ♦ **Arrows:** Select a resource in the list by clicking on it, then click the up-arrow or down-arrow buttons to move the resource up or down in the list. The page refreshes between each click on an arrow.

- ♦ **Edit:** Open the Resource Priority Edit function by clicking the *Edit* pen icon, list the resources in the preferred load order with one resource per line, then click *OK* to accept the revised order.

Be careful to not alter the names of the resources while editing the list.



- 6 Click *Apply* or *OK* to save changes.

9.10 Configuring Resource Mutual Exclusion Groups

The Resource Mutual Exclusion (RME) Groups feature allows you to define sets of resources that must not run on the same node at the same time. You might need to use RME Groups for resources when an application does not want its resource to run concurrently on the same node as other applications. You can also use RME Groups if you run multiple instances of a service in the same cluster and they should run on different nodes. This Clusters plug-in feature is available only for OES 2 SP3 and later servers.

For example, because of the designs and supported configurations of Novell iPrint and GroupWise, the GroupWise and iPrint resources should not run on the same node at the same time. In addition, multiple iPrint server instances cannot be hosted on the same node. To avoid resource assignment conflicts, you could create an RME Group that consists of the GroupWise resource and the multiple iPrint resources.

- ♦ [Section 9.10.1, “Understanding Resource Mutual Exclusion,” on page 143](#)
- ♦ [Section 9.10.2, “Setting Up RME Groups,” on page 146](#)
- ♦ [Section 9.10.3, “Viewing RME Group Settings,” on page 146](#)

9.10.1 Understanding Resource Mutual Exclusion

Resource Mutual Exclusion is the practice of letting a node run only one of the resources in an RME Group at a time. The resources can run concurrently on different nodes.

- ♦ [“How RME Groups Work” on page 144](#)
- ♦ [“Rules for Node Assignment” on page 144](#)
- ♦ [“Examples” on page 144](#)
- ♦ [“Planning for RME Groups” on page 145](#)

How RME Groups Work

When RME Groups exist, Novell Cluster Services does the following:

- ♦ It does not allow resources that are members of the same RME Group to run on the same node at the same time.
- ♦ It evaluates the exclusions together, but separately enforces each group of exclusions.

Rules for Node Assignment

When a resource is brought online, Cluster Services honors the resource settings in the following order:

1. Resource Follows Master
2. Resource Mutual Exclusion Group
3. Preferred Nodes list

When onlining, migrating, or failing over a resource onto a node, Cluster Services evaluates RME settings to check if the resource is a member of any RME Groups, then assesses whether any one of the other resources in the RME Groups that the resource belongs to is already running there. If another resource from the same RME Group is running on the node, the situation is handled as if the node was not on the resource's Preferred Nodes list. The checking process repeats itself with another node as a candidate until a suitable node is found, or until the resource's Preferred Nodes list is exhausted.

Resources that are not members of any of the RME Groups are not restricted. They are allowed to run at any time and in any combination on any node in their Preferred Nodes lists.

Examples

If you try to cluster migrate a resource to a node where another member of its RME group is running, the resource is not migrated and remains in the Running state on the node where it is currently loaded. Only one resource in the group is allowed to run on the node at a time.

If you try to cluster migrate multiple resources in an RME Group to the same node, the resources are brought online on different nodes. Only one of the resources (randomly selected) is brought online on the specified node. The other resources are migrated to different nodes in their Preferred Nodes lists.

Planning for RME Groups

An RME Group can contain any combination of the resources that are available to the cluster. Resources that are members of the same group are not allowed to run concurrently on a node. A resource can be a member of more than one group.

You can define up to four groups (Group A, Group B, Group C, and Group D). The group names are fixed; they cannot be customized.

For example, the following image shows an RME Groups page where three groups (A, B, and D) have been defined. The selected check boxes in a column indicate which resources belong to that group. Group A has four member resources, and Group B and Group D have three members each. The AUTO_POOL_12_SERVER resource belongs to Group A and Group B. It has mutually exclusive relationships with two different sets of resources, and each exclusion is managed separately.

Figure 9-1 Sample RME Groups

Clusters > Cluster Options

Cluster Properties: oes2_sales_cluster.novell

Policies | Priorities | Protocols | **RME Groups** | Business Continuity

View or change Resource Mutual Exclusion (RME) group settings.

Name	Group A	Group B	Group C	Group D
AUTO_POOL_14_SERVER	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
AUTO_POOL_13_SERVER	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
AUTO_POOL_12_SERVER	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
AUTO_POOL_11_SERVER	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
AUTO_POOL_10_SERVER	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
AUTO_POOL_09_SERVER	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
AUTO_POOL_08_SERVER	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
AUTO_POOL_07_SERVER	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
AUTO_POOL_06_SERVER	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
AUTO_POOL_05_SERVER	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
AUTO_POOL_04_SERVER	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
AUTO_POOL_03_SERVER	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
AUTO_POOL_02_SERVER	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
AUTO_POOL_01_SERVER	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Master_IP_Address_Resource	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>

OK Cancel Apply

RME Groups are supported when all cluster nodes are running a supported platform. If you are managing an older version of Novell Cluster Services (without the RME code) from a new Clusters plug-in in iManager, the tables on the RME Groups page are shown as *No items*. You cannot set up RME Groups for clusters running on older platforms.

Figure 9-2 Sample RME Groups Page for an Unsupported Platform

Clusters > Cluster Options

Cluster Properties: oes2_engr_cluster.novell

Policies | Priorities | Protocols | **RME Groups** | Business Continuity

View or change resource mutual exclusion (RME) group settings.

Name	Group A	Group B	Group C	Group D
No items				

OK Cancel Apply

If you are upgrading from an OES 2 SP2 or NetWare cluster, the RMS Groups are not available. You should wait until the rolling cluster upgrade has been completed on all nodes before you define RME Groups. If you define RME Groups while the cluster is in a mixed-mode condition, the RME Groups are honored only if the resource fails over to a cluster node where the new OS is running.

9.10.2 Setting Up RME Groups

To define sets of resources that must not be assigned to the same node at the same time:

- 1 In iManager, select *Clusters > Cluster Options*.
- 2 Browse to locate and select the Cluster object of the cluster you want to manage.
- 3 Click *Properties*, then click the *RME Groups* tab.

The RME Groups page lists the resources in the cluster and columns for four possible RME Groups for the cluster. In each Group's column, the check boxes correspond to the resources listed at the left.

If you are managing an older version of Novell Cluster Services (without the RME code) from a new Clusters plug-in in iManager, the tables are shown as *No Items*. You cannot set up RME Groups for clusters running on older platforms.

- 4 To set up the members in a group, select the check boxes in the same column for two or more resources that must not be assigned on the same node at the same time.

The selected resources in a column are not allowed to run concurrently on a node. A resource can be a member of more than one group. You can define up to four groups (A, B, C, and D).

- 5 Click *OK* or *Apply* to save your changes.

9.10.3 Viewing RME Group Settings

The RME Groups settings are displayed in the following locations:

- ♦ [“Cluster Report” on page 146](#)
- ♦ [“RME Groups Configuration Page” on page 147](#)

Cluster Report

The Cluster Report has an RME Groups section that lists the member resources of each group.

- 1 In iManager, select *Clusters > Cluster Manager*.
- 2 Browse to locate and select the Cluster object of the cluster you want to manage.
- 3 Click *Run Report*, then scroll down to view the *RME Groups*.

The RME Groups section lists only the resources in the cluster that are members in each of the RME Groups.

For example, the following report shows the members for Groups A, B, and D. Group C has no defined members. The `AUTO_POOL_12_SERVER` resource belongs to two RME Groups. It has mutually exclusive relationships with two different groups of resources. These exclusions are managed separately.

Resource Mutual Exclusion (RME) Groups			
Group A	Group B	Group C	Group D
AUTO_POOL_12_SERVER	AUTO_POOL_14_SERVER		AUTO_POOL_10_SERVER
AUTO_POOL_09_SERVER	AUTO_POOL_12_SERVER		AUTO_POOL_07_SERVER
AUTO_POOL_06_SERVER	AUTO_POOL_05_SERVER		Master_IP_Address_Resource
AUTO_POOL_02_SERVER			

RME Groups Configuration Page

The RME Groups configuration page displays all possible resources, and indicates membership by the check boxes that are selected in the same column under the RME Group Name.

- 1 In iManager, select *Clusters > Cluster Options*.
- 2 Browse to locate and select the Cluster object of the cluster you want to manage.
- 3 Click *Properties*, then click the *RME Groups* tab.

The RME Groups configuration page lists the all of the resources in the cluster. The selected check boxes in the same group column indicate their memberships in the four possible RME Groups.

9.11 Controlling Resource Monitoring

You can use the `cluster monitor` command to start, stop, or view the status of monitoring for a specified cluster resource. Issue the command as the `root` user from the node where the resource is currently loaded:

```
cluster monitor <resourcename> {start | stop | status}
```

Monitoring must be enabled for the resource. For instructions, see [Section 9.6, “Enabling Monitoring and Configuring the Monitor Script,”](#) on page 135.

9.12 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. Ensure that you 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.

9.13 Renaming a Cluster Resource

A `cluster rename` option is available that allows you to rename a cluster resource. Renaming the resource does not modify the virtual server name (the NCS:NCP Server object).

Your setup must meet the following prerequisites:

- ♦ This command must be issued from the master node.
- ♦ The resource must be in the *Offline* state in order to be renamed.
- ♦ The new name must not exist prior to the renaming.
- ♦ Novell eDirectory must be running when you attempt to rename cluster resource.

To rename the cluster resource:

- 1 Log in to the master node in the cluster as the `root` user, and open a terminal console.
- 2 Offline the cluster resource that you want to rename by entering

```
cluster offline resource_name
```

For example:

```
cluster offline POOL1_SERVER
```

3 Rename the cluster resource by entering

```
cluster rename <resource_name> <new_resource_name>
```

For example:

```
cluster rename POOL1_SERVER customized_name22
```

4 Online the cluster resource by entering

```
cluster online new_resource_name
```

For example:

```
cluster online customized_name22
```

9.14 Deleting Cluster Resources

Ensure that you 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 9.14.1, “Deleting a Cluster Resource on a Master Node,” on page 148](#)
- ♦ [Section 9.14.2, “Deleting a Cluster Resource on a Non-Master Node,” on page 149](#)

9.14.1 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 9.14.2, “Deleting a Cluster Resource on a Non-Master Node,” on page 149](#).

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 LVM volume groups, use the YaST Partitioner.
- 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.

9.14.2 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 9.14.1, “Deleting a Cluster Resource on a Master Node,” on page 148](#).

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 LVM volume group, use the YaST Partitioner.
- 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 [seconds]
```

For *seconds*, specify a value of 60 seconds or more.

For example:

```
cluster restart 120
```

9.15 Additional Information for Creating Cluster Resources

Cluster resources can be configured for storage, services, or virtual machines.

- [Section 9.15.1, “Creating Storage Cluster Resources,” on page 150](#)
- [Section 9.15.2, “Creating Service Cluster Resources,” on page 150](#)
- [Section 9.15.3, “Creating Virtual Machine Cluster Resources,” on page 151](#)

9.15.1 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 11, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 155
Linux POSIX volumes	Chapter 13, “Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers,” on page 229
NCP volumes	“Configuring NCP Volumes with Novell Cluster Services” in the <i>OES 11: 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” in the <i>OES 11: Dynamic Storage Technology Administration Guide</i>

9.15.2 Creating Service Cluster Resources

For information about creating cluster resources for various services on your OES 11 server, see [Chapter 10, “Creating and Managing Service Cluster Resources,” on page 153](#).

9.15.3 Creating Virtual Machine Cluster Resources

If you install Novell Cluster Services at the host level of an OES 11 (Xen) server, you can create cluster resources for the virtual machines. For information, see [Section 14.2, “Virtual Machines as Cluster Resources,”](#) on page 250.

10 Creating and Managing Service Cluster Resources

The individual service guides contain information about creating and managing service cluster resources with Novell Cluster Services on your Novell Open Enterprise Server (OES) 11 server. See [Table 10-1](#) for links to those resources. You can also find this list by going to “[Clustering Linux Services](#)” on the [Clustering \(High Availability\) Documentation Web site](#) (<http://www.novell.com/documentation/oes11/cluster-services.html#clust-config-resources>).

Table 10-1 *Clustering OES 11 Services with Novell Cluster Services*

OES 11 Service	For information, see
Apache Web Server	“ Apache Web Server ” in the OES 11: Novell Cluster Services NetWare to Linux Conversion Guide
Archive and Version Services	“ Configuring Archive and Version Service for Novell Cluster Services ” in the OES 11: Novell Archive and Version Services 2.1 Administration Guide
Certificate Server	<p>The eDirectory Certificate Server is not cluster-enabled. The Certificate Server service issues Server Certificate objects that might need to reside on each node in a cluster, depending on the service that is clustered.</p> <p>See “eDirectory Server Certificates” in the OES 11: Novell Cluster Services NetWare to Linux Conversion Guide.</p>
DFS VLDB (Distributed File Services volume location database)	“ Clustering Novell Distributed File Services ” in the OES 11 SP1: Novell Distributed File Services Administration Guide for Linux
DHCP Server	<p>In the OES 11: Novell DNS/DHCP Services for Linux Administration Guide, see:</p> <ul style="list-style-type: none">◆ “Configuring DHCP with Novell Cluster Services for the NSS File System”◆ “Configuring DHCP with Novell Cluster Services for the Linux File System”
DNS Server	“ Configuring DNS with Novell Cluster Services ” in the OES 11: Novell DNS/DHCP Services for Linux Administration Guide
eDirectory	Novell eDirectory is not clustered because it has its own replica system.
File, AFP (Apple Filing Protocol)	“ Configuring AFP with Novell Cluster Services for an NSS File System ” in the OES11: Novell AFP Administration Guide
File, CIFS (Windows File Services)	“ Configuring CIFS with Novell Cluster Services for an NSS File System ” in the OES 11: Novell CIFS for Linux Administration Guide

OES 11 Service	For information, see
File, FTP	“Cluster Enabling Pure-FTPd in an OES 11 Environment” in the <i>OES 11: Planning and Implementation Guide</i>.
File, NetStorage	“Configuring NetStorage with Novell Cluster Services” in the <i>OES 11: NetStorage Administration Guide for Linux</i>.
File, Samba	“Configuring Samba for LVM Volume Groups and Novell Cluster Services” in the <i>OES 11: Novell Samba Administration Guide</i>
iFolder 3.9	“Clustering iFolder Servers with Novell Cluster Services for Linux” in the <i>Novell iFolder 3.9 Administration Guide</i>
iPrint	“Configuring iPrint with Novell Cluster Services” in the <i>OES 11: iPrint Linux Administration Guide</i>
MySQL	A MySQL template is available that uses a shared LVM volume group that you have already created.
QuickFinder (Server Synchronization Feature)	“Configuring QuickFinder Server for Novell Cluster Services” in the <i>OES 11: Novell QuickFinder Server 5.0 Administration Guide</i>
Storage, DST shadow volumes built with NSS volumes	In the OES 11: Dynamic Storage Technology Administration Guide , see: <ul style="list-style-type: none"> ◆ “Configuring DST Shadow Volumes with Novell Cluster Services”
Storage, Linux POSIX volumes	Chapter 13, “Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers,” on page 229
Storage, NCP volumes	“Configuring NCP Volumes with Novell Cluster Services” in the <i>OES 11: NCP Server for Linux Administration Guide</i>
Storage, NSS pools and volumes	Chapter 11, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 155
Xen virtual machines	Section 14.2, “Virtual Machines as Cluster Resources,” on page 250

11 Configuring Cluster Resources for Shared NSS Pools and Volumes

After you have installed and configured Novell Cluster Services on a server, you can configure Novell Storage Services (NSS) pool cluster resources. This section describes how to create and cluster-enable shared NSS pools volumes as pool cluster resources with Novell Cluster Services.

- ♦ [Section 11.1, “Requirements for Creating Pool Cluster Resources,” on page 155](#)
- ♦ [Section 11.2, “Guidelines for Using Pool Cluster Resources,” on page 159](#)
- ♦ [Section 11.3, “Creating Cluster-Enabled Pools and Volumes,” on page 161](#)
- ♦ [Section 11.4, “Cluster-Enabling an Existing NSS Pool and Its Volumes,” on page 167](#)
- ♦ [Section 11.5, “Configuring a Load Script for the Shared NSS Pool,” on page 171](#)
- ♦ [Section 11.6, “Configuring an Unload Script for the Shared NSS Pool,” on page 173](#)
- ♦ [Section 11.7, “Configuring a Monitor Script for the Shared NSS Pool,” on page 174](#)
- ♦ [Section 11.8, “Adding Advertising Protocols for NSS Pool Cluster Resources,” on page 174](#)
- ♦ [Section 11.9, “Adding NFS Export for a Clustered Pool Resource,” on page 175](#)
- ♦ [Section 11.10, “Mirroring and Cluster-Enabling Shared NSS Pools and Volumes,” on page 176](#)
- ♦ [Section 11.11, “Renaming a Pool for a Pool Cluster Resource,” on page 180](#)
- ♦ [Section 11.12, “Adding a Volume to a Clustered Pool,” on page 181](#)
- ♦ [Section 11.13, “Changing an Assigned Volume ID,” on page 181](#)
- ♦ [Section 11.14, “Expanding the Size of a Clustered Pool,” on page 182](#)
- ♦ [Section 11.15, “Deleting NSS Pool Cluster Resources,” on page 194](#)

11.1 Requirements for Creating Pool Cluster Resources

Your system must meet the requirements in this section in addition to the cluster requirements described in [Chapter 4, “Planning for Novell Cluster Services,” on page 37](#).

- ♦ [Section 11.1.1, “Novell Cluster Services,” on page 156](#)
- ♦ [Section 11.1.2, “Resource IP Address,” on page 156](#)
- ♦ [Section 11.1.3, “Novell eDirectory,” on page 156](#)
- ♦ [Section 11.1.4, “Novell Storage Services,” on page 156](#)
- ♦ [Section 11.1.5, “Shared Storage,” on page 157](#)
- ♦ [Section 11.1.6, “NCP Server for Linux,” on page 158](#)
- ♦ [Section 11.1.7, “Novell CIFS for Linux,” on page 158](#)
- ♦ [Section 11.1.8, “Novell AFP for Linux,” on page 159](#)

- ♦ [Section 11.1.9, “Novell Samba,” on page 159](#)
- ♦ [Section 11.1.10, “Domain Services for Windows,” on page 159](#)

11.1.1 Novell Cluster Services

Novell Cluster Services must be installed, configured, and running when you create and manage the shared NSS pools and volumes. The cluster must be active.

11.1.2 Resource IP Address

Each NSS pool cluster resource requires a unique static IP address. The IP address is used to provide access to and failover capability for the pool cluster resource. Users access the pool by using the resource IP address instead of the server IP address where the pool is active. The IP address you assign to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.

IMPORTANT: The IP address for the virtual server must be in the same IP subnet as the server nodes in the cluster where you plan to use it.

11.1.3 Novell eDirectory

Novell eDirectory must be running and working properly when you cluster-enable a pool and when you manage the pool.

11.1.4 Novell Storage Services

NSS must be installed and running on each server in the cluster. For information about installing NSS and managing NSS pools and volumes, see the [OES 11: NSS File System Administration Guide for Linux](#). For information about using NLVM commands, see [OES 11: NLVM Reference](#).

In addition, the following requirements must be met:

- ♦ [“Pool” on page 156](#)
- ♦ [“Volume” on page 157](#)
- ♦ [“Naming Conventions” on page 157](#)
- ♦ [“Pool Cluster Resource Name” on page 157](#)
- ♦ [“Location of Cluster, Server, and Storage Objects” on page 157](#)

Pool

We recommend that your pool cluster resource be one device, one pool, one volume. You can create a pool up to 8 TB in size on a single device by formatting the device in GPT format.

You can create and cluster-enable a shared NSS pool by using the Storage plug-in for Novell iManager, the server-based NSS Management Utility (nssmu), or the Novell Linux Volume Manager (NLVM) command line interface. You can also use these tools to create NSS volumes on the shared pool. These tools interact with Novell Cluster Services to create the pool cluster resource.

You can cluster-enable an existing NSS pool by using the Clusters plug-in for Novell iManager.

Volume

You must create at least one volume in the shared pool. Typically, you create all volumes for a shared pool when you set up the pool cluster resource and before you need to cluster migrate or fail over the resource to a different node in the cluster.

Naming Conventions

Pool names and volume names must comply with the same naming conventions. Names must be 2 to 15 characters in length. Novell Cluster Services supports pool names with characters A to Z, 0 to 9, and underscores (_). The name cannot begin or end with an underscore; it cannot contain multiple adjacent underscores (__). It does not allow the following special characters in a shared pool name:

!@#\$\$%& ()

IMPORTANT: Before you cluster-enable an existing NSS pool, you must rename the pool and its volumes to use names that do not contain special characters.

Pool Cluster Resource Name

The pool name is used to form the default name of the pool cluster resource and the virtual server (the NCS:NCP Server object). If you modify the resource name, ensure that the name conforms to the resource naming conventions as described in [Section 9.1.1, “Naming Conventions for Cluster Resources,” on page 128](#).

Storage Object	Default Name
Resource name	<cluster_name>_<poolname>
Clustered Volume object name	<cluster_name>_<volume_name> \\<cluster_name>-<poolname>-SERVER\<volume_name>
Virtual server name	<cluster_name>-<poolname>-SERVER

Location of Cluster, Server, and Storage Objects

The Server, Pool, Volume, Cluster Resource, and Cluster objects are recommended to be in the same context (such as ou=ncs,o=novell). It is supported for the objects to be in the same context or different contexts. If the objects are in different contexts, you might need to cluster migrate the pool cluster resource back to the node where the pool was created in order to modify the pool or volume, or to perform other tasks like setting up Distributed File Services junctions or home directories. You receive an eDirectory error if the operation cannot find the information that it needs in the same context.

11.1.5 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 4.8, “Shared Disk Configuration Requirements,” on page 53](#)

- ♦ [Section 4.9, “SAN Rules for LUN Masking,” on page 55](#)
- ♦ [Section 4.10, “Multipath I/O Configuration Requirements,” on page 56](#)

11.1.6 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. This is necessary to provide authenticated access to data using the Novell Trustee model.

Novell Storage Services requires that the NCP Server for Linux service be installed and running on each node in the server. NCP Server must be running even if users access volumes on the shared NSS pool only via other protocols.

WARNING: Cross-protocol file locking is required when using multiple protocols for data access on the same volume. This helps prevent possible data corruption that might occur from cross-protocol access to files.

The NCP Cross-Protocol File Lock parameter is enabled by default when you install NCP Server. If you modify the Cross-Protocol File Lock parameter, you must modify the setting on all nodes in the cluster.

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

For information, see “[Configuring Cross-Protocol File Locks for NCP Server](#)” in the *OES 11: 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 11: NCP Server for Linux Administration Guide*.

11.1.7 Novell CIFS for Linux

Common Internet File System (CIFS) is the Windows networking protocol. Novell CIFS allows you to give clients access via CIFS to volumes on the shared NSS pool.

WARNING: To prevent possible data corruption, enable the NCP Cross-Protocol File Locks parameter for NCP Server on all nodes in the cluster before you allow users to access the data. For information, see “[Configuring Cross-Protocol File Locks for NCP Server](#)” in the *OES 11: NCP Server for Linux Administration Guide*.

Novell CIFS must be installed, configured, and working properly before you can specify CIFS as an advertising protocol when you cluster-enable 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 11: Novell CIFS for Linux Administration Guide*.

11.1.8 Novell AFP for Linux

Apple Filing Protocol (AFP) is the Macintosh networking protocol. 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, enable the NCP Cross-Protocol File Locks parameter for NCP Server on all nodes in the cluster before you allow users to access the data. For information, see [“Configuring Cross-Protocol File Locks for NCP Server”](#) in the *OES 11: NCP Server for Linux Administration Guide*.

Novell AFP must be installed, configured, and working properly before you can specify AFP as an advertising protocol when you cluster-enable 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 *OES11: Novell AFP Administration Guide*.

11.1.9 Novell Samba

Novell Samba is supported for NSS pool cluster resources as an alternative to Novell CIFS. The setup is not integrated in the Advertising Protocols options when you cluster-enable a pool. If you do not use Novell CIFS, Novell Samba can be set up after you cluster-enable the pool. For information about setting up Novell Samba for storage resources, see [“Configuring Samba for LVM Volume Groups and Novell Cluster Services”](#) in the *OES 11: Novell Samba Administration Guide*.

Novell Samba requires that users are enabled for Linux with Linux User Management (LUM). For information, see *OES 11: Novell Linux User Management Administration Guide*.

WARNING: To prevent possible data corruption, enable the NCP Cross-Protocol File Locks parameter for NCP Server on all nodes in the cluster before you allow users to access the data. For information, see [“Configuring Cross-Protocol File Locks for NCP Server”](#) in the *OES 11: NCP Server for Linux Administration Guide*.

11.1.10 Domain Services for Windows

Cluster-enabled NSS volumes can be used in a Domain Services for Windows environment.

11.2 Guidelines for Using Pool Cluster Resources

Consider the guidelines in this section when working with shared NSS pools and volumes in the cluster:

- ♦ [Section 11.2.1, “Guidelines for Shared Pools,”](#) on page 159
- ♦ [Section 11.2.2, “Guidelines for Shared Volumes,”](#) on page 160

11.2.1 Guidelines for Shared Pools

- ♦ When the pool cluster resource is brought online, the pool is automatically activated by the resource load script. You typically do not activate the pool at the terminal console.
- ♦ The unit of failover is the device (disk or LUN) that contains the pool. If you use multiple devices in the pool, all of those devices must be failed over together.

IMPORTANT: As a best practice, we recommend a single LUN (or device) per shared pool, and one volume per pool. (one LUN=>one pool=>one volume)

- ♦ If you delete a cluster-enabled pool, Novell Cluster Services automatically removes the Pool Resource object and the virtual server object from Novell eDirectory.

Ensure that you offline the cluster resource before you attempt to delete the cluster resource or the pool. For information, see [Section 9.14, “Deleting Cluster Resources,” on page 148](#).

- ♦ Before you rename a cluster-enabled pool, ensure that you offline the pool resource, activate the pool by using iManager or NSSMU, rename the pool, deactivate the pool, then online the pool resource.

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.

11.2.2 Guidelines for Shared Volumes

You must create at least one volume on a shared pool before you migrate it. As a best practice, we recommend that you create only one volume per shared pool.

When you create a volume, commands are added to the pool resource load and unload scripts to automatically mount and dismount the volume when the scripts run. You can modify the load script to comment out the mount command so that you can manually mount the volume on a node in the cluster where the pool resource has been activated.

Volumes on cluster-enabled pools do not appear as independent cluster resources. A cluster-enabled volume does not have an associated load and unload script or an assigned IP address. If you want each volume to be in a separate cluster resource, each volume must have its own pool.

When a server fails, the cluster resources fail over to other servers in the cluster. Because the cluster-enabled pool fails over, all volumes in the pool also fail over, but only the volumes that have been added to the load script are automatically mounted. Any volumes in the pool that are not in the load script must be mounted manually. For this reason, volumes that you do not want to fail over should be in separate pools that are not cluster-enabled.

When you create an encrypted NSS volume in a shared pool, you must mount the volume manually by using NSSMU and enter the password. NSS uses the password to create a key. Instead of storing it in the server memory as it does for non-shared volumes, NSS asks 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.

If you delete a volume from a cluster-enabled pool, Novell Cluster Services automatically removes the volume mount command from the pool cluster resource load script.

11.3 Creating Cluster-Enabled Pools and Volumes

After you have installed and configured Novell Cluster Services, you can create a pool cluster resource by creating a cluster-enabled NSS pool and volume.

- ♦ [Section 11.3.1, “Initializing and Sharing a Device,” on page 161](#)
- ♦ [Section 11.3.2, “Creating a Cluster-Enabled Pool and Volume with iManager,” on page 163](#)
- ♦ [Section 11.3.3, “Creating a Cluster-Enabled Pool and Volume with NSSMU,” on page 165](#)

11.3.1 Initializing and Sharing a Device

Devices (disks or LUNs) that you want to use for your cluster-enabled pool must be initialized in order to be visible to the NSS management tools.

In order to be able to cluster-enable the pool as you create it, each of the devices that you want to use for the pool must be marked as shareable for clustering. Novell Cluster Services must be installed and running on the server in order to share the device.

You can perform the device management tasks by using the Novell Storage Services (NSS) Management Utility or the Storage plug-in to iManager.

- ♦ [“Initializing and Sharing a Disk with iManager” on page 161](#)
- ♦ [“Initializing and Sharing a Disk with NSSMU” on page 162](#)

Initializing and Sharing a Disk with iManager

- 1 In iManager, select *Storage > Devices*, then select the server where the storage is currently assigned in the cluster.
- 2 From the *Devices* list, select the device where you want to create the pool.
- 3 If the device has not been previously initialized or if you want to delete the current partitioning structures on the device, click *Initialize Disk*, then click *OK* to confirm and continue.

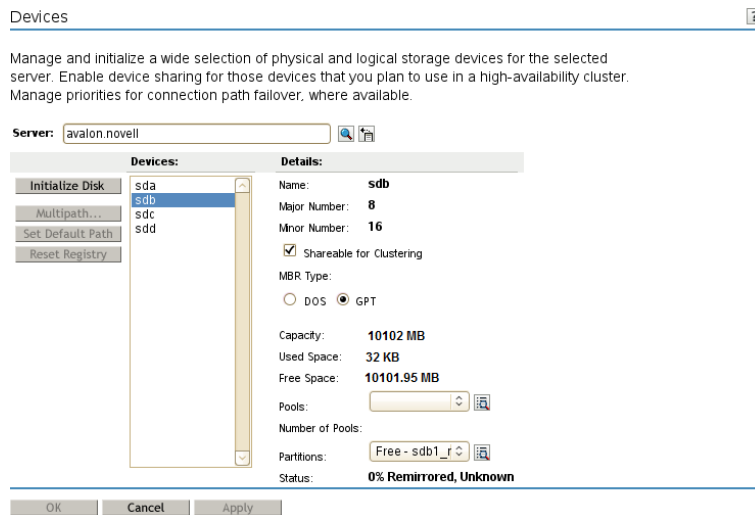
WARNING: Initializing a disk destroys all of the data on it.

- 4 If you are prompted, specify whether to use the GPT or DOS partitioning scheme, then click *OK* to complete the initialization.

GPT can be used for any size device. DOS supports devices up to 2 TB in size.

Wait for the page to refresh before continuing.

- 5 In the *Details* area, select the *Shareable for Clustering* check box, then click *OK* to apply the change.



- 6 Repeat [Step 2](#) through [Step 5](#) for each device that you want to use in the pool.
- 7 Exit iManager.
- 8 Continue with [Section 11.3.2, “Creating a Cluster-Enabled Pool and Volume with iManager,”](#) on [page 163](#).

Initializing and Sharing a Disk with NSSMU

- 1 Log in as the `root` user to the server in the cluster where the storage is currently assigned, then open a terminal console.
- 2 At the console prompt, enter


```
nssmu
```
- 3 In the NSSMU, select *Devices* and press Enter.
- 4 In the *Devices* list, select the device where you want to create the pool.
- 5 If the device has not been initialized or if you want to delete the current partitioning structures on the device, press F3 to initialize the selected device, then press Y (Yes) to confirm and continue.

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

- 6 Specify whether to use the GPT or DOS partitioning scheme, then click *OK* to complete the initialization.

GPT can be used for any size device. DOS supports devices up to 2 TB in size.

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

The *Shareable for Clustering* value changes from *No* to *Yes*.

NFS Management Utility	
NFS Shared Pools, RAID 5/10	
Devices	Device Information
sda	Name: sdc
sdb	ID: 8:32
* sdc	Partitioning type: GPT
sdd	RAID: No
sde	LVM: No
	Capacity: 10102.00 MB
	Partitioned Space: 32 KB
	Unpartitioned Space: 10101.95 MB
	Sharable for Clustering: Yes
F3=Init F5=Refresh F6=UnShare SPACE=Select/Unselect F8=More	

- 8 Repeat [Step 4](#) through [Step 7](#) for each device that you want to use in the pool.
- 9 Exit NSSMU.
- 10 Continue with [Section 11.3.3, “Creating a Cluster-Enabled Pool and Volume with NSSMU,”](#) on [page 165](#).

11.3.2 Creating a Cluster-Enabled Pool and Volume with iManager

- 1 Log in to iManager as an administrator user.
- 2 In *Roles and Tasks*, select *Storage > Pools*.
- 3 Browse to select the Cluster object (🔴🔴🔴) of the cluster.
This automatically selects the server that is currently the master node in the cluster. The shared device will be assigned to this server and the pool will be created there.
- 4 Click the *New* link to open the New Pool wizard.
- 5 Specify the new pool name, then click *Next*.
- 6 Select the check box next to the shared device where you want to create the pool, then specify the size of the pool. We recommend one that you use one device per pool.
- 7 Deselect the *Mount On Creation* option.
The *Mount On Creation* option determines if the pool you are creating is to be activated as soon as it is created. For cluster-enabled pools, activation is done by bringing the pool cluster resource online. Deselect the option so that you can configure the resource and modify the scripts before you bring the resource online for the first time.
- 8 Select the *Cluster Enable on Creation* check box.
This option is selected by default if the device is shared. You can deselect this option and cluster-enable the pool later as described in [Section 11.4, “Cluster-Enabling an Existing NSS Pool and Its Volumes,”](#) on [page 167](#).

9 On the Cluster Pool Information page, specify the following information:

Parameter	Action
Virtual Server Name	<p>(Optional) Change the default name of the Virtual Server for the cluster resource.</p> <p>The default virtual server name for the resource is the cluster name plus the cluster resource name. For example, if the cluster name is <code>cluster1</code> and the pool cluster resource name is <code>POOL1_SERVER</code>, then the default virtual server name is <code>CLUSTER1-POOL1-SERVER</code>.</p>
CIFS Server Name	<p>If CIFS is enabled as an advertising protocol, specify the name of the CIFS virtual server that CIFS clients see when they browse the network. The name can be up to 15 characters.</p> <p>A default name is specified in the form of <code>clustername_poolname_w</code>. If this exceeds 15 characters, characters are dropped from the left.</p> <p>If Novell CIFS is not installed and running, this field value is <code>NOT_SUPPORTED</code>.</p>
IP Address	<p>Specify an IP address for the pool cluster resource. Tab between the address fields.</p> <p>Each pool cluster resource requires its own unique IP address. The IP address assigned to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.</p>
Advertising Protocols	<p>Select the check boxes of the advertising protocols (AFP, CIFS, NCP) that you want to enable for data requests to this shared pool. NCP is required to support authenticated access to data via the Novell Trustee model.</p> <p>Selecting a protocol causes commands to be added to the pool cluster resource's load and unload scripts to activate the protocol for the resource. This lets you ensure that the cluster-enabled pool is highly available to users via the specified protocol.</p> <p>If Novell CIFS or Novell AFP are not installed and running, selecting the CIFS or AFP check box has no effect.</p>
Online Resource after Create	<p>The check box is deselected by default and dimmed so that you cannot change the setting.</p> <p>The pool is currently active on the server. You must deactivate the pool from the server before attempting to bring the resource online. You should also configure the resource load, unload, and monitoring scripts before you bring the resource online.</p>
Define Additional Properties	<p>Select the <i>Define Additional Properties</i> check box.</p> <p>This allows you to configure the resource policies for the start, failover, and failback modes, and to configure the preferred nodes.</p>

10 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.

- 11 Repeat the following procedure for each cluster volume that you want to create on the shared pool. We recommend using only one volume per shared pool.
 - 11a In iManager, select *Storage*, then select the *Volumes*.
 - 11b Click *New*.
 - 11c Specify the new volume name, then click *Next*.
Each shared volume in the cluster must have a unique name across all nodes.
 - 11d Select the check box next to the cluster pool where you want to create the volume, select *Allow the volume to grow to the size of the pool*, then click *Next*.
 - 11e 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 11: NSS File System Administration Guide for Linux*.
 - 11f Choose whether you want the volume activated and mounted when it is created, then click *Finish*.
- 12 Verify that the pool cluster resource was created and is online.
 - 12a In iManager, select *Clusters > Cluster Options*.
 - 12b Browse to select the Cluster.
 - 12c In the Cluster Objects list, view the pool cluster resource, such as POOL1_SERVER.
 - 12d In iManager, select *Clusters > Cluster Manager*.
 - 12e If the resource is online, the state is *Running*.
- 13 Continue with [Section 11.5, “Configuring a Load Script for the Shared NSS Pool,”](#) on page 171.

11.3.3 Creating a Cluster-Enabled Pool and Volume with NSSMU

- 1 On the master node, log in as the `root` user, then open a terminal console.
- 2 Start NSSMU by entering `nssmu` at the terminal console prompt.
- 3 From the NSSMU main menu, select *Pools*.
- 4 On the Pools page, press `Insert`, type a name for the new pool you want to create, then press `Enter`.
- 5 From the list of available devices, select the shared device where you want the pool created (such as `sdC`), then press `Enter`.
- 6 Specify the amount of space (in MB) to use, then press `Enter`.
- 7 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.

- 8 Specify *Yes* for the *Cluster Enable on Creation* option.

This option is selected by default if the device is shared. You can specify *No* for this option and cluster-enable the pool later as described in [Section 11.4, “Cluster-Enabling an Existing NSS Pool and Its Volumes,”](#) on page 167.

- 9 On the Cluster Pool Information page, specify the following information:

Parameter	Action
Virtual Server Name	(Optional) Change the default name of the Virtual Server for the cluster resource. The default virtual server name for the resource is the cluster name plus the cluster resource name. For example, if the cluster name is <code>cluster1</code> and the pool cluster resource name is <code>POOL1_SERVER</code> , then the default virtual server name is <code>CLUSTER1-POOL1-SERVER</code> .
CIFS Server Name	If CIFS is enabled as an advertising protocol, specify the name of the CIFS virtual server that CIFS clients see when they browse the network. The name can be up to 15 characters. A default name is specified in the form of <code>clustername_poolname_w</code> . If this exceeds 15 characters, characters are dropped from the left. If Novell CIFS is not installed and running, this field value is <code>NOT_SUPPORTED</code> .
IP Address	Specify an IP address for the pool cluster resource. Tab between the address fields. Each pool cluster resource requires its own unique IP address. The IP address assigned to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.
Advertising Protocols	Select the check boxes of the advertising protocols (AFP, CIFS, NCP) that you want to enable for data requests to this shared pool. NCP is required to support authenticated access to data via the Novell Trustee model. Selecting a protocol causes commands to be added to the pool cluster resource's load and unload scripts to activate the protocol for the resource. This lets you ensure that the cluster-enabled pool is highly available to users via the specified protocol. If Novell CIFS or Novell AFP are not installed and running, selecting the CIFS or AFP check box has no effect.

The default name of a pool cluster resource is the pool name plus the word “SERVER”, such as `POOL1_SERVER`. You can modify the resource name after the resource has been created by using the `cluster rename` command. For information, see [Section 9.13, “Renaming a Cluster Resource,”](#) on page 147. Changing the resource name does not modify the pool name or the virtual server name.

- 10 Select *Apply*, then press Enter to create and cluster-enable the pool.
- 11 Repeat the following procedure for each cluster volume that you want to create on the shared NSS pool. We recommend using only one volume per shared pool.
- 11a On the NSSMU main menu, select *Volumes*.
- 11b On the Volumes page, press Insert, type a name for the new volume you want to create, then press Enter.
- Each shared volume in the cluster must have a unique name across all nodes.

- 11c Specify Y(es) to encrypt volume or N(o) to create a regular volume.
- 11d 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.
- 11e (Optional) From the list of volumes, select the newly created volume, press F8 to view more options, press Enter to open the *Properties* page to review and change volume attributes, then select *Apply* and press Enter to save any changes.
The Backup and Salvage Files attributes are selected by default.
For information about volume attributes, see “[Volume Attributes](#)” in the *OES 11: NSS File System Administration Guide for Linux*.
- 11f Exit NSSMU.
- 12 Verify that the pool cluster resource was created.
 - 12a In iManager, select *Clusters > Cluster Options*.
 - 12b Browse to select the Cluster.
 - 12c In the Cluster Objects list, view the pool cluster resource, such as POOL1_SERVER.
 - 12d In iManager, select *Clusters > Cluster Manager*.
 - 12e If the resource is online, the state is *Running*.
- 13 Continue with [Section 11.5, “Configuring a Load Script for the Shared NSS Pool,”](#) on page 171.

11.4 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. This makes the data highly available to users.

IMPORTANT: You can cluster-enable pools when you create them by following the procedures in [Section 11.3, “Creating Cluster-Enabled Pools and Volumes,”](#) on page 161.

The procedure in this section describes how to enable clustering for an existing pool and its volumes. It assumes the following:

- ♦ The pool contains at least one volume.
- ♦ The pool name meets the naming conventions supported by Novell Cluster Services. For information, see “[Naming Conventions](#)” on page 157.

If the pool name contains special characters, rename the pool by using NSS management tools. For information, see “[Renaming a Pool](#)” in the *OES 11: NSS File System Administration Guide for Linux*.

- ♦ The pool and its volumes have Storage objects in the eDirectory tree where you are setting up the cluster.

If the 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 11: NSS File System Administration Guide for Linux*.

- ♦ The pool resides on a disk that has been enabled as *Sharable for Clustering*. If it is not shared, go to the Devices page of NSSMU, select the device, then press F6 to share the device.

- ♦ Each pool resource requires its own unique static IP address that is in the same IP subnet as the cluster IP address.
- ♦ To add CIFS as an advertising protocol when you cluster-enable an existing pool, you must install and configure Novell CIFS on the server, and ensure that the CIFS daemon is running and functioning properly before you begin. The CIFS daemon should also be running before trying to online the resource.
- ♦ To add AFP as an advertising protocol when you cluster-enable an existing pool, you must install and configure Novell AFP on the server, and ensure that the AFP daemon is running and functioning properly before you begin. The AFP daemon should also be running before trying to online the resource.

Before you cluster-enable the pool, activate it on the master node in the cluster. This allows the pool and volume information to be automatically added to the load, unload, and monitor scripts for the pool cluster resource. In addition, the volume entry is automatically removed from the `/etc/fstab` so that the pool cluster resource load script controls when the volume is mounted. If you use special settings for mounting the volume, you must manually add those settings to the load script.

The default name of a pool cluster resource is the pool name plus the word “SERVER”, such as `POOL1-SERVER`. You can modify the resource name after the resource has been created by using the `cluster rename` command. For information, see [Section 9.13, “Renaming a Cluster Resource,” on page 147](#). Changing the resource name does not modify the pool name or the virtual server name.

After the pool is cluster-enabled, users can access the cluster-enabled pool by using the pool cluster resource’s IP address or virtual server name. The default virtual server name of a pool cluster resource is the cluster name plus the default cluster resource name, such as `CLUSTER1-POOL1-SERVER`. You can specify the name of the resource’s virtual server name when you cluster-enable the pool.

To cluster-enable an existing NSS pool:

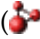
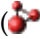

- 1** Log in to iManager as an administrator user.
- 2** Verify that the device is marked as shareable for clustering.
 - 2a** In *Roles and Tasks*, select *Storage > Devices*.
 - 2b** Browse to select the server where the pool is currently assigned.
 - 2c** Select the device to view its Details page.
 - 2d** If the *Shareable for Clustering* check box is deselected, select the check box, click *Apply*, read the warning, then click *Continue* to enable sharing.


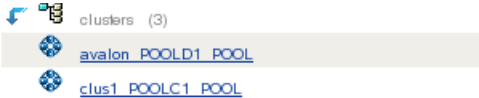
Marking the device as shareable automatically modifies the Shared attribute on the eDirectory objects for the pools and volumes on the device.

Wait a couple of minutes before continuing to allow the attribute change to synchronize in eDirectory.

- 3** Deactivate the shared pool on the current server.
 - 3a** In *Roles and Tasks*, select *Storage > Pools*.
 - 3b** Browse to select the server where the pool is currently assigned.
 - 3c** In the Pools list, select the pool, then click *Deactivate*.

Wait until the page refreshes and confirms that the pool is deactive.

- 4 Activate the pool on the master node in the cluster.
 - 4a In *Roles and Tasks*, select *Storage > Pools*.
 - 4b In the Server field, browse to select the Cluster object () of the cluster.
The Cluster object is automatically associated with the node that is currently the master node in the cluster.
 - 4c In the Pools list, select the pool, then click *Activate*.
Wait until the page refreshes and confirms that the pool is active. This automatically activates the pool and mounts its volumes on the server that is currently the master node in the cluster.
- 5 In *Roles and Tasks*, select *Clusters > Cluster Options*, then browse to select the Cluster object () of the cluster.
- 6 On the Cluster Options page, click the *New* link in the Cluster Objects toolbar.
- 7 On the Resource Type page, specify Pool () as the resource type you want to create by clicking the *Pool* radio button, then click *Next*.
- 8 On the Cluster Pool Information page, specify the following information:

Parameter	Action
Pool Name	<p>Specify the <i>Pool Name</i> by browsing to select the Pool object () of the pool you want to cluster-enable, such as <code>servername_POOLC1_POOL</code>.</p> <p>After the pool is cluster enabled, the Pool object name is changed to <code>clus1_POOLC1_POOL</code>, such as <code>clus1_POOLC1_POOL</code>.</p> 
Virtual Server Name	<p>(Optional) Change the default name of the Virtual Server for the cluster resource.</p> <p>The default virtual server name for the resource is the cluster name plus the cluster resource name. The default name of a pool cluster resource is the pool name plus the word “SERVER”, such as <code>POOL1_SERVER</code>.</p> <p>For example, if the cluster name is <code>cluster1</code> and the pool cluster resource name is <code>POOL1_SERVER</code>, then the default virtual server name is <code>CLUSTER1-POOL1-SERVER</code>.</p>
CIFS Server Name	<p>If CIFS is enabled as an advertising protocol, specify the name of the CIFS virtual server that CIFS clients see when they browse the network. The name can be up to 15 characters.</p> <p>A default name is specified in the form of <code>clus1_POOLC1_POOL</code>. If this exceeds 15 characters, characters are dropped from the left.</p> <p>If Novell CIFS is not installed and running, this field value is <code>NOT_SUPPORTED</code>.</p>
IP Address	<p>Specify an IP address for the pool cluster resource.</p> <p>Each pool cluster resource requires its own unique IP address. The IP address assigned to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.</p>

Parameter	Action
Advertising Protocols	<p>Select the check boxes of the advertising protocols (AFP, CIFS, NCP) that you want to enable for data requests to this shared pool. NCP is required to support authenticated access to data via the Novell Trustee model.</p> <p>Selecting a protocol causes commands to be added to the pool cluster resource's load and unload scripts to activate the protocol for the resource. This lets you ensure that the cluster-enabled pool is highly available to users via the specified protocol.</p> <p>If Novell CIFS or Novell AFP are not installed and running, selecting the CIFS or AFP check box has no effect.</p>
Online Resource after Create	<p>The check box is deselected by default and dimmed so that you cannot change the setting.</p> <p>The pool is currently active on the master node. You must deactivate the pool from the server before attempting to bring the resource online. You should also configure the resource load, unload, and monitoring scripts before you bring the resource online.</p>
Define Additional Properties	<p>Select the <i>Define Additional Properties</i> check box.</p> <p>This allows you to configure the resource policies for the start, failover, and failback modes, and to configure the preferred nodes.</p>

9 Click *Next*.

10 On the Resource Policies page, configure the resource policies for the start, failover, and failback modes.

For information, see [“Configuring the Start, Failover, and Failback Modes for Cluster Resources” on page 139](#).

11 Click *Next*.

12 On the Resource Preferred Nodes page, assign the preferred nodes to use for the resource.

For information, see [“Configuring Preferred Nodes for a Resource” on page 140](#).

13 Click *Finish*.

The pool cluster resource appears in the *Cluster Objects* list on the Cluster Options page, such as POOL1_SERVER.

The load, unload, and monitoring scripts are automatically created with information about the pools and volumes. The volume entries are removed from the `/etc/fstab` file so that the resource load script can be used to control when the pools and volumes are mounted. The resource is offline.

14 Verify that the pool cluster resource was created and that it is offline.

14a Select *Clusters > Cluster Options*, then view the pool cluster resource in the Cluster Objects list, such as POOL1_SERVER.

14b Select *Clusters > Cluster Manager*, then view the pool cluster resource state as *Offline*.

- 15** Manually deactivate the pool on its currently assigned node.

You must deactivate the pool from the server before attempting to bring the pool cluster resource online.

15a In *Roles and Tasks*, select *Storage > Pools*.

15b In the Server field, browse to select the Cluster object (🔗) of the cluster.

The Cluster object is automatically associated with the node that is currently the master node in the cluster.

15c In the Pools list, select the pool, then click *Deactivate*.

Wait until the page refreshes and confirms that the pool is deactive.

- 16** Verify the cluster resource scripts for the pool cluster resource, and modify them if desired.

16a Select *Clusters > Cluster Options*, then select the cluster.

16b Click the pool cluster resource link to open its Properties page.

16c Configure the following resource settings:

Property	Action
Monitoring	(Optional) Select the <i>Monitoring</i> tab, select <i>Enable Resource Monitoring</i> , and configure monitoring as described in Section 9.6, “Enabling Monitoring and Configuring the Monitor Script,” on page 135.
Load script	Select the <i>Scripts</i> tab, select <i>Load Script</i> , then modify the script as described in Section 11.5, “Configuring a Load Script for the Shared NSS Pool,” on page 171.
Unload Script	Select the <i>Scripts</i> tab, select <i>Unload Script</i> , then modify the unload script as described in Section 11.6, “Configuring an Unload Script for the Shared NSS Pool,” on page 173.
Monitoring Script	Select the <i>Scripts</i> tab, select <i>Monitor Script</i> , then modify the unload script as described in Section 9.6, “Enabling Monitoring and Configuring the Monitor Script,” on page 135.

16d At the bottom of the page, click *Apply* or *Ok* to save your changes.

The changes are not applied until the next time the resource is brought online.

- 17** Select *Clusters > Cluster Manager*, select the check box next to the pool cluster resource, then click *Online* to bring the resource online.

The pool is activated and its volumes are mounted on the primary preferred node that is configured for the pool cluster resource.

11.5 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 9.4, “Configuring a Load Script for a Cluster Resource,”](#) on page 133.

IMPORTANT: Do not comment out commands that are automatically generated for parameters that define the cluster resource, such as the IP address, pool name, and volume name.

If you need to modify the IP address, administrator credentials, or other attributes of an existing resource, follow the procedure in [Section 8.11, “Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster,”](#) on page 108.

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/ncsfuns

exit_on_error add_secondary_ipaddress 10.10.10.43
exit_on_error nss /poolact=SHPOOL43
exit_on_error ncpcon mount SHVOL43=252

exit_on_error ncpcon bind --ncpservname=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 11: NSS File System Administration Guide for Linux*.

For example, to specify the Long name space, add the `/opt=ns=long` switch as follows:

```
ncpcon mount <VOLUMENAME>=<VOLUMEID> /opt=ns=long
```

For example, to specify the Unix name space, add the `/opt=ns=unix` switch as follows:

```
ncpcon mount <VOLUMENAME>=<VOLUMEID> /opt=ns=unix
```

If you enable Novell AFP as a advertising protocol on the pool, add the following command after the `bind` command:

```
exit_on_error cluster_afp.sh add CLUSTER-POOL1-SERVER 10.10.10.43
```

If you enable Novell CIFS as a advertising protocol on the pool, add the following command after the `bind` command:

```
exit_on_error novcifs --add
'--vserver=".cn=CLUSTER-POOL1-SERVER.o=novell.t=AVALON_TREE."'
--ip-addr=10.10.10.43
```

11.6 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 9.5, “Configuring an Unload Script for a Cluster Resource,”](#) on page 134.

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/ncsfncs

ignore_error ncpcon unbind --ncpservname=NCS1-SHPOOL43-SERVER --
ipaddress=10.10.10.43

ignore_error nss /pooldeact=SHPOOL43
ignore_error del_secondary_ipaddress 10.10.10.43

exit 0
```

By default, `ncpcon dismount` commands are not added to the unload script. The pool deactivation automatically dismounts the volumes on the pool. However, you can add `ncpcon dismount` commands for each volume before the pool deactivation line. For example:

```
ignore_error ncpcon dismount SHVOL43
ignore_error nss /pooldeact=SHPOOL43
```

Adding the `ncpcon dismount` command provides predictable and clean dismount behavior. It allows the volume dismount to occur and be logged independently of the pool deactivation, which can help with troubleshooting dismount issues. However, the `ncpcon dismount` commands are not automatically maintained and updated if you change a volume name or delete a volume. You must modify the unload script yourself.

If Novell AFP is an advertising protocol for the pool, add the following line above the unbind command line:

```
ignore_error cluster_afp.sh del CLUSTER-POOL1-SERVER 10.10.10.43
```

If Novell CIFS is an advertising protocol for the pool, add the following line above the unbind command:

```
ignore_error novcifs
--remove '--vserver=".cn=CLUSTER-POOL1-SERVER.o=novell.t=AVALON_TREE."'
--ip-addr=10.10.10.43
```

11.7 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 9.6.2, “Configuring Resource Monitoring,” on page 137](#).

The following is a sample monitoring script for an NSS pool cluster resource:

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

exit_on_error status_fs /dev/pool/POOL1 /opt/novell/nss/mnt/.pools/POOL1 nsspool

exit_on_error status_secondary_ipaddress 10.10.10.41
exit_on_error ncpcon volume VOL1

exit 0
```

You can add other monitoring commands to the script. For example, the CIFS `monitor` command helps to keep CIFS up and running. It checks to see if the `cifs` daemon is running and takes action if it is not running. If CIFS is running, it returns its status. If CIFS is dead or not running, it starts the `cifs` daemon and returns its status. If the restart fails, the resource’s failover settings are invoked, such as failing over the resource to its preferred node, or going comatose. Add the following line above the `exit` command:

```
exit_on_error rcnovell-cifs monitor
```

Pool cluster resources that were created on OES 2 SP3 and earlier clusters might use the CIFS `status` command in the monitoring script. It checks the status and immediately invokes the resource’s failover settings if the `cifs` daemon is not running. To take advantage of the restart capability that is offered by the `monitor` command, you can modify the line to use “`monitor`” instead of “`status`”.

You can also add lines to the monitor script that check the status of processes used by clustering. For information, see [Section 9.6.4, “Monitoring Services that Are Critical to Clustering,” on page 138](#).

11.8 Adding Advertising Protocols for NSS Pool Cluster Resources

Shared NSS pools support three advertising protocols: NCP (default, required), 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, ensure that the pool is cluster migrated back to the server where it was originally created.

- 1 Log in to iManager as an administrator user.
- 2 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 104](#).
- 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 > Cluster Options*, then browse to locate and select the Cluster object of the cluster you want to manage.

- 5 On the Cluster Options page, use one of the following methods to open the Cluster Pool Properties page for the pool cluster resource that you want to manage:
 - ♦ Select the check box next to the pool cluster resource, then click *Details*.
 - ♦ Click the *Name* link for the pool cluster resource.

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		

- 6 On the Cluster Pool Properties page, click the *Protocols* tab.
- 7 If you are enabling the resource for Novell AFP, select the *AFP* check box, then click *Apply*.
- 8 If you are enabling the resource for Novell CIFS, select the *CIFS* check box and view the *CIFS Virtual Server Name*, then click *Apply*.

A default name (up to 15 characters) is suggested for the CIFS Virtual Server Name when you first add the advertising protocol.

If you later want to modify the CIFS Virtual Server Name, you must use the CIFS plug-in for iManager to manage the CIFS share name of an existing share. For information, see “[Setting CIFS General Server Parameters](#)” in the *OES 11: Novell CIFS for Linux Administration Guide*.

- 9 Online the cluster resource.

For information, see [Section 8.5, “Onlining and Offlining \(Loading and Unloading\) Cluster Resources from a Cluster Node,”](#) on page 104.

11.9 Adding NFS Export for a Clustered Pool Resource

You can add NFS support for a clustered NSS volume by adding an `exportfs(8)` command to the load script and unload script for a clustered pool resource. Basically, the order in load script should be NSS, NFS, IP, and NCP. The order is reversed in the unload script.

- ♦ [Section 11.9.1, “Sample NFS Load Script,”](#) on page 175
- ♦ [Section 11.9.2, “Sample NFS Unload Script,”](#) on page 176

11.9.1 Sample NFS Load Script

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

# Activate the NSS pool and mount its volume
exit_on_error nss /poolact=POOL_16
exit_on_error ncpcon mount VOL_161=224

# Export the volume for NFS
exit_on_error exportfs -o rw,sync,no_root_squash,fsid=216 */media/nss/VOL_161

exit_on_error add_secondary_ipaddress 192.168.188.16

exit_on_error ncpcon bind --ncpservname=CLUS1_POOL_16_SERVER --ipaddress=192.168.188.16

exit 0
```

11.9.2 Sample NFS Unload Script

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

ignore_error ncpcon unbind --ncpservname=CLUS1_POOL_16_SERVER --ipaddress=192.168.188.16

ignore_error del_secondary_ipaddress 192.168.188.16

#Unexport a volume for NFS
ignore_error exportfs -u */media/nss/VOL_161

# Deactivate the pool and volume
ignore_error nss /pooldeact=POOL_16

exit 0
```

11.10 Mirroring and Cluster-Enabling Shared NSS Pools and Volumes

- ♦ [Section 11.10.1, “Understanding NSS Mirroring,” on page 176](#)
- ♦ [Section 11.10.2, “Requirements for NSS Mirroring,” on page 178](#)
- ♦ [Section 11.10.3, “Creating and Mirroring NSS Pools on Shared Storage,” on page 178](#)
- ♦ [Section 11.10.4, “Verifying the NSS Mirror Status in the Cluster,” on page 180](#)

11.10.1 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 11-1](#) 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 11-1 Single Cluster with Mirrored NSS Volumes in Separate Buildings

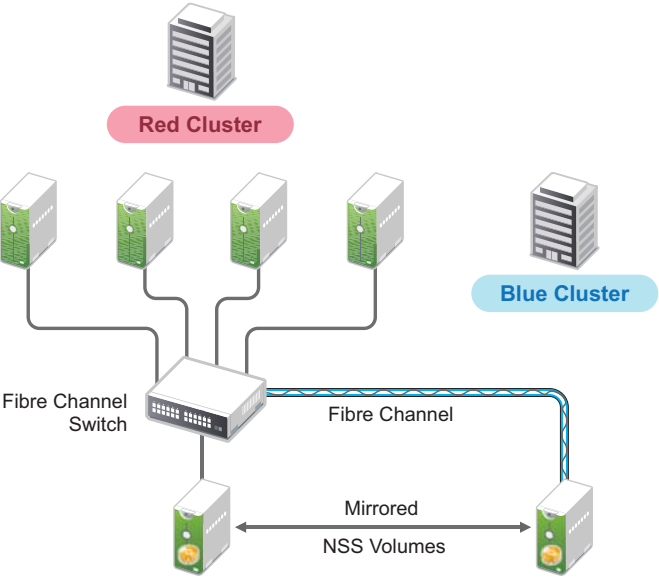
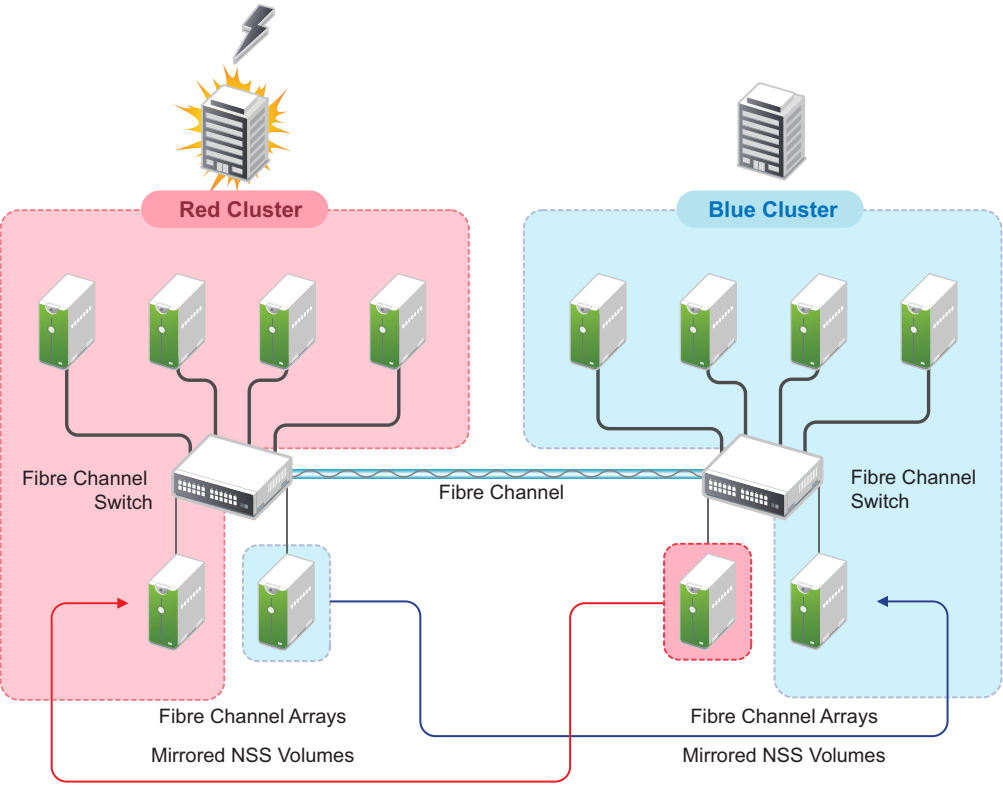


Figure 11-2 depicts a more complex hardware configuration in which two clusters are placed in separate buildings, but each has one side of a mirrored NSS volume in the building with the other cluster. If a disaster occur in one building, data is still safe and immediately accessible on the mirrored NSS volume in the other building.

Figure 11-2 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.

11.10.2 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.

11.10.3 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 11 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 178](#)
- ♦ [“Mirroring the Partition of an Existing NSS Pool” on page 179](#)

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 as described in [Section 11.3.3, “Creating a Cluster-Enabled Pool and Volume with NSSMU,”](#) on page 165.

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%.

11.10.4 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.

11.11 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. Renaming the Pool automatically changes the Pool Resource object name in eDirectory. It does not change the virtual server name (the NCS:NCP Server object).

- 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.
- 6 If you previously used the `cluster rename` command to use a customized name for the cluster resource, you must repeat that process to return to the customized name. For information, see [Section 9.13, "Renaming a Cluster Resource," on page 147](#).

11.12 Adding a Volume to a Clustered Pool

When you use NSSMU or the Storage plug-in to iManager to add an NSS volume to a clustered pool, a `mount` command is automatically added to the load script for the pool cluster resource. Offline and online the primary pool cluster resource to apply the modified load script.

If you add a volume to the primary pool for a clustered DST shadow volume, the `mount` command is added twice in the primary pool's cluster load script, once after the primary pool's activation command and once after the secondary pool's activation command. You must manually delete the instance that occurs after the secondary pool's activation, then offline and online the primary pool cluster resource to apply the modified load script.

- 1 In iManager under *Roles and Tasks*, select *Storage > Volumes*.
- 2 Browse to select the Cluster object of the cluster where the pool is currently online.
You can also select the server where the pool cluster resource is currently mounted.
- 3 Click *New* to open the *New Volume* wizard, then create the new volume on the shared pool.
For information, see "[Creating Unencrypted NSS Volumes](#)" in the *OES 11: NSS File System Administration Guide for Linux*.
- 4 Under *Roles and Tasks*, select *Clusters > Cluster Manager*, then browse to select the Cluster object of the cluster.
- 5 If you added a volume to the primary pool for a DST shadow volume, modify the load script to remove the instance of the `mount` command that immediately follows the pool activation line for the secondary pool.
 - 5a Select *Clusters > Cluster Options*.
 - 5b Click the name link of the pool cluster resource to open its Properties page, then click the *Scripts* tab.
 - 5c Delete the `mount` command for the new volume that follows the pool activation command for the secondary pool.
 - 5d Click *OK* to save your changes.
 - 5e Click *Clusters > Cluster Manager* to return to the Cluster Manager page.
- 6 In the *Cluster Resource* list, select the check box next to the pool cluster resource, then click *Offline* to take the resource offline.
- 7 In the *Cluster Resource* list, select the check box next to the pool cluster resource, then click *Online* to bring the resource online.

This applies the modified load script, which mounts the newly created volume.

11.13 Changing an Assigned 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. You can view the volume IDs assigned on a node by using the `ncpcon volumes` command.

In older operating systems, there was a mounted volume limit of 64 volumes (values 0 to 63). Some older applications might have hardcoded the old maximum limit of 64 mounted volumes, and might not be able to handle volume IDs greater than 63. You can use the Clusters plug-in to iManager to modify the volume ID in the scripts for a given cluster resource in order to specify a value that works for the application.

Changing the volume ID does not affect the ability to log in to, back up, or access the data. However, there is a brief disruption of service as the cluster resource is offlined and online to apply the script changes. If you modify the volume ID for a volume in the cluster resource scripts, ensure that you do the following:

- ☐ Volume IDs that you manually assign must be unique across every volume on all servers in cluster.
- ☐ After the value is changed, you must offline and online the cluster resource for the volume in order to mount the volume with its new volume ID.
- ☐ After the volume is mounted with its new ID, the clients must log out and log in to the volume in order to reconnect to the volume with its new volume ID. Automatic reconnection after cluster resource failovers or migrations occurs properly after this one-time reset.

Some clients might cache the volume IDs. To reset the cached value, the client must be rebooted and reconnected to the volume.

- ☐ After the volume is mounted with its new ID, if the backup software is running on a client connection, you might need to restart the backup to reconnect to the volume with its new volume ID. Automatic reconnection after cluster resource failovers or migrations occurs properly after this one-time reset.

11.14 Expanding the Size of a Clustered Pool

NSS pools are made up of segments of space from one or more devices up to a maximum total size of 8 TB (terabytes). Each device can be up to 2 TB in size if a DOS partitioning scheme is used, or up to 8 TB in size if a GPT partitioning scheme is used. You can expand a pool by adding more space from the same device or from a different shared device. If you extend the size of a LUN device that is currently being used by the pool, you must make the extended size visible to all nodes in the cluster before you add the newly available space to the clustered pool. If you add a new shared LUN device, you must make the device visible to all nodes in the cluster before you add space from the device to the clustered pool.

- ♦ [Section 11.14.1, “Planning to Expand a Clustered Pool,” on page 182](#)
- ♦ [Section 11.14.2, “Increasing the Size of a Clustered Pool by Extending the Size of Its LUN,” on page 184](#)
- ♦ [Section 11.14.3, “Increasing the Size of a Clustered Pool by Adding Space from a Different Shared LUN,” on page 189](#)
- ♦ [Section 11.14.4, “Verifying the Expanded Pool Cluster Resource on All Nodes,” on page 193](#)

11.14.1 Planning to Expand a Clustered Pool

Consider the guidelines in this section when you plan to expand the size of a clustered NSS pool.

- ♦ [“Free Space on a Shared Device” on page 183](#)
- ♦ [“Extending a LUN Might Require a Service Outage” on page 183](#)

- ♦ [“Adding a LUN Might Require a Service Outage” on page 183](#)
- ♦ [“Multipath I/O Devices” on page 183](#)

Free Space on a Shared Device

To expand the size of a pool, you need free unpartitioned space on the same shared device or on other shared devices that can be failed over to the same node in a cluster. To extend the size of an existing LUN, the free unpartitioned space must immediately follow the LUN on the same shared device.

Extending a LUN Might Require a Service Outage

Novell Cluster Services and NSS allow you to perform an online extension of a LUN that is used by a clustered pool without taking the resource offline or stopping Cluster Services. However, not all SAN storage arrays and vendor-specific storage drivers fully support transparent online extension of LUNs. Refer to the third-party vendor documentation for information about how to extend a LUN. We recommend that you test the specific hardware environment and hardware configuration to confirm that it behaves correctly before you perform an online LUN extension.

IMPORTANT: To prevent any chance of data loss or corruption, we recommend that you back up the volumes before you expand a LUN.

After you extend the LUN size on the SAN storage array, you must scan for devices on each node to update the node's storage map with its new size information before you expand the pool size on the active node, or allow the pool cluster resource to fail over to other nodes. If the device has multiple I/O paths, you must also update each node's multipath map. Depending on your SAN storage array, device driver, and server hardware, a server restart might be required to force the node to recognize the extended LUN size.

Adding a LUN Might Require a Service Outage

Novell Cluster Services and NSS allow you to add a LUN to a clustered pool without taking the resource offline or stopping Cluster Services. The LUN must be able to be failed over to the same node in a cluster as the pool cluster resource. Refer to the third-party vendor documentation for information about how to add a LUN.

After you create a new LUN on the SAN storage array and assign it to all nodes, you must scan for devices on each node to update the node's storage map with the new device before you expand the pool size on the active node, or allow the pool cluster resource to fail over to other nodes. If the device has multiple I/O paths, you must also update each node's multipath map. Depending on your SAN storage array, device driver, and server hardware, a server restart might be required to force the node to recognize the new LUN.

Multipath I/O Devices

Consider the following requirements and guidelines as you work with devices that have multiple I/O paths:

- ♦ The procedures in this section assume that you use the Linux Device Mapper - Multipath I/O (DM-MPIO) software to manage the multiple I/O paths for devices. You must modify the instructions accordingly if you are using a third-party multipath solution.

- ♦ Before you expand the pool size on the active node or allow the pool cluster resource to fail over to other nodes, you must update the multipath map on each node in the cluster so that DM-MPIO recognizes the LUN changes:
 - ♦ **Extend the LUN size:** On each node, you can restart the `multipathd` daemon, or you can use the `multipathd` command to resize the multipath map for the LUN.
 - ♦ **New LUN:** On each node, use the `multipath -v2` command to rebuild the node's multipath map.
- ♦ For a device with multiple I/O paths, use the device's multipath device name (such as `mpathb`) or device ID.
- ♦ Steps in this section that pertain only to multipath I/O devices are preceded by the keyword "MPIO".

11.14.2 Increasing the Size of a Clustered Pool by Extending the Size of Its LUN

One way to increase the size of a clustered pool is to add space to its LUN device on the SAN storage array, and then increase the size of the pool to use the newly available space. Before you extend the size of a LUN, ensure that you understand the cautions in ["Extending a LUN Might Require a Service Outage"](#) on page 183.

- ♦ ["Extending the Size of a LUN"](#) on page 184
- ♦ ["Increasing the Pool Size by Using Free Space on the Same LUN"](#) on page 187

Extending the Size of a LUN

You can extend the size of the LUN if there is free unpartitioned space on the device that immediately follows the LUN. If the LUN has multiple I/O paths, special handling is required to update the device size reported in the multipath map on each node. These steps are marked with the keyword "MPIO".

- 1 Identify the device that is used by the pool cluster resource:
 - 1a In iManager, go to *Storage > Pools*, then select the Cluster object of the cluster.
 - 1b In the Pools list, select the pool, then view the device name of the LUN under the pool *Details* area.

For example, a device with a single path has a device node name such as `sdb`. A device with multiple I/O paths has a multipath device name such as `mpathb`.
- 2 (Offline LUN extension only) If the SAN storage array or vendor-specific storage driver does not support online LUN extension, take the pool cluster resource offline. Open a terminal as the root user, then at the command prompt, enter:


```
cluster offline <resource_name>
```
- 3 On the SAN storage array, use the third-party vendor tools to extend the size of the existing LUN.

Devices can be up to 2 TB in size with the DOS partitioning scheme. Devices can be up to 8 TB in size with the GPT partitioning scheme.

For example, increase the size of the existing LUN from 100 GB to 150 GB.

- 4 Beginning with the active node, perform the following steps on each node in turn in order to update the node's storage map with the new LUN size:

4a Log in to the node as the root user, then open a terminal console.

4b Scan for devices to recognize the new size for the LUN. At the command prompt, enter:

```
/bin/rescan-scsi-bus.sh -forcerescan [--luns=XX]
```

Use the `--luns=XX` option if you want to scan only for the extended LUN. You can use the `lsscsi(8)` command to locate the LUN number for the device. Information for each device is output on a single line. The first entry is four numbers separated by colons that represent the `Host:Bus:Target:LUN` of the SCSI device.

For information about other `rescan-scsi-bus.sh` options, see “Scanning for New Devices without Rebooting” (http://www.suse.com/documentation/sles11/stor_admin/data/scandev.html) in the *SUSE Linux Enterprise Server 11 Storage Administration Guide* (http://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html).

4c Scan for storage objects. At the command prompt, enter:

```
nlvm rescan
```

4d (MPIO) If the device has multiple I/O paths, do either of the following to update the multipath map with the new size information for the extended LUN:

- ♦ **Resize the multipath map for the LUN:** At the command prompt, enter:

```
multipathd -k'resize map <mpio_map_name>'
```

There is no space between `-k` and the `'resize map <mpio_map_name>'` option. You can use the `multipath -ll` command to locate the multipath map name for the device.

For example, the following command resizes the map entry for a LUN device with the multipath map ID of `36001c230ce31da000eb0fa8d1ccb0c02`:

```
multipathd -k'resize map 36001c230ce31da000eb0fa8d1ccb0c02'
```

You can alternatively use the `multipathd -k` command to enter interactive mode and issue the `resize map <mpio_map_name>` command. From this mode, the available commands can be viewed by entering `help`. When you are finished entering commands, press `Ctrl+D` to quit.

- ♦ **Restart the multipathd daemon:** At the command prompt, enter one of the following commands:

```
rcmultipathd restart
```

```
/etc/multipathd restart
```

4e (MPIO) Verify that the new size of the LUN is correctly reported to DM-MPIO. At the command prompt, enter:

```
multipath -ll
```

The `-ll` option shows the current multipath topology from all available information (`sysfs`, the device mapper, path checkers, and so on).

For example, the `mpathb` device reports a new size of 150 GB.

```

mpathb (36001438005de9b8d0000800007520000) dm-1 HP,HSV450
[size=150G][features=1 queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=100][active]
  \_ 4:0:1:7 sdae      65:224 [active][ready]
  \_ 2:0:1:7 sdu       65:64  [active][ready]
\_ round-robin 0 [prio=20][enabled]
  \_ 4:0:0:7 sdv       65:80  [active][ready]
  \_ 2:0:0:7 sdg       8:96   [active][ready]

```

- 4f** Repeat [Step 4a](#) through [Step 4e](#) on each node in the cluster.

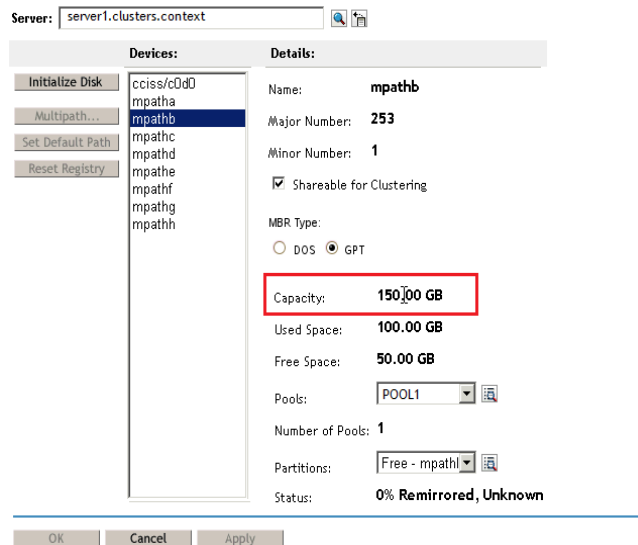
Depending on your SAN storage array, device driver, and server hardware, a server restart might be required to force a node to recognize the extended LUN size.

- 5** In iManager, verify that the new device size is reported correctly to NSS:

5a In iManager, go to *Storage > Devices*, then select the Cluster object of the cluster.

5b Select the device (such as `mpathb`), then view the device capacity under the device's *Details* area.

For example, the `mpathb` device's *Capacity* field reports a size of 150 GB.



5c If the new LUN size is reported correctly, continue to [Step 6](#). If the size is not reported correctly, repeat [Step 4](#), then check again.

- 6** (Offline LUN extension only) If you took the pool cluster resource offline in [Step 2](#), you can now bring pool cluster resource online. Open a terminal as the `root` user, then at the command prompt, enter:

```
cluster online <resource_name>
```

- 7** After the new LUN size has been recognized by all nodes, continue with [“Increasing the Pool Size by Using Free Space on the Same LUN”](#) on page 187.

Increasing the Pool Size by Using Free Space on the Same LUN

The pool size is not automatically expanded to use the new free space on the LUN. After you have expanded the size of the existing LUN and the new size is recognized by all nodes in the cluster, you can increase the size of the pool by allocating space from the newly available free space.

- 1 In iManager, go to *Storage > Pools*, then select the Cluster object of the cluster.
- 2 On the Pool page, select the pool, then view the pool's current size in the *Pool Details* area.

The *Total Space* for the pool reports only the amount of space that has been allocated to the pool. The pool size is not automatically expanded to use the new free space on the LUN.

New...
Delete
Rename...
Activate
Deactivate
Increase size...
Snapshot...
Properties...
Update eDirectory
Deleted Volumes...
Offline

Pools:

POOL1

Details:

Name:

POOL1

Mount Point:

/opt/novell/nss/mnt/.pools/POOL1

Partitions:

mpathb1.1

Number of Partitions:

1

State:

Active

LSS Type:

ZLSS

Share State:

Sharable for Clustering

Volumes:

VOL1

Number of Volumes:

1

Devices:

mpathb

Number of Devices:

1

Total Space:

100.00 GB

Free Space:

160.00 KB

Used Space:

100.00 GB

Purgeable Space:

0.00 GB

Other in-use space:

100.00 GB

Block Size:

4.00 KB

Creation Date:

Sep 16, 2010 2:58:27 pm

Last Update:

Sep 16, 2010 2:58:27 pm

- 3 Increase the size of the pool by adding space from the same LUN:

- 3a On the Pool page, select the pool, then click *Increase Size*.

The Expand a Pool wizard opens and presents all devices with available free space.

- 3b In the list of devices, select the check box for the newly expanded LUN, such as mpathb, then specify the amount of space to add in MB.

Typically, this is the reported amount of free space available. For example:

	Used Size (MB)	Device Name	Free Size (MB)
<input type="checkbox"/>	0	mpatha	1015
<input checked="" type="checkbox"/>	51201	mpathb	51201
<input type="checkbox"/>	0	mpathc	102399
<input type="checkbox"/>	0	mpathd	20479
<input type="checkbox"/>	0	mpathe	20479
<input type="checkbox"/>	0	mpathf	20479
<input type="checkbox"/>	0	mpathg	20479
<input type="checkbox"/>	0	mpathh	20479

Pool Size: 153599

- 3c Click *Finish* to save and apply the change.

- 4 On the Pool page, select the pool, then verify that the *Total Space* field reports the new size. For example, the new size of POOL1 is 150 GB.

Pools:	Details:
New... Delete Rename... Activate Deactivate Increase size... Snapshot... Properties... Update eDirectory Deleted Volumes... Offline	<p>POOL1</p> <p>Name: POOL1</p> <p>Mount Point: /opt/novell/nss/mnt/pools/POOL1</p> <p>Partitions: mpathb1.1</p> <p>Number of Partitions: 2</p> <p>State: Active</p> <p>LSS Type: ZLSS</p> <p>Share State: Sharable for Clustering</p> <p>Volumes: VOL1</p> <p>Number of Volumes: 1</p> <p>Devices: mpathb</p> <p>Number of Devices: 1</p> <p>Total Space: 150.00 GB</p> <p>Free Space: 49.85 GB</p> <p>Used Space: 100.15 GB</p> <p>Purgeable Space: 0.00 GB</p> <p>Other in-use space: 100.15 GB</p> <p>Block Size: 4.00 KB</p> <p>Creation Date: Sep 16, 2010 2:58:27 pm</p> <p>Last Update: Sep 16, 2010 2:58:27 pm</p>

- 5 Go to *Storage > Partitions*, then verify that a new partition exists for the LUN and is assigned to the pool.

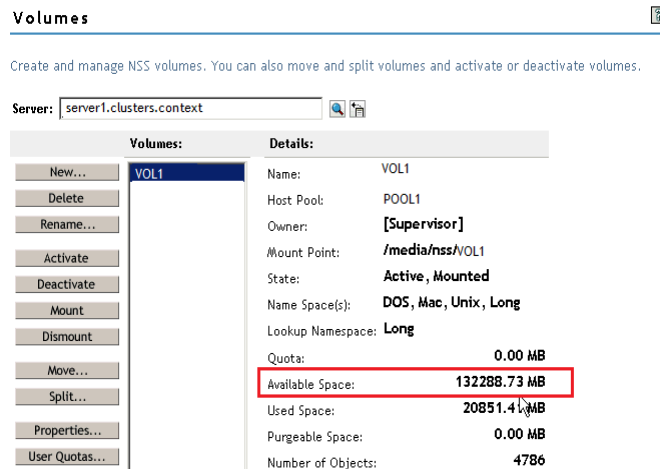
For example, the partitions mpathb1.1 and mpathb1.2 are allocated to POOL1.

Server: server1.clusters.context

Partitions							
New Edit Delete Details							
<input type="checkbox"/>	Name	Type	Status	Device Name	RAID Name	Pool Name	Size
<input type="checkbox"/>	cciss/c0d0_freespace1	Free Space	Available	cciss/c0d0			27.07 MB
<input type="checkbox"/>	cciss/c0d0p1	Linux	In Use	cciss/c0d0			509.84 MB
<input type="checkbox"/>	cciss/c0d0p2	Linux Swap	In Use	cciss/c0d0			15.99 GB
<input type="checkbox"/>	cciss/c0d0p3	Linux	In Use	cciss/c0d0			256.88 GB
<input type="checkbox"/>	cluster1.sbd	Clustering	In Use	mpatha			7.98 MB
<input type="checkbox"/>	mpatha1.nwfreespace1	Free Space	Available	mpatha			7.66 MB
<input type="checkbox"/>	mpatha_freespace1	Free Space	Available	mpatha			1008.31 MB
<input type="checkbox"/>	mpathb1.1	NSS	In Use	mpathb		POOL1	100.00 GB
<input type="checkbox"/>	mpathb1.2	NSS	In Use	mpathb		POOL1	50.00 GB
<input type="checkbox"/>	mpathb_freespace1	Free Space	Available	mpathb			2.34 MB

- 6 For each volume on the pool, if the volume quota is set to *No Quota*, the volume can grow to the size of the pool. Go to *Storage > Volumes*, select the volume, then verify that its *Available Space* field has increased by the amount of space that was added to the pool.

For example, by increasing the pool size by 50 GB, the available space for VOL1 increased from about 82 GB to about 132 GB.



7 Continue with “[Verifying the Expanded Pool Cluster Resource on All Nodes](#)” on page 193.

11.14.3 Increasing the Size of a Clustered Pool by Adding Space from a Different Shared LUN

You can increase the size of a clustered pool by adding space to the pool from a different shared LUN device that can be failed over with the resource. Before you add a LUN, ensure that you understand the cautions in “[Adding a LUN Might Require a Service Outage](#)” on page 183.

- ♦ “[Creating and Sharing a New LUN Device](#)” on page 189
- ♦ “[Increasing the Pool Size by Adding Free Space from a Different LUN](#)” on page 191

Creating and Sharing a New LUN Device

If you do not have a shared device that can be failed over with the resource, you can create a new LUN and share it with all nodes in the cluster. If the LUN has multiple I/O paths, special handling is required to add the device to the multipath map on each node. These steps are marked with the keyword “MPIO”.

- 1 Use third-party tools for the SAN storage array to create a new shared LUN device, and assign it to each cluster node.
- 2 Beginning with the active node, perform the following steps on each node in turn in order to update the node’s storage map and multipath map with the new device:
 - 2a Log in to the node as the root user, then open a terminal console.
 - 2b Scan for devices to help the OS recognize the new LUN. At the command prompt, enter:

```
/bin/rescan-scsi-bus.sh -forcerescan
```

For information about other `rescan-scsi-bus.sh` options, see “[Scanning for New Devices without Rebooting](#)” (http://www.suse.com/documentation/sles11/stor_admin/data/scandev.html) in the *SUSE Linux Enterprise Server 11 Storage Administration Guide* (http://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html).

- 2c** Use the `lsscsi` command to verify that the LUN is seen by the OS. At the command prompt, enter:

```
lsscsi
```

Information for each device is output on a single line. The first entry is four numbers separated by colons that represent the Host:Bus:Target:LUN of the SCSI device. In the sample output below, the `/dev/sdc` and `/dev/sde` devices are two individual paths for LUN 2.

```
# lsscsi
[4:0:0:0]    disk      DGC        RAID 5          0216  /dev/sdb
[4:0:0:1]    disk      DGC        RAID 5          0216  /dev/sdf
[4:0:0:2]    disk      DGC        RAID 5          0216  /dev/sdc
[4:0:1:0]    disk      DGC        RAID 5          0216  /dev/sdd
[4:0:1:1]    disk      DGC        RAID 5          0216  /dev/sdg
[4:0:1:2]    disk      DGC        RAID 5          0216  /dev/sde
```

If the LUN is not seen, repeat this step using different `lsscsi` command line parameters to scan for all devices (such as `-w -c -l`). For information about command options, see the `lsscsi(8)` man page.

If the LUN is still not seen by the OS, a server restart might be required.

- 2d** Scan for storage objects. At the command prompt, enter:

```
nlvm rescan
```

- 2e** (MPIO) If the device has multiple I/O paths, rebuild the multipath map. At the command prompt, enter:

```
multipath -v2
```

- 2f** (MPIO) Verify that the new LUN is reported to DM-MPIO. At the command prompt, enter:

```
multipath -ll
```

The `-ll` option shows the current multipath topology from all available information (`sysfs`, the device mapper, path checkers, and so on).

- 2g** Verify that the LUN is now visible to NSS on the node. In NSSMU, go to the Devices page, then verify that the new LUN is listed.

For example, the new device is `mpathj` on the active node. Remember that a multipath device might have a different multipath device name on each node.

NSS version 4.12a Build 3018 NSS Management Utility	
Devices	Device Information
cciss/c0d0	
mpatha	Name: mpathj
mpathb	objectID: mpathj
mpathc	Removable: No
mpathd	
mpathe	Capacity: 25.00 GB
mpathf	Partitioned Space: 0.00 KB
mpathg	Unpartitioned Space: 0.00 KB
mpathh	
* mpathj	Shareable for clustering: No

- 2h** Repeat [Step 2a](#) through [Step 2g](#) on each node in the cluster.

Depending on your SAN storage array, device driver, and server hardware, a server restart might be required to force a node to recognize the new LUN.

3 In iManager, initialize the new LUN, then mark it as shareable for clustering:

3a In iManager, go to *Storage > Devices*.

3b On the Devices page, select the Cluster object of the cluster.

Selecting the Cluster object connects you to the current master node of the cluster.

3c In the Devices list, select the newly added LUN device.

WARNING: Ensure that you select the correct device. Initializing a device destroys all data on it.

3d Click *Initialize*, then click *OK* to confirm that you want to initialize the device.

If you are prompted for the partitioning scheme, specify your preferred partitioning scheme as DOS (up to 2 TB in size) or GPT (up to 8 TB in size).

3e In the Devices list, select the new LUN device.

3f In the *Details* area, select the *Shareable for Clustering* check box, then click *Apply*.

The screenshot shows the iManager interface for the 'Storage > Devices' page. The 'Server' field is set to 'server1.clusters.context'. On the left, the 'Devices' list shows several multipath devices, with 'mpathj' selected. The 'Details' pane on the right shows the configuration for 'mpathj': Name: mpathj, Major Number: 253, Minor Number: 40, MBR Type: GPT (selected), Capacity: 25.00 GB, Used Space: 16.00 KB, Free Space: 25.00 GB, Pools: (empty), Number of Pools: (empty), Partitions: Free - mpathj, Status: 0% Remirrored, Unknown. The 'Shareable for Clustering' checkbox is checked. At the bottom are 'OK', 'Cancel', and 'Apply' buttons.

4 After the new LUN is recognized by all nodes, and has been initialized and shared, continue with [“Increasing the Pool Size by Adding Free Space from a Different LUN”](#) on page 191.

Increasing the Pool Size by Adding Free Space from a Different LUN

After you have set up a different shared device that can be failed over with the pool cluster resource, you can increase the size of the pool by allocating space from it.

1 In iManager, go to *Storage > Pools*, then select the Cluster object of the cluster.

Selecting the Cluster object automatically connects you to the current master node of a cluster.

2 Select the pool, then view the pool’s current size.

The *Total Space* for the pool reports only the amount of space that has been allocated to the pool. For example, Pool7 is 20 GB in size.

Server:

New...
Delete
Rename...
Activate
Deactivate
Increase size...
Snapshot...
Properties...
Update eDirectory
Deleted Volumes...
Offline

POOL1
POOL2
POOL3
POOL4
POOL5
POOL6
POOL7

Name: POOL7
Mount Point: /opt/novell/nss/mnt/.pools/POOL7
Partitions: mpathh1.1
Number of Partitions: 1
State: Active
LSS Type: ZLSS
Share State: Sharable for Clustering
Volumes: VOL7
Number of Volumes: 1
Devices: mpathh
Number of Devices: 1
Total Space: 20.00 GB
Free Space: 19.93 GB
Used Space: 65.89 MB
Purgeable Space: 20.00 KB
Other in-use space: 65.87 MB

3 Increase the size of the pool by adding space from a different LUN:

3a On the Pool page, select the pool, then click *Increase Size*.

The Expand a Pool wizard opens and presents all devices with available free space.

3b In the list of devices, select the check box for the new LUN, such as `mpathj`, then specify the amount of space to add in MB.

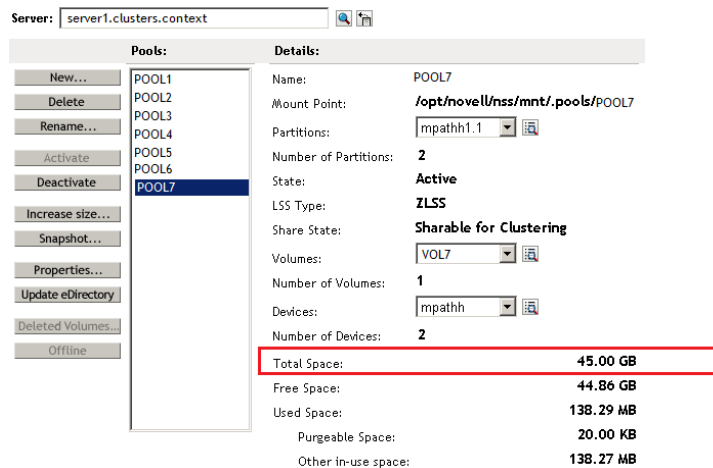
Typically, this is the reported amount of free space available. For example, `mpathj` has 25599 MB available:

	Used Size (MB)	Device Name	Free Size (MB)
<input type="checkbox"/>	<input type="text" value="0"/>	mpatha	1015
<input type="checkbox"/>	<input type="text" value="0"/>	mpathb	2
<input type="checkbox"/>	<input type="text" value="0"/>	mpathc	1
<input checked="" type="checkbox"/>	<input type="text" value="25599"/>	mpathj	25599

Pool Size: 46077

3c Click *Finish* to save and apply the change.

- On the Pool page, select the pool, then verify that the *Total Space* field reports the new size. For example, POOL7 has increased in size from 20 GB to about 45 GB.



- For each volume on the pool, if the volume quota is set to *No Quota*, the volume can grow to the size of the pool. Go to *Storage > Volumes*, select the volume, then verify that a volume's *Available Space* field has increased by the amount of space that was added to the pool.
- After the pool size is increased, continue with [Section 11.14.4, "Verifying the Expanded Pool Cluster Resource on All Nodes,"](#) on page 193.

11.14.4 Verifying the Expanded Pool Cluster Resource on All Nodes

Verify that the expanded pool size is reported correctly on each node in the cluster.

- Log in to a node as the *root* user, then open a terminal console.
- Cluster migrate the pool cluster resource to each node in turn, and use NSSMU to verify that the devices, partitions, pool, and volumes show the correct information on that server:
 - Cluster migrate the pool cluster resource.


```
cluster migrate resource_name node_name
```
 - Launch NSSMU by entering


```
nssmu
```
 - On the *Devices* page, select the device and verify that the new size is reported.
 - On the *Partitions* page, select each partition for the device, then verify that the partitions are shown for the expanded pool.
 - On the *Pools* page, select the pool, then verify the total space available.
 - On the *Volumes* page, select each volume that has no quota set, then verify that the available space has increased by the amount of space added to its pool.
 - If the information is incorrect, wait a few minutes to allow time for the system to recognize the new partitions, then try again.
- Repeat the previous validation for each node in the cluster.

11.15 Deleting NSS Pool Cluster Resources

Ensure that you offline the cluster resource before attempting to delete either the cluster resource or the clustered pool. For example, if you want to unshare a pool, offline the cluster resource for the pool before you mark the pool or the device as Not Shareable for Clustering, then you can delete the eDirectory object for the cluster resource.

WARNING: If you attempt to delete a cluster resource without first offlining it, deletion errors occur, and the data associated with the clustered pool is not recoverable.

All resource configuration must happen from the master node. On the Cluster Options page for iManager, 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 9.14, “Deleting Cluster Resources,” on page 148](#).

12 Configuring and Managing Cluster Resources for LVM Volume Groups

After you have installed and configured Novell Cluster Services, you can create shared cluster resources for Linux Logical Volume Manager (LVM) volume groups. You create an LVM logical volume on the volume group, and add a Linux POSIX file system such as Ext2, Ext3, ReiserFS, and XFS.

This section describes how to configure the LVM volume group cluster resource, logical volume, and file system. You can also create a virtual server name for the resource, and bind it to the resource's IP address.

- ♦ [Section 12.1, "Requirements for Creating LVM Cluster Resources," on page 195](#)
- ♦ [Section 12.2, "Initializing a SAN Device," on page 197](#)
- ♦ [Section 12.3, "Configuring an LVM Cluster Resource with NSSMU," on page 198](#)
- ♦ [Section 12.4, "Configuring an LVM Cluster Resource with LVM Commands and the Generic File System Template," on page 208](#)
- ♦ [Section 12.5, "Creating a Virtual Server Object for an LVM Volume Group Cluster Resource," on page 220](#)
- ♦ [Section 12.6, "Renaming the Mount Point Path for a Clustered LVM Volume," on page 223](#)
- ♦ [Section 12.7, "Disabling Clustering for an LVM Volume," on page 225](#)
- ♦ [Section 12.8, "Deleting a Clustered LVM Volume Group and Logical Volume," on page 226](#)
- ♦ [Section 12.9, "LVM Management Tools," on page 227](#)

12.1 Requirements for Creating LVM Cluster Resources

Your system must meet the requirements in this section in addition to the cluster requirements described in [Chapter 4, "Planning for Novell Cluster Services," on page 37](#).

- ♦ [Section 12.1.1, "Novell Cluster Services," on page 196](#)
- ♦ [Section 12.1.2, "Linux Logical Volume Manager 2 \(LVM2\)," on page 196](#)
- ♦ [Section 12.1.3, "Clustered Logical Volume Manager Daemon \(CLVMD\)," on page 196](#)
- ♦ [Section 12.1.4, "Resource IP Address," on page 196](#)
- ♦ [Section 12.1.5, "Shared Storage Devices," on page 196](#)
- ♦ [Section 12.1.6, "All Nodes Must Be Present," on page 196](#)
- ♦ [Section 12.1.7, "Mixed Mode OES Clusters \(Not Supported\)," on page 197](#)
- ♦ [Section 12.1.8, "NCP File Access with Novell NCP Server," on page 197](#)
- ♦ [Section 12.1.9, "CIFS/Samba File Access with Novell Samba," on page 197](#)

12.1.1 Novell Cluster Services

Novell Cluster Services must be installed, configured, and running when you create and manage the shared LVM volume group and logical volume. The cluster must be active.

12.1.2 Linux Logical Volume Manager 2 (LVM2)

The Linux Logical Volume Manager (LVM) 2 is the software that supports LVM volume groups and logical volumes. LVM2 must be installed and running on each node in the cluster. LVM2 runs on OES 11 nodes automatically; no separate installation or setup is required.

12.1.3 Clustered Logical Volume Manager Daemon (CLVMD)

The Linux Clustered Volume Manager Daemon (CLVMD, `clvmd`) is the software that allows you to exclusively mount a shared volume group on one node at a time in a cluster. It distributes LVM metadata updates around a cluster. CLVM must be installed and running on each node in the cluster. CLVMD runs on OES 11 nodes automatically; no separate installation or setup is required.

IMPORTANT: Ensure that you have installed the latest patches for SUSE Linux Enterprise Server 11 SP1. Clustered LVM volume groups require Linux kernel version 2.6.32.45-0.3 or later.

12.1.4 Resource IP Address

Each cluster resource needs a unique static IP address that is in the same subnet as the IP addresses that are used for the cluster and cluster nodes. The IP address is used to provide access and failover capability for the cluster-enabled volume.

12.1.5 Shared Storage Devices

The shared storage device that you use for an LVM volume group cluster resource must be initialized and have no partitions on it. When the device is used in a cluster resource, LVM uses the entire device for the volume group. Ensure that you size your LUNs accordingly.

IMPORTANT: Do not mark the device as shareable for clustering because doing so adds a 4 KB partition to the device, which makes it unavailable to LVM.

12.1.6 All Nodes Must Be Present

LVM requires the presence of all the nodes to modify metadata on shared storage. This allows LVM to get the exclusive locks it needs to perform actions on shared storage.

Before you attempt to create or modify LVM volume group cluster resources:

- ♦ All of the nodes must be joined in the cluster and running properly.
- ♦ The `clvmd` daemon must be running on all nodes.

12.1.7 Mixed Mode OES Clusters (Not Supported)

LVM volume group cluster resources are not supported in mixed-mode OES clusters. For information, see [Section 6.3, “Requirements and Guidelines for Upgrading Clusters from OES 2,” on page 87](#).

12.1.8 NCP File Access with Novell NCP Server

Novell NCP Server can be used to provide NCP file access to Linux POSIX file systems on OES 11 servers. Its NCP volumes feature can be used to provide NCP access to files on an LVM volume group cluster resource.

NCP Server must be installed, configured, and running on each node in the cluster. You must modify the Generic File System resource template scripts to add commands that mount and dismount the NCP volume, and create a NCP Virtual Server object for the volume group.

For information about setting up NCP volumes on a resource, see “[Configuring NCP Volumes with Novell Cluster Services](#)” in the *OES 11: NCP Server for Linux Administration Guide*.

12.1.9 CIFS/Samba File Access with Novell Samba

Novell Samba is the Linux Samba software that has been modified to work with Novell eDirectory. It can be used to provide CIFS/Samba access to files on an LVM volume group cluster resource.

Novell Samba must be installed and configured on each node in the cluster. The cluster load script starts the service when you online the Samba cluster resource, and the unload script stops the service when you offline it.

For information about using the Samba resource template to create a Samba cluster resource that is based on an LVM volume group, see “[Configuring Samba for LVM Volume Groups and Novell Cluster Services](#)” in the *OES 11: Novell Samba Administration Guide*.

12.2 Initializing a SAN Device

When you initialize the SAN device that you want to use for the LVM volume group, do not mark it as shareable for clustering.

- 1 Ensure that the SAN device is attached to all nodes in the cluster.
- 2 Log in to the master node of the cluster as the `root` user, then open a terminal console.
- 3 Launch NSSMU by entering

```
nssmu
```

- 4 In the NSSMU main menu, select *Devices*, then press Enter.
- 5 In the *Devices* list, select the SAN device (such as `sdd`), then view information about it.

A device that has never been initialized reports a partition type of *Uninitialized*. If the device contains partitions or data, be prepared to lose all data on the device when it is initialized. The clustered volume group requires the entire device.

WARNING: Initializing a device removes all of the data on it.

- 6 Press F3 to initialize the selected device. Read the advisory that this action destroys all data on the device, then press Y (yes) to confirm, or press N (no) or press Esc to cancel and choose a different device.
- 7 Specify the partitioning scheme to use as DOS or GPT, then press Enter.
DOS supports devices up to 2 TB in size. GPT supports devices of any size.
- 8 Verify that the device is initialized and that it is unshared (that is, *Shareable for Clustering* is set to *No.*)
- 9 Press Esc twice to exit NSSMU.

12.3 Configuring an LVM Cluster Resource with NSSMU

This section describes how to use NSSMU to create and cluste-enable an LVM volume group. NSSMU automatically uses the Generic File System template (`Generic_FS_Template`) to create a volume group cluster resource.

After you create the resource, you can add lines to its load script, unload script, and monitor script to customize the resource for other uses. Compare the `Generic_FS_Template` to the resource template for your product to determine which lines need to be added or modified.

You can alternatively use the Novell Linux Volume Manager `create linux volume` command with the `lvm` and `shared` options to automatically create and cluster-enable an LVM volume group. For information, see “[Create Linux Volume](#)” in the *OES 11: NLVM Reference*. The NLVM command allows you to specify a group name that is different than the volume name.

The examples in this section use following setup. Ensure that you replace the sample values with information for your configuration.

Parameter	Sample Value
Device name for the shared device	<div>/dev/sdd</div> <div>The device is initialized and contains no partitions. It is not enabled as shareable for clustering.</div>
Volume group name	<div>vol44</div> <div>By default, NSSMU uses the logical volume name as the LVM volume group name. If you use the NLVM <code>create linux volume</code> command to create the LVM volume group cluster resource, you can specify a different name for the volume group, such as <code>vg44</code>.</div>
Volume name	vol44
Linux POSIX file system type	<div>ext3</div> <div>Valid values are <code>ext2</code>, <code>ext3</code>, <code>reiserfs</code>, and <code>xfs</code>.</div>
Make options for the file system	<div>None (do not specify a value). Press Enter to continue.</div> <div>For a list of the supported file system options for the file system type you are making, see the <code>mkfs(8)</code> man page and the man page for the specific file system: <code>mkfs.ext2(8)</code>, <code>mkfs.ext3(8)</code>, <code>mkfs.reiserfs(8)</code>, or <code>mkfs.xfs(8)</code>.</div> <div>IMPORTANT: The file system creation fails if you specify a make option that is not supported by the file system type.</div>

Parameter	Sample Value
Mount options	<p>rw</p> <p>The Read/Write (rw) option is specified by default. For a list of available options that work with the file system type you are using, see the <code>mount(8)</code> man page.</p>
Volume size	<p>100 GB</p> <p>A 100 GB LUN is prepared in the shared storage subsystem. It is attached to the nodes in the cluster.</p> <p>You are not prompted to enter a volume size. The clustered LVM volume group and logical volume use the entire device. When you select the device, all of the device's free available space is displayed in the <i>Free Size</i> field.</p>
Resource IP address	<p>10.10.10.44</p> <p>This is the IP address of the virtual server for the cluster resource. The address must be unique and in the same subnet as the cluster's IP address. Specify the IP address in IPv4 format.</p>
Mount device	<p><code>/dev/vol44/vol44</code></p> <p>The mount device path format is</p> <p><code>/dev/<volume_group_name>/<logical_volume_name></code></p> <p>If you specify a different name for the volume group, such as <code>vg44</code>, by using the <code>NLVM create linux volume</code> command, the mount device path is <code>/dev/vg44/vol44</code>.</p>
Mount point	<p><code>/mnt/vol44</code></p> <p>You can use any valid Linux path as the mount point.</p> <p>NSSMU automatically creates the mount point path if it does not exist on this node. However, you must manually create the mount point on each of the other nodes in the cluster. If the path does not exist on a node when you fail over the resource to it, the resource goes comatose. You can alternatively add the following line to the load script to create the path on a node if it does not exist:</p> <p><code>ignore_error mkdir -p \$MOUNT_PATH</code></p>

- ♦ [Section 12.3.1, “Creating an LVM Volume Group Cluster Resource,” on page 200](#)
- ♦ [Section 12.3.2, “Configuring the LVM Cluster Resource Settings,” on page 202](#)
- ♦ [Section 12.3.3, “Viewing or Modifying the LVM Resource Scripts,” on page 205](#)
- ♦ [Section 12.3.4, “Sample LVM Resource Scripts,” on page 206](#)

12.3.1 Creating an LVM Volume Group Cluster Resource

Use the procedure in this section to create and cluster-enable an LVM volume group cluster resource by using NSSMU. The disk will contain one LVM volume group and one logical volume.

- 1 Log in to the master node of the cluster as the `root` user, then open a terminal console.
- 2 If you have not already done so, initialize the SAN device that you want to use for the LVM volume group.

For information, see [Section 12.2, “Initializing a SAN Device,” on page 197](#).

- 3 Launch NSSMU by entering

```
nssmu
```

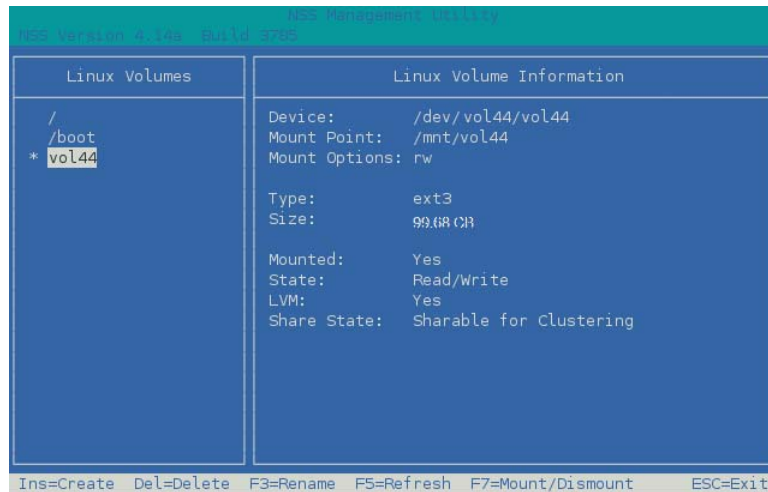
- 4 In the NSSMU main menu, select *Linux Volumes*.
- 5 Press *Insert* to begin creating an LVM volume group, specify the following information.

Parameter	Action
LVM type	Select <i>Cluster Enabled LVM2 Volume</i> .
Volume name	Type the name for the LVM logical volume (such as <code>vol144</code>), then press Enter. The name is also used for the LVM volume group, which is used in the name of the LVM volume group cluster resource.
Volume IP address	Type the resource IP address for the LVM logical volume (such as <code>10.10.10.44</code>), then press Enter.
Volume file system type	Select one of the following the Linux POSIX file systems, then press Enter: <ul style="list-style-type: none">♦ <code>ext3</code>♦ <code>resiserfs</code>♦ <code>ext2</code>♦ <code>xfs</code>
Mount point path	Type the full mount point path for the LVM logical volume (such as <code>/mnt/vol144</code>), then press Enter.
Make options	Press Enter. Specifying a make option is optional.
Mount options	Press Enter to accept the default read and write options (<code>rw</code>). You can add other mount options.
Device	From the list of devices, select the device that you initialized, such as <code>sdd</code> , then press <i>Insert</i> or <i>Enter</i> to select the device. All of the device's free available space is automatically displayed in the <i>Free Size</i> field.

- 6 Press *F3* to accept the setup you entered for the volume group cluster resource.

The resource is created and brought online on the node where you created it. The resource is named `<volume_group_name>_reference`. In the example, the name of the volume group is the same as the logical volume, so the reference name is `vol144_reference`.

7 In the *Linux Volumes* list, select the newly created volume and view information about it.



Parameter	Description
Device	Specifies the full device node path for the LVM logical volume. Example: /dev/vol44/vol44
Mount Point	Specifies the path on the root file system where this volume is mounted. Examples: /mnt/vol44
Mount options	Specifies the mount options that are applied whenever this volume is automatically mounted after a reboot. Example: rw
Type	Specifies the file system type. Examples: ext2, ext3, reiserfs, xfs
Size	Specifies the amount of space reserved for this volume. Example: 99.58 GB
Mounted	Specifies whether the volume is mounted or unmounted. The cluster-enabled resource is automatically brought online on creation, and the load script mounts the logical volume. Value: Yes or No
State	Specifies the availability for the file system. Example: Read/Write
LVM	Specifies whether the specified volume is an LVM logical volume. Value: Yes or No
Share State	Specifies whether the LVM logical volume is cluster enabled for a Novell Cluster Services cluster. Value: Shareable for Clustering

- 8 Press Escape twice to exit NSSMU.
- 9 Continue with [Section 12.3.2, “Configuring the LVM Cluster Resource Settings,”](#) on page 202.

12.3.2 Configuring the LVM Cluster Resource Settings

Use the procedure in this section to verify that the LVM volume group cluster resource was created and is online. Customize the resource policies, monitoring, and preferred nodes settings.

- 1 Open Novell iManager in a Web browser, then log in as an administrator user.
- 2 In *Roles and Tasks*, select *Clusters > Cluster Manager*.
- 3 Browse to select the Cluster object of the cluster where you created the volume group cluster resource.
- 4 In the list of resources, locate the new resource, such as `vol44_resource`, and notice the state of the resource. It should be online and running.

Type	Name	State	Location	Lives	Up Since
	Master_IP_Address_Resource	Running	avalon	1	Dec 8, 2011 1:33:11 PM
	vol44_resource	Running	avalon	1	Dec 8, 2011 2:01:41 PM

- 5 Click the resource's name link to open its Properties page.

- 6 On the Resource Policies page, view and modify the resource's Policy settings if needed. Click *Apply* if you make changes.

Clusters > Cluster Manager

Cluster Resource Properties: vol44_resource.clus1.ncs.novell

Policies Monitoring Preferred Nodes Scripts Business Continuity

Set Start, Failover and Failback modes for the new resource.

Resource Behavior	Failover Mode
<input type="checkbox"/> Resource Follows Master	<input checked="" type="radio"/> Auto
<input type="checkbox"/> Ignore Quorum	<input type="radio"/> Manual

Start Mode	Failback Mode
<input checked="" type="radio"/> Auto	<input type="radio"/> Auto
<input type="radio"/> Manual	<input checked="" type="radio"/> Disable
	<input type="radio"/> Manual

OK Cancel Apply

- 6a (Optional) Select the *Resource Follows Master* check box if you want to ensure that the resource runs only on the master node in the cluster.

If the master node in the cluster fails, the resource fails over to the node that becomes the master.

- 6b (Optional) Select the *Ignore Quorum* check box if you don't want the cluster-wide timeout period and node number limit enforced.

The quorum default values were set when you installed 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.

- 6c By default, the Generic File System resource template sets the Start mode and Failover mode to *Auto* and the Failback Mode to *Disable*. You can change the default settings as needed.

- ♦ **Start Mode:** If the Start mode is set to *Auto*, the resource automatically loads on a designated server when the cluster is first brought up. If the Start mode is set to *Manual*, you can manually start the resource on a specific server when you want, instead of having it automatically start when servers in the cluster are brought up.
- ♦ **Failover Mode:** If the Failover mode is set to *Auto*, the resource automatically moves to the next server in the Assigned Nodes list if there is a hardware or software failure. If the Failover mode is set to *Manual*, you can intervene after a failure occurs and before the resource is started on another node.
- ♦ **Failback Mode:** If the Failback mode is set to *Disable*, the resource continues running on the node it has failed to. If the Failback mode is set to *Auto*, the resource automatically moves back to its preferred node when the preferred node is brought back online. Set the Failback mode to *Manual* to prevent the resource from moving back to its preferred node when that node is brought back online, until you are ready to allow it to happen.

- 7** Enable and configure monitoring for the resource, then click *Apply*.

For information, see [Section 9.6, “Enabling Monitoring and Configuring the Monitor Script,” on page 135](#).

Clusters > Cluster Manager

Cluster Resource Properties: vol44_resource.clus1.ncs.novell

Policies | **Monitoring** | Preferred Nodes | Scripts | Business Continuity

To monitor the health of this resource, enable Resource monitoring. You can set the interval to poll the resource's health, and set an action if it fails to successfully load on the maximum number of local restarts.

☒ Enable Resource Monitoring

Polling Interval: 0 Minutes

Failure Rate

Maximum Local Failures: 0

Time Interval: 0 Minutes
(Example: 3 Failures in 5 minutes)

Failure Action

If the Failure rate settings are reached, perform the following action:

☒ Set Resource as Comatose

☐ Migrate the Resource based on the Preferred Nodes List

☐ Reboot the Hosting Node without Syncing or Unmounting Disks

OK Cancel Apply

- 7a** In the Properties page, select the *Monitoring* tab.

- 7b** Select the *Enable Resource Monitoring* check box to enable resource monitoring for the selected resource.

Resource monitoring is disabled by default.

- 7c** Specify the *Polling Interval* to control how often you want the resource monitoring script for this resource to run.

You can specify the value in minutes or seconds. For information, see [“Polling Interval” on page 136](#).

- 7d** Specify the number of failures (*Maximum Local Failures*) for the specified amount of time (*Time Interval*).

For information, see [“Failure Rate” on page 136](#).

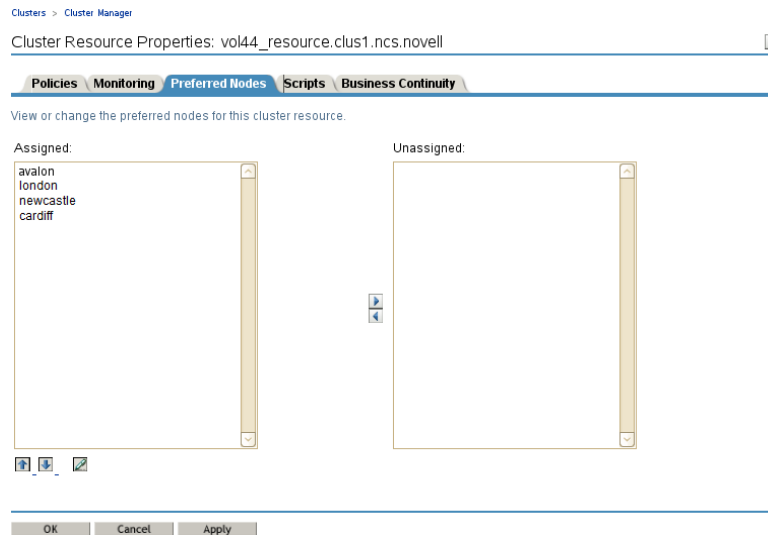
- 7e** Specify the *Failover Action* by indicating whether you want the resource to be set to a comatose state, to migrate to another server, or to reboot the hosting node (without synchronizing or unmounting the disks) if a failure action initiates. The reboot option is normally used only for a mission-critical cluster resource that must remain available.

For information, see [“Failure Action” on page 136](#).

- 8 Click the *Preferred Nodes* tab, assign preferred nodes for the resource by moving them from the *Unassigned* list to the *Assigned* list, then click *Apply*.

When you configure a volume group cluster resource with NSSMU or with NLVM commands, the node where you create it is automatically assigned as a preferred node for the resource.

When you bring a resource online, it is automatically loaded on the most preferred node in the list. If the node is not available, the other nodes are tried in the order that they appear in the list. You can modify the order of the nodes by clicking the Edit (pen) icon to open the list in a text editor. In the editor, click *OK* to close the editor, then click *Apply* to save your changes.



- 9 At the bottom of the page, click *OK* to close the Properties page and save your changes.
The changes do not take effect until the resource is taken offline and brought online again.
- 10 If you modified the settings on any of the pages, you must take the resource offline and bring it online in order for the changes to take effect.
 - 10a In Roles and Tasks, select *Clusters > Cluster Manager*.
 - 10b Select the check box next to the resource, then click *Offline*. Wait for the status to report that it is offline, then continue.
 - 10c Select the check box next to the resource, then click *Online*.
- 11 Continue with [Section 12.3.3, “Viewing or Modifying the LVM Resource Scripts,” on page 205](#).

12.3.3 Viewing or Modifying the LVM Resource Scripts

You can customize the scripts by adding lines for other products that use a shared LVM volume group resource. Compare the generic script with the templates for those products to identify what lines need to be added or modified.

- 1 In iManager, select *Clusters > Cluster Manager*.
- 2 Browse to select the Cluster object of the cluster where you created the volume group cluster resource.
- 3 From the list of Cluster objects, click the name link to open the Properties page, then click the *Scripts* tab.

The *Scripts* tab opens to the load script.

- 4 On the *Load Script* page, view or modify the load script. Click *Apply* if you make changes.
See the [“Sample LVM Resource Load Script” on page 207](#).
- 5 Click the *Unload Script* link to view or modify the unload script. Click *Apply* if you make changes.
See the [“Sample LVM Resource Unload Script” on page 207](#).
- 6 Click the *Monitor Script* link to view or modify the monitor script. Click *Apply* if you make changes.
See the [“Sample LVM Resource Monitor Script” on page 208](#).
- 7 At the bottom of the page, click *OK* to close the Properties page and save your changes.
The changes do not take effect until the resource is taken offline and brought online again.
- 8 If you modified the settings on any of the pages, you must take the resource offline and bring it online in order for the changes to take effect.
 - 8a In Roles and Tasks, select *Clusters > Cluster Manager*.
 - 8b Select the check box next to the resource, then click *Offline*. Wait for the status to report that it is offline, then continue.
 - 8c Select the check box next to the resource, then click *Online*.
 - 8d Verify that the resource comes online and reports a *Running* state.
If the resource goes into a Comatose state, it is probably because you made a mistake in the lines you added or modified in the scripts. Take the resource offline, then go back to correct the scripts, and try to bring it online again.

12.3.4 Sample LVM Resource Scripts

This section contains sample scripts for the LVM resource. The sample values are based on the values used in NSSMU to create the generic LVM cluster resource in [Section 12.3.1, “Creating an LVM Volume Group Cluster Resource,” on page 200](#).

- ♦ [“Sample LVM Resource Load Script” on page 207](#)
- ♦ [“Sample LVM Resource Unload Script” on page 207](#)
- ♦ [“Sample LVM Resource Monitor Script” on page 208](#)

Sample LVM Resource Load Script

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name (nssmu uses volume name for group name)
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# activate the volume group
exit_on_error vgchange -a ey $VOLGROUP_NAME

# create the mount point if it does not exist on the node
ignore_error mkdir -p $MOUNT_POINT

# mount the file system
exit_on_error mount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

exit 0
```

Sample LVM Resource Unload Script

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name (nssmu uses volume name for group name)
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# unmount the volume
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error umount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# deactivate the volume group
exit_on_error vgchange -a n $VOLGROUP_NAME

exit 0
```

Sample LVM Resource Monitor Script

```
#!/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 volume group name
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# check the logical volume
exit_on_error status_lv $MOUNT_DEV

# test the file system
exit_on_error status_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# status the IP address
exit_on_error status_secondary_ipaddress $RESOURCE_IP

exit 0
```

12.4 Configuring an LVM Cluster Resource with LVM Commands and the Generic File System Template

This section describes how to use Linux Logical Volume Manager (LVM) commands to create a shared LVM volume group. You use the Generic File System template (`Generic_FS_Template`) to create a volume group cluster resource that cluster-enables the existing shared volume group.

After you create the resource, you can add lines to its load script, unload script, and monitor script to customize the resource for other uses. Compare the `Generic_FS_Template` to the resource template for your product to determine which lines need to be added or modified.

- ♦ [Section 12.4.1, “Creating a Shared LVM Volume with LVM Commands,” on page 209](#)
- ♦ [Section 12.4.2, “Creating a Generic File System Cluster Resource for an LVM Volume Group,” on page 213](#)
- ♦ [Section 12.4.3, “Sample LVM Resource Load Scripts,” on page 218](#)

12.4.1 Creating a Shared LVM Volume with LVM Commands

You can create a shared LVM volume group and logical volume by using native Linux LVM commands. Before you begin, initialize the SAN device as described in [Section 12.2, “Initializing a SAN Device,”](#) on page 197.

The procedures in this section uses the following sample parameters. Ensure that you replace the sample values with your values.

Parameter	Sample Value
LVM physical volume	<code>/dev/sdd</code> For clustering, we recommend only one device in the LVM volume group.
LVM volume group name	<code>myclustervg01</code>
LVM logical volume	<code>myclustervol01</code>
File system type	<code>ext3</code> This is the file system type that you make on the LVM logical volume, such as <code>ext2</code> , <code>ext3</code> , <code>reiserfs</code> , or <code>xfs</code> . For information about <code>mkfs</code> command options for each file system type, see the <code>mkfs(8)</code> man page.
Logical volume path	<code>/dev/myclustervg01/myclustervol01</code>
Mount point for the logical volume	<code>/mnt/myclustervol01</code>

- ♦ [“Creating a Shared LVM Volume Group \(Quick Reference\)”](#) on page 210
- ♦ [“Creating a Shared LVM Volume Group \(Detailed Instructions\)”](#) on page 210

Creating a Shared LVM Volume Group (Quick Reference)

You can create the volume group and logical volume by issuing the following LVM commands as the `root` user on the cluster node. When you are done, continue with [Section 12.4.2, “Creating a Generic File System Cluster Resource for an LVM Volume Group,”](#) on page 213.

Command Action	Command
Create the LVM physical volume.	<code>pvcreate <device></code>
Create the clustered LVM volume group.	<code>vgcreate -c y <vg_name> <device></code>
Activate the volume group exclusively on the node.	<code>vgchange -a ey <vg_name></code>
Create the LVM logical volume.	<code>lvcreate -n <lv_name> -L size <vg_name></code>
Add a file system to the LVM logical volume.	<code>mkfs -t <fs_type> /dev/<vg_name>/<lv_name> [fs_options]</code>
If it does not exist, create the mount point path.	<code>mkdir -p <full_mount_point_path></code>
Mount the LVM logical volume.	<code>mount -t <fs_type> /dev/<vg_name>/<lv_name> <mount_point></code>
Dismount the LVM logical volume and deactivate the volume group. This allows the cluster resource to control when the volume is available.	<code>umount /dev/<vg_name>/<lv_name> vgchange -a n <vg_name></code>

Creating a Shared LVM Volume Group (Detailed Instructions)

For detailed instructions, use the following procedure to create the LVM volume group and logical volume:

- 1 Log in as the Linux `root` user to the first node of the cluster, then open a terminal console.
- 2 Create an LVM physical volume on the device (such as `/dev/sdd`) by entering:

```
pvcreate <device>
```

For example:

```
pvcreate /dev/sdd
No physical volume label read from /dev/sdd
Physical volume "/dev/sdd" successfully created
```

- 3 Display information about the physical volume by entering:

```
pvdiskdisplay [device]
```

For example, to view information about /dev/sdd:

```
pvddisplay /dev/sdd
"/dev/sdd" is a new physical volume of "512 MB"
--- NEW Physical volume ---
PV Name                /dev/sdd
VG Name
PV Size                512 MB
Allocatable            NO
PE Size (KByte)        0
Total PE               0
Free PE                0
Allocated PE           0
PV UUID                dg5WG9-IJvq-MkEM-iRrC-hMKv-Zmwq-YY0N3m
```

- 4 Create an LVM volume group (such as myclustervg01) on the physical volume by entering:

```
vgcreate -c y <vg_name> <device>
```

For example:

```
vgcreate -c y "myclustervg01" /dev/sdd
Clustered volume group "myclustervg01" successfully created
```

The volume group is automatically activated.

- 5 Activate the volume group exclusively on the current server by entering:

```
vgchange -a ey <vg_name>
```

The -a option activates the volume. The ey parameter specifies the values exclusively and yes.

For example:

```
vgchange -a ey myclustervg01
logical volume(s) in volume group "myclustervg01" now active
```

- 6 View information about the volume group by using the vgdisplay command.

```
vgdisplay <vg_name>
```

Notice that 4 MB of the device is used for the volume group's Physical Extent (PE) table. You must consider this reduction in available space on the volume group when you specify the size of the LVM logical volume in the next step ([Step 7](#)).

For example:

```
vgdisplay myclustervg01
--- Volume group ---
VG Name                myclustervg01
System ID
Format                 lvm2
Metadata Areas         1
Metadata Sequence No   1
VG Access               read/write
VG Status               resizable
MAX LV                 0
Cur LV                 0
Open LV                 0
Max PV                 0
Cur PV                 1
Act PV                 1
VG Size                 508.00 MB
PE Size                 4.00 MB
Total PE                127
Alloc PE / Size        0 / 0
Free PE / Size          127 / 508.00 MB
VG UUID                rqyAd3-U2dg-HYLw-0SyN-1o07-jBH3-qHvySe
```

- 7** Create an LVM logical volume (such as myclustervol01) on the volume group by entering:

```
lvcreate -n <lv_name> -L size <vg_name>
```

Specify the logical volume name, size, and the name of the volume group where you want to create it. The size is specified in megabytes by default.

The logical volume full path name is /dev/<vg_name>/<lv_name>.

For example:

```
lvcreate -n "myclustervol01" -L 500 "myclustervg01"
Logical volume "myclustervol01" created
```

This volume's full path name is /dev/myclustervg01/myclustervol01.

- 8** View information about the logical volume by entering:

```
lvdisplay -v /dev/<vg_name>/<lv_name>
```

For example:

```
lvdisplay -v /dev/myclustervg01/myclustervol01
Using logical volume(s) on command line
--- Logical volume ---
LV Name                /dev/myclustervg01/myclustervol01
VG Name                myclustervg01
LV UUID                nIfsMp-aLRr-i4Lw-Wwdt-v5io-2hDN-qrWTLH
LV Write Access        read/write
LV Status              available
# open                 0
LV Size                500.00 MB
Current LE             125
Segments              1
Allocation             inherit
Read ahead sectors     auto
- currently set to    1024
Block device           253:1
```

- 9** Create a file system (such as Ext2, Ext3, ReiserFS, or XFS) on the LVM logical volume by entering:

```
mkfs -t <fs_type> /dev/<vg_name>/<lv_name> [fs_options]
```

For example:

```
mkfs -t ext3 /dev/myclustervg01/myclustervol01
mke2fs 1.41.9 (22-Aug-2009)
Filesystem label=
OS type: Linux
Block size=1024 (log=0)
Fragment size=1024 (log=0)
128016 inodes, 512000 blocks
25600 blocks (5.00%) reserved for the super user
First data block=1
Maximum filesystem blocks=67633152
63 block groups
8192 blocks per group, 8192 fragments per group
2032 inodes per group
Superblock backups stored on blocks:
    8193, 24577, 40961, 57345, 73729, 204801, 221185, 401409

Writing inode tables: done
Creating journal (8192 blocks): done
Writing superblocks and filesystem accounting information: done

This filesystem will be automatically checked every 29 mounts or
180 days, whichever comes first.  Use tune2fs -c or -i to override.
```

- 10** If the mount point does not exist, create the full directory path for the mount point.

```
mkdir -p <full_mount_point_path>
```

For example:

```
mkdir -p /mnt/myclustervol01
```

- 11** Mount the logical volume on the desired mount point by entering:

```
mount -t <fs_type> /dev/<vg_name>/<lv_name> <mount_point>
```

For example:

```
mount -t ext3 /dev/myclustervg01/myclustervol01 /mnt/myclustervol01
```

- 12** Dismount the volume and deactivate the LVM volume group by entering:

```
umount /dev/<vg_name>/<lv_name>  
vgchange -a n "vg_name"
```

This allows you to use the load script and unload script to control when the volume group is activated or deactivated in the cluster.

For example, to unmount the volume and deactivate the myclustervg01 volume group, enter

```
umount /dev/myclustervg01/myclustervol01  
vgchange -a n "myclustervg01"  
0 logical volume(s) in volume group "myclustervg01" now active
```

- 13** Continue with [Section 12.4.2, “Creating a Generic File System Cluster Resource for an LVM Volume Group,”](#) on page 213.

12.4.2 Creating a Generic File System Cluster Resource for an LVM Volume Group

You use the Generic File System template (Generic_FS_Template) to create a cluster resource for the LVM volume group.

The sample scripts in this section use the following sample parameters. Ensure that you replace the sample values with your values.

Parameter	Sample Value
RESOURCE_IP	10.10.189.136
MOUNT_FS	ext3
VOLGROUP_NAME	myclustervg01
MOUNT_DEV	/dev/\$VOLGROUP_NAME/myclustervol01
MOUNT_POINT	/mnt/myclustervol01

To configure the generic file system resource for the LVM volume group:

- 1 In iManager, select *Clusters > Cluster Options*.
- 2 Browse to locate and select the Cluster object of the cluster you want to manage.

- 3 Click the *New* link.
- 4 Click the *Resource* radio button, then click *Next*.

Clusters > Cluster Options

New Resource ?

Resource Type Select the type of cluster resource to create.

Type

☐ Pool

☒ Resource

☐ Template

- 5 Specify information to define your new cluster resource:

- 5a Specify the name of the resource you want to create, such as `myclustervg01`.

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.

- 5b In the *Inherit From Template* field, browse to select the *Generic_FS_Template*.

Clusters > Cluster Options

New Resource ?

Cluster Resource Information Create a new cluster resource or cluster resource template.

Cluster Resource Name:

Inherit From Template: 

☐ Online Resource after Create

☒ Define Additional Properties

- 5c Do not select *Online Resource after Create*.
 - You want to define the scripts and set the resource policies before you bring the resource online for the first time.
 - 5d Select the *Define Additional Properties* check box.
 - 5e Click *Next*.
- 6 Configure the load script for the resource by replacing the variables with your own values, specify the *Load Script Timeout* value, then click *Next*.

The following is the default Generic_FS template load script. The variable values appear in bold font. A sample script is available in [“Sample LVM Resource Load Script” on page 219](#). You can add a `mkdir` command for the mount point, or create the directory manually on each node before you allow the resource to fail over to other nodes.

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

# define the IP address
RESOURCE_IP=a.b.c.d
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=name
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/volume_name
# define the mount point
MOUNT_POINT=/mnt/mount_point

# activate the volume group
exit_on_error vgchange -a ey $VOLGROUP_NAME

# create the mount point on the node if it does not exist ignore_error mkdir -
p $MOUNT_POINT

# mount the file system
exit_on_error mount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

exit 0
```

- 7 Configure the unload script for the resource by replacing the variables with your own values, specify the *Unload Script Timeout* value, then click *Next*.

The following is the default Generic_FS template unload script. The variable values appear in bold font. A sample script is available in [“Sample LVM Resource Unload Script” on page 219](#).

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

# define the IP address
RESOURCE_IP=a.b.c.d
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=name
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/volume_name
# define the mount point
MOUNT_POINT=/mnt/mount_point

# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# unmount the volume
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error umount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# deactivate the volume group
exit_on_error vgchange -a n $VOLGROUP_NAME

# return status

exit 0
```

- 8 Configure the monitor script for the resource by replacing the variables with your own values, specify the *Monitor Script Timeout* value, then click *Next*.

The following is the default Generic_FS template monitor script. The variable values appear in bold font. A sample script is available in “[Sample LVM Resource Monitor Script](#)” on page 220.

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

# define the IP address
RESOURCE_IP=a.b.c.d
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=name
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/volume_name
# define the mount point
MOUNT_POINT=/mnt/mount_point

# check the logical volume
exit_on_error status_lv $MOUNT_DEV

# test the file system
exit_on_error status_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# status the IP address
exit_on_error status_secondary_ipaddress $RESOURCE_IP

exit 0
```

- 9 On the Resource Policies page, view and modify the resource’s Policy settings:

- 9a (Optional) Select the *Resource Follows Master* check box if you want to ensure that the resource runs only on the master node in the cluster.

If the master node in the cluster fails, the resource fails over to the node that becomes the master.

- 9b (Optional) Select the *Ignore Quorum* check box if you don't want the cluster-wide timeout period and node number limit enforced.

The quorum default values were set when you installed 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.

- 9c By default, the Generic File System resource template sets the Start mode and Failover mode to *Auto* and the Failback Mode to *Disable*. You can change the default settings as needed.

- ♦ **Start Mode:** If the Start mode is set to *Auto*, the resource automatically loads on a designated server when the cluster is first brought up. If the Start mode is set to *Manual*, you can manually start the resource on a specific server when you want, instead of having it automatically start when servers in the cluster are brought up.
- ♦ **Failover Mode:** If the Failover mode is set to *Auto*, the resource automatically moves to the next server in the Assigned Nodes list if there is a hardware or software failure. If the Failover mode is set to *Manual*, you can intervene after a failure occurs and before the resource is started on another node.
- ♦ **Failback Mode:** If the Failback mode is set to *Disable*, the resource continues running on the node it has failed to. If the Failback mode is set to *Auto*, the resource automatically moves back to its preferred node when the preferred node is brought

back online. Set the Failback mode to *Manual* to prevent the resource from moving back to its preferred node when that node is brought back online, until you are ready to allow it to happen.

9d Click *Next*.

- 10** On the Preferred Nodes page, assign preferred nodes for the resource by moving them from the *Unassigned* list to the *Assigned* list.

When you bring a resource online, it is automatically loaded on the most preferred node in the list. If the node is not available, the other nodes are tried in the order that they appear in the list. You can modify the order of the nodes by clicking the Edit (pen) icon to open the list a text editor. In the editor, click *OK* to close the editor.

Clusters > Cluster Manager

Cluster Resource Properties: vol44_resource.clus1.ncs.novell

Policies Monitoring Preferred Nodes Scripts Business Continuity

View or change the preferred nodes for this cluster resource.

Assigned:

- avalon
- london
- newcastle
- cardiff

Unassigned:

OK Cancel Apply

- 11** Click *Finish*.

- 12** Verify that the resource was created by viewing its entry in the *Cluster Objects* list on the Cluster Options page.

Clusters

Cluster Options

Cluster: cluster1.clusters.city.mycompany

Properties... Repair...

View cluster resource configuration information and administer cluster resources for the selected cluster.

Type	Name	IP Address	Distinguished Name	Pool Name
Master IP Address Resource	Master IP Address Resource	10.10.10.44	cn=cluster1,ou=clusters,ou=city,o=mycompany	
avalon	avalon	10.10.10.37	cn=avalon,cn=cluster1,ou=clusters,ou=city,o=mycompany	
POOLC1_SERVER	POOLC1_SERVER	10.10.10.41	cn=CLUSTER1_POOLC1_SERVER,ou=clusters,ou=city,o=mycompany	POOLC1
myclustervg01	myclustervg01			

- 13** (Optional) Enable and configure resource monitoring:

For information, see [Section 9.6, “Enabling Monitoring and Configuring the Monitor Script,”](#) on page 135.

13a Select the resource name link to open its Properties page, then select the *Monitoring* tab.

13b Select the *Enable Resource Monitoring* check box to enable resource monitoring for the selected resource.

Resource monitoring is disabled by default.

- 13c** Specify the *Polling Interval* to control how often you want the resource monitoring script for this resource to run.

You can specify the value in minutes or seconds. For information, see [“Polling Interval” on page 136](#).

- 13d** Specify the number of failures (*Maximum Local Failures*) for the specified amount of time (*Time Interval*).

For information, see [“Failure Rate” on page 136](#).

- 13e** Specify the *Failover Action* by indicating whether you want the resource to be set to a comatose state, to migrate to another server, or to reboot the hosting node (without synchronizing or unmounting the disks) if a failure action initiates. The reboot option is normally used only for a mission-critical cluster resource that must remain available.

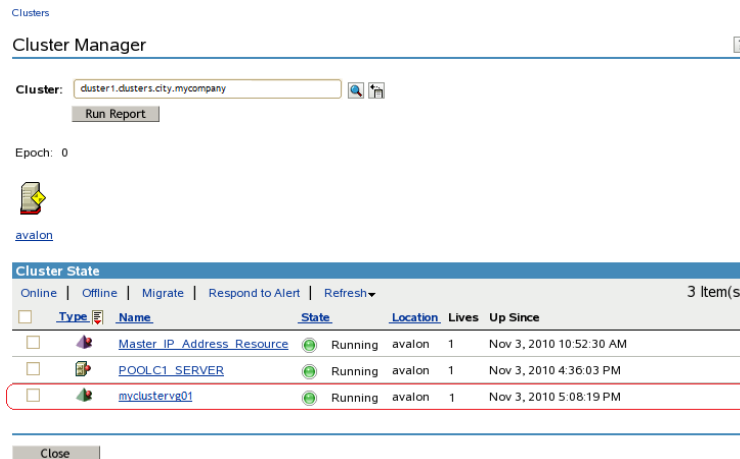
For information, see [“Failure Action” on page 136](#).

- 14** Bring the resource online:

- 14a** Select *Clusters > Cluster Manager*.

- 14b** Select the check box next to the new CSM cluster resource, then click *Online*.

Ensure that the resource successfully enters the *Running* state. If the resource goes comatose instead, it is probably because you made an error when typing values in the script definitions. Take the resource offline, go to the resource’s *Properties > Scripts* page to review and modify its scripts as needed to correct errors, then try again to online the resource.



12.4.3 Sample LVM Resource Load Scripts

This section contains sample scripts for the LVM resource. The sample values are based on the values used to create the resource with LVM commands and the Generic File System template in [Section 12.4.2, “Creating a Generic File System Cluster Resource for an LVM Volume Group,” on page 213](#).

- ♦ [“Sample LVM Resource Load Script” on page 219](#)
- ♦ [“Sample LVM Resource Unload Script” on page 219](#)
- ♦ [“Sample LVM Resource Monitor Script” on page 220](#)

Sample LVM Resource Load Script

Use the following sample load script to complete the fields for your LVM volume group cluster resource:

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

# define the IP address
RESOURCE_IP=10.10.189.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=myclustervg01
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/myclustervol01
# define the mount point
MOUNT_POINT=/mnt/myclustervol01

# activate the volume group
exit_on_error vgchange -a ey $VOLGROUP_NAME

# create the mount point on the node if it does not exist
ignore_error mkdir -p $MOUNT_POINT

# mount the file system
exit_on_error mount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

exit 0
```

Sample LVM Resource Unload Script

Use the following sample unload script to complete the fields for your LVM volume group cluster resource:

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

# define the IP address
RESOURCE_IP=10.10.189.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=myclustervg01
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/myclustervol01
# define the mount point
MOUNT_POINT=/mnt/myclustervol01

# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# unmount the volume
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error umount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# deactivate the volume group
exit_on_error vgchange -a n $VOLGROUP_NAME

exit 0
```

Sample LVM Resource Monitor Script

Use the following sample monitor script to complete the fields for your LVM volume group cluster resource:

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

# define the IP address
RESOURCE_IP=10.10.189.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=myclustervg01
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/myclustervol01
# define the mount point
MOUNT_POINT=/mnt/myclustervol01

# check the logical volume
exit_on_error status_lv $MOUNT_DEV

# test the file system
exit_on_error status_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# status the IP address
exit_on_error status_secondary_ipaddress $RESOURCE_IP

exit 0
```

12.5 Creating a Virtual Server Object for an LVM Volume Group Cluster Resource

When you cluster-enable the Linux LVM volume group, a virtual server object is not created automatically as it is when you cluster-enable an NSS pool. You can assign a virtual server name to the cluster resource by using the `/opt/novell/ncs/bin/ncs_ncpserv.py` script to create an NCS:NCP Server object for it. Having a virtual server object allows the Linux POSIX volume cluster resource to be viewed in the *Browse* panel in iManager. You can also bind the virtual server name to the IP address to allow users to access the resource via an assigned name in addition to its IP address.

The NCS:NCP Server object does not give users NCP access to the data on a Linux POSIX volume. An NCP volume is required to do that.

IMPORTANT: If you are creating NCP volumes on the shared Linux POSIX volume, do not follow the procedure in this section. Go to [“Creating a Shared NCP Volume on the Linux POSIX Cluster Resource”](#) in the *OES 11: NCP Server for Linux Administration Guide*, then follow the instructions there.

The virtual server name is stored in Novell eDirectory in an NCS:NCP Server object under the Cluster object where you created the resource. You must add a line to the load and unload scripts that identifies the name of this virtual server and a line that binds or unbinds the name to the IP address of the Linux POSIX cluster resource.

- ♦ [Section 12.5.1, “Creating the Virtual Server Object,” on page 221](#)
- ♦ [Section 12.5.2, “Modifying the Load Script,” on page 221](#)
- ♦ [Section 12.5.3, “Modifying the Unload Script,” on page 222](#)
- ♦ [Section 12.5.4, “Activating the Script Changes,” on page 222](#)
- ♦ [Section 12.5.5, “Verifying the Virtual Server,” on page 222](#)

12.5.1 Creating the Virtual Server Object

You use the `/opt/novell/ncs/bin/ncs_ncpserv.py` script to create a virtual server object (NCS:NCP Server) in eDirectory for the LVM volume group cluster resource. If the resource does not have NCP volumes on it, you do not use the `-v` option. For information about the `ncs_ncpserv.py` script, see [Section A.6, “Virtual NCS NCP Server Object \(ncs_ncpserv.py\) Script,” on page 286](#).

- 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 vg_name -i resource_ip_address
```

Replace the `vg_name` and `resource_ip_address` with the information for your particular solution.

Do not use periods in cluster resource names. 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, to create the NCS:NCP Server object for the `myclustervg01` cluster resource where the IP address is 10.10.10.44 and the cluster context is `ou=clusters,ou=city,o=mycompany`, enter

```
./ncs_ncpserv.py -c myclustervg01 -i 10.10.10.44
```

The confirmation message is displayed:

```
NCP Server 'cn=cluster1_myclustervg01_server,ou=clusters,ou=city,o=mycompany'
created.
```

- 4 Continue with [Section 13.6.2, “Modifying the Load Script,” on page 246](#).

12.5.2 Modifying the Load Script

After you have created an NCS:NCP Server object, you must modify the load script so that it binds the NCS:NCP Server object to the resource.

- 1 In iManager, select *Clusters > Cluster Options*, then select the cluster.
- 2 Click the name link of the LVM volume group cluster resource to open its Cluster Resource Properties page.
- 3 Click the *Scripts* tab to open the *Load* script.
- 4 In the definition area, add the following lines to define the virtual NCP server name:

```
# define NCP server name
NCP_SERVER=cluster1_myclustervg01_server
```

Replace the NCP server name with the name for your virtual NCP server.

- 5 Under the mount line, add a line to bind the NCP server name to the resource IP address:

```
# bind the NCP server name
exit_on_error ncpcon bind --ncpservname=$NCP_SERVER --ipaddress=$RESOURCE_IP
```

- 6 Click *Apply* to save your changes.

The script changes are not active until the next time the cluster resource is taken offline, and then brought online. Do not active the script changes at this time.

- 7 Continue with [Section 13.6.3, “Modifying the Unload Script,” on page 247](#).

12.5.3 Modifying the Unload Script

After you have created an NCS:NCP Server object, you must modify the unload script so that it unbinds the NCS:NCP Server object from the resource.

- 1 In iManager, select *Clusters > Cluster Options*, then select the cluster.
- 2 Click the name link of the LVM volume group cluster resource to open its Cluster Resource Properties page.
- 3 Click the *Scripts* tab, then click *Unload* to open the *Unload* script.
- 4 In the definition area, add the following lines to define the virtual NCP server name:

```
# define NCP server name
NCP_SERVER=cluster1_myclustervg01_server
```

Replace the NCP server name with the name for your virtual NCP server. Use the same value for variable that you did in the load script.

- 5 Under the definition, add a line to unbind the NCP server name from the resource IP address:

```
# unbind the NCP server name
ignore_error ncpcon unbind --ncpservname=$NCP_SERVER --
ipaddress=$RESOURCE_IP
```

- 6 Click *Apply* to save your changes.
The script changes are not active until the next time the cluster resource is taken offline, and then brought online.
- 7 Continue with [Section 13.6.4, “Activating the Script Changes,” on page 247](#).

12.5.4 Activating the Script Changes

The script changes are not active until the next time the cluster resource is taken offline, and then brought online.

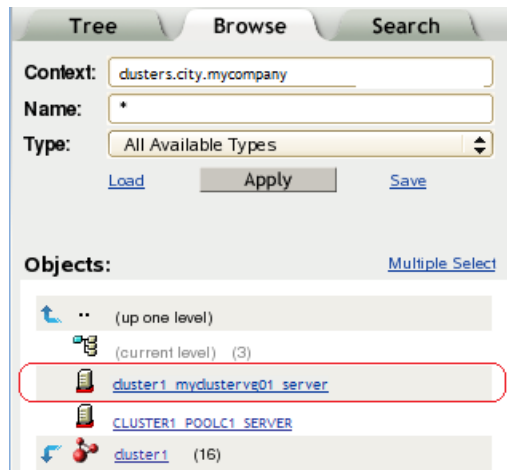
- 1 In iManager, select *Clusters > Cluster Manager*, then select the cluster.
- 2 Select the check box next to the LVM volume group cluster resource, then click *Offline*.
Wait until the resource is reports an Offline status before continuing.
- 3 Select the check box next to the LVM volume group cluster resource, then click *Online*.
Wait until the resource is reports an Online status before continuing.
- 4 Continue with [Section 13.6.5, “Verifying the Virtual Server,” on page 247](#).

12.5.5 Verifying the Virtual Server

You can verify that an NCS:NCP Server object appears in the *Browse* panel in iManager.

- 1 In the iManager toolbar, click the *View Objects* icon.
- 2 In the left panel, click *Browse*.

- 3 Browse to the Cluster container to see the virtual server object for the cluster resource, such as `cluster1_myclustervg01_server`.



12.6 Renaming the Mount Point Path for a Clustered LVM Volume

For a clustered LVM volume, you can rename the mount point path by modifying the mount point variable in the cluster load, unload, and monitor scripts.

- 1 Open Novell iManager in a Web browser, then log in as an administrator user.
- 2 In *Roles and Tasks*, select *Clusters > Cluster Manager*.
- 3 Browse to select the Cluster object of the cluster where you created the volume group cluster resource.
- 4 Select the check box next to the resource, then click *Offline*. Wait for the status to report that it is offline, then continue.
- 5 Modify the mount point path value in the load, unload, and monitor scripts for the LVM volume group cluster resource:

- 5a On the Cluster Manager page, select the resource's name link to open its Cluster Properties page, then click the *Scripts* tab.

The Scripts tab automatically displays the load script.

- 5b Modify the load script:

- 5b1 In the load script, type the new value for the mount point in the `MOUNT_POINT` variable:

```
MOUNT_POINT=/media/ext3/vol44
```

- 5b2 Ensure that the following command is added above the mount command line in the load script in order to create the path on nodes if it does not exist.

```
# create the mount point if it does not exist on the node ignore_error
mkdir -p $MOUNT_POINT
```

You can alternatively make the new path by using the `mkdir` command in a terminal console on each node. If the master node is not the most preferred node, ensure that you make the path before you bring the resource online.

- 5b3 Click *Apply*.

- 5c** Click the *Unload Script* link, type the new value for the mount point in the MOUNT_POINT variable, then click *Apply*.

```
MOUNT_POINT=/media/ext3/vol44
```

- 5d** Click the *Monitor Script* link, type the new value for the mount point in the MOUNT_POINT variable, then click *Apply*.

```
MOUNT_POINT=/media/ext3/vol44
```

- 5e** At the bottom of the page, click *OK* to close the Properties page and save your changes.
The changes do not take effect until the resource is brought online.

- 6** Bring the resource online in order for the script changes to take effect.

- 6a** In *Roles and Tasks*, select *Clusters > Cluster Manager*.

- 6b** Select the check box next to the resource, then click *Online*.

- 6c** Verify that the resource comes online and reports a *Running* state.

If the resource goes into a Comatose state, it is probably because you made a mistake in the lines you added or modified in the scripts. Take the resource offline, then go back to correct the scripts, and try to bring it online again.

- 7** In NSSMU, verify that the new mount point is used when the clustered LVM volume resource is brought online:

- 7a** Log in as the root user to the node that is hosting the resource, then start NSSMU by entering:

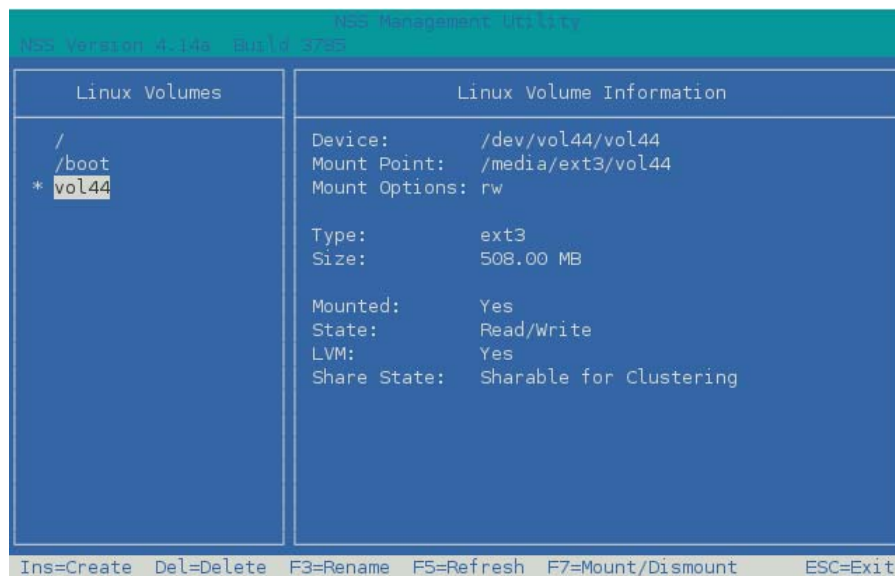
```
nssmu
```

- 7b** From the NSSMU main menu, select *Linux Volumes*, then press Enter.

- 7c** In the *Linux Volumes* list, select the clustered LVM volume.

- 7d** View the volume details to verify that the mount point has changed.

For example, the mount point is /media/ext3/vol44.



- 7e** Click Escape twice to exit NSSMU.

12.7 Disabling Clustering for an LVM Volume

Use the procedure in this section if you want to disable clustering for an LVM volume. Afterwards, you can mount and dismount the volume only as a local LVM volume.

- 1 Delete the LVM volume group cluster resource:
 - 1a Log in to Novell iManager as a cluster administrator user.
 - 1b Select *Clusters > Cluster Manager*, then browse to select the Cluster object.
 - 1c On the Cluster Manager page, select the check box next to the LVM volume group cluster resource (such as `vol44_resource`), then click *Offline*. Wait until the resource is offline to continue.
 - 1d Select *Clusters > Cluster Options*.
 - 1e On the Cluster Options page, select the check box next to the LVM volume group cluster resource, click *Delete*, then click *OK* to confirm.
 - 1f Exit iManager.

- 2 Log in as the `root` user on the server where the resource was online, then open a terminal console.

- 3 Remove the shareable state from the clustered LVM volume group. Enter

```
vgchange -c n <vg_name>
```

For example:

```
vgchange -c n vol44
```

- 4 Activate the volume group on the current node. Enter

```
vgchange -a ey <vg_name>
```

For example:

```
vgchange -a ey vol44
```

- 5 Mount the LVM logical volume. Enter

```
mount -t <fstype> <device_name> <full_mount_point_path>
```

Replace *fstype* with the file system type of the volume.

Replace *device_name* with the full device path of the logical volume, such as `/dev/<vg_name>/<lv_name>`.

Replace *full_mount_point_path* with the volume's mount point.

For example, enter

```
mount -t ext3 /dev/vol44/vol44 /media/ext3/vol44
```

- 6 In a text editor, modify the `/etc/fstab` file to specify the mount point information and file system type.

This entry allows the volume to be mounted automatically on reboot. It also provides the automatic mount information for NSSMU to use to mount and dismount the volume locally.

For example, complete the line for the volume's device path, such as:

```
/dev/vol44/vol44    /media/ext3/vol44    ext3    rw    0 0
```

File	Mount Point	Filesystem	Options	Dump	Pass
/dev/sda2	swap	swap	defaults	0	0
/dev/sda3	/	ext3	acl,user_xattr	1	1
/dev/sda1	/boot	ext2	acl,user_xattr	1	2
proc	/proc	proc	defaults	0	0
sysfs	/sys	sysfs	noauto	0	0
debugfs	/sys/kernel/debug	debugfs	noauto	0	0
usbfs	/proc/bus/usb	usbfs	noauto	0	0
devpts	/dev/pts	devpts	mode=0620,gid=5	0	0
/dev/vol44/vol44	/media/ext3/vol44	ext3	rw	0	0

- 7 In NSSMU, go to the Linux Volumes page, then select the volume to view its details.
The volume is no longer cluster enabled. You can mount and dismount the volume only on the current node.

Linux Volumes	Linux Volume Information
/	Device: /dev/vol44/vol44
/boot	Mount Point: /media/ext3/vol44
* vol44	Mount Options: rw
	Type: ext3
	Size: 508.00 MB
	Mounted: No
	State: Read/Write
	LVM: Yes
	Share State: Not Sharable for Clustering

Ins=Create Del=Delete F3=Rename F5=Refresh F7=Mount/Dismount ESC=Exit

12.8 Deleting a Clustered LVM Volume Group and Logical Volume

Before you can delete a clustered LVM volume group, you must take the volume group cluster resource offline.

- 1 Log in to the server as the root user, then open a terminal console.
- 2 Take the resource offline by entering

```
cluster offline <resource_name>
```

For example, enter

```
cluster offline vol44_resource
```

- 3 Delete the LVM logical volume, the clustered LVM volume group, and the Cluster Resource objects, then re-initialize the device:
 - 3a In the terminal console, launch NSSMU by entering

nssmu

- 3b** In the NSSMU main menu, select *Linux Volumes*, then press Enter.
 - 3c** In the *Linux Volumes* list, select the clustered linux volume, then press Delete.
Deleting the volume and volume group sets the device in an uninitialized state.
 - 3d** When you are prompted to confirm the delete action, read the warning message, then press Y (yes) to continue, or press N (no) to cancel the delete action.
 - 3e** Press Esc to return to the main menu, then select *Devices* and press Enter.
 - 3f** In the *Devices* list, select the device, then press F3 to initialize it.
 - 3g** When you are prompted to confirm the initialization action, read the warning message, then press Y (yes) to continue, or press N (no) to cancel the action.
 - 3h** Press Esc twice to exit NSSMU.
- 4** Verify that the LVM volume group is deleted by entering the following the native LVM command:

`vgdisplay`

12.9 LVM Management Tools

LVM tools are available in the YaST Expert Partitioner and in command line commands.

- ♦ [Section 12.9.1, “Using the YaST Expert Partitioner,” on page 227](#)
- ♦ [Section 12.9.2, “Using LVM Commands,” on page 227](#)

12.9.1 Using the YaST Expert Partitioner

You can access the LVM tools in the YaST Expert Partitioner by using the desktop menus, or enter the following in a terminal console:

`yast2 disk`

For information about using the Partitioner, see “LVM Configuration” (http://www.suse.com/documentation/sles11/stor_admin/data/lvm.html) in the *SLES 11 SP1: Storage Administration Guide* (http://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html).

12.9.2 Using LVM Commands

For information about using LVM commands, see the man pages for the commands described in [Table 12-1](#). Perform the commands as the `root` user.

Table 12-1 LVM Commands

Command	Description
<code>pvcreeate <device></code>	Initializes a device (such as <code>/dev/sdb</code>) for use by LVM as a physical volume.
<code>pvdisplay <device></code>	Displays information about the LVM physical volume, such as whether it is currently being used in a logical volume.

Command	Description
<code>vgcreate -c y <vg_name> <dev1> [dev2...]</code>	Creates a clustered volume group with one or more specified devices.
<code>vgchange -a [ey n] <vg_name></code>	<p>Activates (-a ey) or deactivates (-a n) a volume group and its logical volumes for input/output.</p> <p>IMPORTANT: Ensure that you use the ey option to exclusively activate a volume group on a cluster node. This option is used by default in the load script.</p>
<code>vgremove <vg_name></code>	Removes a volume group. Before using this command, remove the logical volumes, then deactivate the volume group.
<code>vgdisplay <vg_name></code>	<p>Displays information about a specified volume group.</p> <p>To find the total physical extent of a volume group, enter</p> <pre>vgdisplay vg_name grep "Total PE"</pre>
<code>lvcreate -L size -n <lv_name> <vg_name></code>	Creates a logical volume of the specified size.
<code>lvremove </dev/vg_name/lv_name></code>	<p>Removes a logical volume, such as /dev/vg_name/lv_name.</p> <p>Before using this command, close the logical volume by dismounting it with the <code>umount</code> command.</p>
<code>vgextend <vg_name><device></code>	Adds a specified physical volume to an existing volume group
<code>vgreduce <vg_name> <device></code>	<p>Removes a specified physical volume from an existing volume group.</p> <p>IMPORTANT: Ensure that the physical volume is not currently being used by a logical volume. If it is, you must move the data to another physical volume by using the <code>pvmove</code> command.</p>
<code>lvextend -L size </dev/vg_name/lv_name></code>	<p>Extends the size of a specified logical volume. Afterwards, you must also expand the file system to take advantage of the newly available space.</p>
<code>lvreduce -L size </dev/vg_name/lv_name></code>	<p>Reduces the size of a specified logical volume.</p> <p>IMPORTANT: Ensure that you reduce the size of the file system first before shrinking the volume, otherwise you risk losing data.</p>

13 Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers

Novell Cluster Services for Novell Open Enterprise Server (OES) 11 supports cluster resources for Linux POSIX volumes that you created with the Enterprise Volume Management System (EVMS) and the Cluster Segment Manager (CSM) on OES 2 servers. This section describes how to modify the scripts and cluster settings for the existing resources so they can run on OES 11 clusters.

IMPORTANT: For OES 11, Novell Cluster Services creates LVM volume group resources for Linux POSIX file systems. To create new resources, see [Chapter 12, “Configuring and Managing Cluster Resources for LVM Volume Groups,”](#) on page 195.

- ♦ [Section 13.1, “Requirements for Using CSM Cluster Resources on OES 11,”](#) on page 229
- ♦ [Section 13.2, “Modifying the Scripts for a CSM Resource,”](#) on page 230
- ♦ [Section 13.3, “Modifying the Scripts for a CSM Resource with a Segment Manager On It,”](#) on page 234
- ♦ [Section 13.4, “Configuring the Preferred Nodes and Cluster Settings for a CSM Cluster Resource,”](#) on page 243
- ♦ [Section 13.5, “Verifying the Load and Unload Scripts,”](#) on page 245
- ♦ [Section 13.6, “Creating a Virtual Server Object for a CSM Cluster Resource,”](#) on page 245
- ♦ [Section 13.7, “Deleting a Cluster Resource,”](#) on page 248

13.1 Requirements for Using CSM Cluster Resources on OES 11

On OES 2 clusters, a cluster resource for a shared Linux POSIX volume used the Enterprise Volume Management System (EVMS) utility to add a Cluster Segment Manager (CSM) on the device. Because EVMS is deprecated in SUSE Linux Enterprise Server 11, it is not available in OES 11.

In OES 11, Novell Cluster Services provides compatibility support for these CSM-based cluster resources. It has been modified in the following ways to allow you to move clustered Linux POSIX file systems from OES 2 to OES 11:

- ♦ Recognizes and supports the EVMS-based file system structure.
- ♦ Provides a way for managing a resource in the load, unload, and monitor scripts that works with the EVMS device and segment structure.
- ♦ Changes the device path from `/dev/evms/` to `/dev/mapper/`, which is the same location where LVM devices are stored.

The Cluster Segment Manager Import/Export (CSMPORT, `csmport`) utility provides script commands that make it possible to use the CSM resources in OES 11 clusters. The commands activate, deactivate, or check the status of the CSM container. You must add `csmport` commands in the load script, unload script, and monitor scripts before you bring the resource online on OES 11 servers in a mixed-mode cluster. For information about the utility commands that are used in the scripts, see [Section A.5, “Cluster Segment Manager Import/Export \(csmport\) Utility,” on page 280](#), or see the `csmport(8)` man page.

A mixed-mode cluster is supported as a temporary configuration scenario for converting an OES 2 cluster to OES 11. After the CSM-based resource is configured to run on OES 11, it should fail over only to OES 11 nodes in the mixed-mode cluster. For information, see [Section 6.1.2, “Linux POSIX File Systems on an EVMS CSM \(Compatibility Only\),” on page 86](#).

You cannot create new CSM-based cluster resources on OES 11 clusters. For new Linux POSIX volume resources, you must use the Linux Logical Volume Manager (LVM) volume groups as described in [Chapter 12, “Configuring and Managing Cluster Resources for LVM Volume Groups,” on page 195](#).

13.2 Modifying the Scripts for a CSM Resource

Before you attempt to migrate the CSM-based Linux POSIX Volume cluster resources from OES 2 to OES 11, you must take all of the CSM-based resources offline on the OES 2 nodes in the cluster. Modify their scripts to run on OES 11 nodes by adding `csmport` commands to activate, deactivate, or check the status of the CSM container. When you are done, the resources can be mounted successfully only on OES 11 nodes in the cluster.

IMPORTANT: If the CSM container for your cluster resource has a segment manager on it, follow the instructions in [Section 13.3, “Modifying the Scripts for a CSM Resource with a Segment Manager On It,” on page 234](#).

The sample scripts in this section use the following sample parameters. Ensure that you replace the sample values with your values.

Parameter	Sample Value
RESOURCE_IP	10.10.10.44
MOUNT_FS	ext3
CONTAINER_NAME	csm44
MOUNT_DEV	/dev/mapper/csm44
MOUNT_POINT	/mnt/lxvol44

IMPORTANT: Perform the following tasks only after you are ready to move the CSM resource to an OES 11 node in a mixed-mode cluster.

- [Section 13.2.1, “Offlining the CSM Cluster Resources,” on page 231](#)
- [Section 13.2.2, “Configuring the Scripts for the CSM Cluster Resource,” on page 231](#)
- [Section 13.2.3, “Sample Load Script for a CSM Resource,” on page 232](#)

- ♦ [Section 13.2.4, “Sample Unload Script for CSM Resource,” on page 233](#)
- ♦ [Section 13.2.5, “Sample Monitor Script for a CSM Resource,” on page 233](#)

13.2.1 Offlining the CSM Cluster Resources

Before any of the new OES11 nodes join an OES2 cluster, you must offline all resources that use Linux POSIX volumes.

- 1 In iManager, select *Clusters > 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 each of the CSM-based cluster resources, then click *Offline*.
- 4 Continue with [Section 13.2.2, “Configuring the Scripts for the CSM Cluster Resource,” on page 231](#).

13.2.2 Configuring the Scripts for the CSM Cluster Resource

- 1 In iManager, select *Clusters > Cluster Manager*.
- 2 Browse to locate and select the Cluster object of the cluster you want to manage.
- 3 Verify that the status of all CSM resources is *Offline*.
- 4 Select *Clusters > Cluster Options*.
- 5 Select the check box next to the CSM resource, then click the *Details* link to view its properties.
You can also click the CSM resource’s name link to view its properties.

- 6 Click the *Scripts* tab.
- 7 On the *Load Script* page, modify the script to add a `csmport` command to activate the CSM container, then click *Apply*.

```
#activate the container
exit_on_error csmport -i $CONTAINER_NAME
```

For an example, see [Section 13.2.3, “Sample Load Script for a CSM Resource,” on page 232](#).

- 8 Click the *Unload Script* link to go to the Upload Script page, modify the script to add a `csmport` command to deactivate the CSM container, then click *Apply*.

```
#deactivate the container
exit_on_error csmport -e $CONTAINER_NAME
```

For an example, see [Section 13.2.4, “Sample Unload Script for CSM Resource,” on page 233](#).

- 9 Click the *Monitor Script* link to go to the *Monitor Script* page, modify the script to add a `csmport` command to check the status of the CSM container, then click *Apply*.

```
#check the container
exit_on_error csmport -c $CONTAINER_NAME
```

For an example, see [Section 13.2.5, “Sample Monitor Script for a CSM Resource,” on page 233](#).

- 10 Continue with [Section 13.4, “Configuring the Preferred Nodes and Cluster Settings for a CSM Cluster Resource,” on page 243](#).

Do not bring the CSM cluster resources online again until after its *Preferred Nodes* list has been modified to use only OES 11 nodes.

13.2.3 Sample Load Script for a CSM Resource

Use the following sample load script to complete the fields for your CSM cluster resource on OES 11:

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the container name
CONTAINER_NAME=csm44
# define the device
MOUNT_DEV=/dev/mapper/$CONTAINER_NAME
# define the mount point
MOUNT_POINT=/mnt/lxvol44

# activate the container
exit_on_error csmport -i $CONTAINER_NAME

# 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
```

You can modify the `mount_fs` 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
```

If the path to the mount point does not exist on other nodes, you can add a line to create the mount point path in the script before the line that mounts the file system:

```
# create the mount point path when loading on a new node
ignore_error mkdir -p $MOUNT_POINT
```

13.2.4 Sample Unload Script for CSM Resource

Use the following sample unload script to complete the fields for your CSM cluster resource on OES 11:

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the container name
CONTAINER_NAME=csm44
# define the device
MOUNT_DEV=/dev/mapper/$CONTAINER_NAME
# define the mount point
MOUNT_POINT=/mnt/lxvol44

# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# unmount the volume
exit_on_error umount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# deactivate the container
exit_on_error csmport -e $CONTAINER_NAME

# return status
exit 0
```

13.2.5 Sample Monitor Script for a CSM Resource

Use the following sample monitor script to complete the fields for your CSM cluster resource on OES 11:

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the container name
CONTAINER_NAME=csm44
# define the device
MOUNT_DEV=/dev/mapper/$CONTAINER_NAME
# define the mount point
MOUNT_POINT=/mnt/lxvol44

# check the IP address
exit_on_error status_secondary_ipaddress $RESOURCE_IP

# check the volume
exit_on_error status_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# check the container
exit_on_error csmport -c $CONTAINER_NAME

# return status
exit 0
```

13.3 Modifying the Scripts for a CSM Resource with a Segment Manager On It

On OES 2, you might have added a segment manager on the CSM container in order to be able to create multiple volumes in the container as described in “[Adding a Segment Manager to the CSM Container](#)” in the *OES 2 SP3: Novell Cluster Services 1.8.8 Administration Guide for Linux*. This is not a typical configuration. If the CSM container has a segment manager on it, you must convert the volumes in the container to compatibility volumes and deport them before you offline them and modify their scripts for OES 11.

IMPORTANT: If the CSM container does not have a segment manager on it, follow the instructions in [Section 13.2, “Modifying the Scripts for a CSM Resource,”](#) on page 230.

The volume device name for each volume is based on its order in the container and whether the device is multipathed. The partitions for the volumes in the container are named by adding a sequential number to the end of the container name. For example:

```
/dev/mapper/<container_name>1
```

```
/dev/mapper/<container_name>2
```

and so on.

If the container name ends with a number, the partitions are named by adding a p before the sequential number. For example:

```
/dev/mapper/<container_name>p1
```

```
/dev/mapper/<container_name>p2
```

and so on.

You must also modify the definition fields in the scripts to have one entry for each of the volumes.

The sample scripts in this section use the following sample parameters. It assumes that the CSM container has a DOS segment manager on it and two Linux POSIX volumes with different file systems. Ensure that you replace the sample values with your values.

Parameter	Sample Value
volume name for Linux POSIX volume 1	lxvolext3
volume name Linux POSIX volume 2	lxvolxfs
RESOURCE_IP	10.10.10.44
MOUNT_FS1	ext3
MOUNT_FS2	xfs
CONTAINER_NAME	csm4
VOLUME_DEV1	csm44p1
VOLUME_DEV2	csm44p2
MOUNT_DEV1	/dev/mapper/\$VOLUME_DEV1
MOUNT_DEV2	/dev/mapper/\$VOLUME_DEV2
MOUNT_POINT1	/mnt/lxvol44ext3
MOUNT_POINT2	/mnt/lxvol44xfs

IMPORTANT: Perform the following tasks only after you are ready to move the CSM resource to an OES 11 node in a mixed-mode cluster.

- ♦ [Section 13.3.1, “Offlining the CSM Cluster Resources,” on page 235](#)
- ♦ [Section 13.3.2, “Converting Linux POSIX Volumes to Compatibility Volumes,” on page 236](#)
- ♦ [Section 13.3.3, “Deporting the CSM Container,” on page 236](#)
- ♦ [Section 13.3.4, “Offlining the CSM Cluster Resource,” on page 237](#)
- ♦ [Section 13.3.5, “Configuring the OES 11 NLVM Configuration File,” on page 237](#)
- ♦ [Section 13.3.6, “Configuring the Scripts for the CSM Cluster Resource,” on page 237](#)
- ♦ [Section 13.3.7, “Sample Load Script for a CSM Resource with a Segment Manager on It,” on page 240](#)
- ♦ [Section 13.3.8, “Sample Unload Script for a CSM Resource with a Segment Manager on It,” on page 242](#)
- ♦ [Section 13.3.9, “Sample Monitor Script for a CSM Resource with a Segment Manager on It,” on page 243](#)

13.3.1 Offlining the CSM Cluster Resources

Before any of the new OES 11 nodes join an OES 2 cluster, you must offline all resources that use Linux POSIX volumes.

- 1 In iManager, select *Clusters > 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 each of the CSM-based cluster resources, then click *Offline*.
- 4 Continue with [Section 13.2.2, “Configuring the Scripts for the CSM Cluster Resource,” on page 231](#).

13.3.2 Converting Linux POSIX Volumes to Compatibility Volumes

On the OES 2 server, convert each volume in the resource's container to a compatibility volume.

- 1 Log in as the `root` user on the OES 2 SP3 server in the cluster where the Linux POSIX volume cluster resource resides, then open a terminal console.
- 2 Online the Linux POSIX volume cluster resource by entering

```
cluster online <resource_name>
```
- 3 Launch the EVMSGUI utility by entering

```
evmsgui
```
- 4 Convert each Linux POSIX volume in the cluster resource to a compatibility volume:
 - 4a Select the *Volumes* tab.
 - 4b Right-click the volume and select *Convert to Compatibility Volume*.
 - 4c Click *Convert*.
 - 4d Click *Save* to save all of the changes.
 - 4e Repeat [Step 4a](#) through [Step 4d](#) for each volume in the cluster resource.
- 5 Continue with [Section 13.3.3, "Deporting the CSM Container,"](#) on page 236.

13.3.3 Deporting the CSM Container

On the OES 2 server, deport the Cluster Segment Manager container. If you are continuing the setup for a resource migration from OES 2 to OES 11, go to [Step 4](#).

- 1 Log in as the `root` user on the OES 2 server in the cluster where the Linux POSIX volume cluster resource resides, then open a terminal console.
- 2 Online the Linux POSIX volume cluster resource by entering

```
cluster online <resource_name>
```
- 3 Launch the EVMSGUI utility by entering

```
evmsgui
```
- 4 Modify the container's properties to set its type to *Deported*:
 - 4a Select the *Containers* tab.
 - 4b Right-click the container and select *Modify Properties*.
 - 4c Select the container, then click *Next*.
 - 4d Select *Type* and change it to *Deported*.
 - 4e Click *Modify*, then click *OK* to confirm the change.
 - 4f Click *Save* to save the container type change.
- 5 Continue with [Section 13.3.4, "Offlining the CSM Cluster Resource,"](#) on page 237.

13.3.4 Offlining the CSM Cluster Resource

- 1 Log in as the root user on the OES 2 server, then open a terminal console.
- 2 Offline the Linux POSIX volume cluster resource by entering

```
cluster offline <resource_name>
```
- 3 Continue with [Section 13.3.5, “Configuring the OES 11 NLVM Configuration File,”](#) on page 237.

13.3.5 Configuring the OES 11 NLVM Configuration File

Configure Novell Cluster Services on an OES 11 server, join the node to the cluster, then add a line for the CSM container in its `/etc/opt/novell/nss/nlvm.conf` file.

- 1 Prepare an OES 11 server for the cluster, but do not allow it to join the cluster.
For information, see [Step 2 in Section 6.4, “Adding OES 11 Nodes to an OES 2 Cluster \(Rolling Cluster Upgrade\),”](#) on page 88.
- 2 Log in as the root user on the OES 2 11 server, then open a terminal console.
- 3 Add the server to the OES 2 SP3 cluster by entering

```
cluster join
```
- 4 In a text editor, add the following line to the `/etc/opt/novell/nss/nlvm.conf` file, then save the file.

```
Include devices /dev/mapper/<containername>
```

Replace *containername* with the name of the CSM container for the Linux POSIX volume.
- 5 Continue with [Section 13.3.6, “Configuring the Scripts for the CSM Cluster Resource,”](#) on page 237.

13.3.6 Configuring the Scripts for the CSM Cluster Resource

- 1 In iManager, select *Clusters > 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 CSM resource, then click *Offline*.
- 4 Select *Clusters > Cluster Options*.
- 5 Select the check box next to the CSM resource, then click the *Details* link.
You can also click the CSM resource’s name link to view its properties.
- 6 Click the *Scripts* tab to view the Load Script page.

- 7** On the Load Script page, modify the script to handle the CSM container and multiple volumes on it.

For an example, see [Section 13.3.7, “Sample Load Script for a CSM Resource with a Segment Manager on It,”](#) on page 240.

- 7a** Add file system definition entries for each file system used by the volumes.

For example, if you have two volumes and each of them has a different type of file system, you create a definition for each:

```
#define the file system types
MOUNT_FS1=ext3
MOUNT_FS2=xf
```

- 7b** Add a device definition for each volume:

```
#define the volume devices
VOLUME_DEV1=cs
```

- 7c** Add a mount device definition for each volume:

```
#define the devices
MOUNT_DEV1=/dev/mapper/$VOLUME_DEV1
MOUNT_DEV2=/dev/mapper/$VOLUME_DEV2
```

- 7d** Add a mount point definition for each volume:

```
#define the mount points
MOUNT_POINT1=/mnt/lxvol
```

- 7e** Add a `csmpport` command to activate the CSM container:

```
#activate the container
exit_on_error csmpport -i $CONTAINER_NAME
```

- 7f** Add a `kpartx` command to activate the partitions. The command should follow the `csmpport` command in the load script.

```
#activate the partitions
exit_on_error /sbin/kpartx -a /dev/mapper/$CONTAINER_NAME
```

- 7g** If you use a `mkdir` command, create one for each mount point:

```
#if the mount path does not exist, create it
ignore_error mkdir -p $MOUNT_POINT1
ignore_error mkdir -p $MOUNT_POINT2
```

- 7h** Add a mount command for each volume:

```
#mount the file systems
exit_on_error mount_fs $MOUNT_DEV1 $MOUNT_POINT1 $MOUNT_FS1
exit_on_error mount_fs $MOUNT_DEV2 $MOUNT_POINT2 $MOUNT_FS2
```

- 7i** Click *Apply* to save the load script changes.

- 8** Click the *Unload Script* link to go to the Unload Script page, then modify the script to handle the CSM container and multiple volumes on it.

For an example, see [Section 13.3.8, “Sample Unload Script for a CSM Resource with a Segment Manager on It,”](#) on page 242.

- 8a** Modify the definitions as described in [Step 7a](#) through [Step 7d](#) in the load script changes.

```
#define the file system types
MOUNT_FS1=ext3
MOUNT_FS2=xfs

#define the container name
CONTAINER_NAME=csm44

#define the volume devices
VOLUME_DEV1=csm44p1
VOLUME_DEV2=csm44p2

#define the devices
MOUNT_DEV1=/dev/mapper/$VOLUME_DEV1
MOUNT_DEV2=/dev/mapper/$VOLUME_DEV2

#define the mount points
MOUNT_POINT1=/mnt/lxvolext3
MOUNT_POINT2=/mnt/lxvolxfs
```

- 8b** Add an unmount command for each volume:

```
#unmount the volumes
exit_on_error umount_fs $MOUNT_DEV1 $MOUNT_POINT1 $MOUNT_FS1
exit_on_error umount_fs $MOUNT_DEV2 $MOUNT_POINT2 $MOUNT_FS2
```

- 8c** Add a `kpartx` command to deactivate the partitions. The command should come before the `csmport` command in the unload script.

```
#deactivate the partitions
exit_on_error /sbin/kpartx -d /dev/mapper/$CONTAINER_NAME
```

- 8d** Add a `csmport` command to deactivate the CSM container:

```
#deactivate the container
exit_on_error csmport -e $CONTAINER_NAME
```

- 8e** Click *Apply* to save the unload script changes.

- 9 Click the *Monitor Script* link to go to the Monitor Script page, then modify the script to handle the CSM container and multiple volumes on it.

For an example, see [Section 13.3.9, “Sample Monitor Script for a CSM Resource with a Segment Manager on It,” on page 243.](#)

- 9a Modify the definitions as described in [Step 7a](#) through [Step 7d](#) in the load script changes.

```
#define the file system types
MOUNT_FS1=ext3
MOUNT_FS2=xfs

#define the container name
CONTAINER_NAME=csm44

#define the volume devices
VOLUME_DEV1=csm44p1
VOLUME_DEV2=csm44p2

#define the devices
MOUNT_DEV1=/dev/mapper/$VOLUME_DEV1
MOUNT_DEV2=/dev/mapper/$VOLUME_DEV2

#define the mount points
MOUNT_POINT1=/mnt/lxvolext3
MOUNT_POINT2=/mnt/lxvolxfs
```

- 9b Add a check for each volume:

```
#check the volumes
exit_on_error status_fs $MOUNT_DEV1 $MOUNT_POINT1 $MOUNT_FS1
exit_on_error status_fs $MOUNT_DEV2 $MOUNT_POINT2 $MOUNT_FS2
```

- 9c Add a `csmport` command to check the status of the CSM container:

```
#check the container
exit_on_error csmport -c $CONTAINER_NAME
```

- 9d Click *Apply* to save the monitor script changes.

- 10 Continue with [Section 13.4, “Configuring the Preferred Nodes and Cluster Settings for a CSM Cluster Resource,” on page 243.](#)

Do not bring the CSM cluster resources online again until after its *Preferred Nodes* list has been modified to use only OES 11 nodes.

13.3.7 Sample Load Script for a CSM Resource with a Segment Manager on It

Use the following sample load script to complete the fields for your CSM cluster resource on OES 11:

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

#define the IP address
RESOURCE_IP=10.10.10.44

#define the file system types
MOUNT_FS1=ext3
MOUNT_FS2=xfs

#define the container name
CONTAINER_NAME=csm44

#define the volume devices
VOLUME_DEV1=csm44p1
VOLUME_DEV2=csm44p2
```

```

#define the devices
MOUNT_DEV1=/dev/mapper/$VOLUME_DEV1
MOUNT_DEV2=/dev/mapper/$VOLUME_DEV2

#define the mount points
MOUNT_POINT1=/mnt/lxvolext3
MOUNT_POINT2=/mnt/lxvolxfs

#if the mount path does not exist, create it
ignore_error mkdir -p $MOUNT_POINT1
ignore_error mkdir -p $MOUNT_POINT2

#activate the container
exit_on_error csmport -i $CONTAINER_NAME

#activate the partitions
exit_on_error /sbin/kpartx -a /dev/mapper/$CONTAINER_NAME

#mount the file systems
exit_on_error mount_fs $MOUNT_DEV1 $MOUNT_POINT1 $MOUNT_FS1
exit_on_error mount_fs $MOUNT_DEV2 $MOUNT_POINT2 $MOUNT_FS2

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

exit 0

```

You can modify the `mount_fs` command to add mount options that are available for the file system. Check the function `mount_fs` in the `/opt/novell/ncs/lib/ncsfuns` 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

```

If the path to the mount point does not exist on other nodes, you can add a line to create the mount point path in the script before the line that mounts the file system:

```

# create the mount point path when loading on a new node
ignore_error mkdir -p $MOUNT_POINT

```

13.3.8 Sample Unload Script for a CSM Resource with a Segment Manager on It

Use the following sample unload script to complete the fields for your CSM cluster resource on OES 11:

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

# define the IP address
RESOURCE_IP=10.10.10.44

# define the file system types
MOUNT_FS1=ext3
MOUNT_FS2=xfs

#define the container name
CONTAINER_NAME=csm44

#define the volume devices
VOLUME_DEV1=csm44p1
VOLUME_DEV2=csm44p2

#define the devices
MOUNT_DEV1=/dev/mapper/$VOLUME_DEV1
MOUNT_DEV2=/dev/mapper/$VOLUME_DEV2

#define the mount points
MOUNT_POINT1=/mnt/lxvoext3
MOUNT_POINT2=/mnt/lxvolxfs

#del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

#unmount the volumes
exit_on_error umount_fs $MOUNT_DEV1 $MOUNT_POINT1 $MOUNT_FS1
exit_on_error umount_fs $MOUNT_DEV2 $MOUNT_POINT2 $MOUNT_FS2

#deactivate the partitions
exit_on_error /sbin/kpartx -d /dev/mapper/$CONTAINER_NAME

#deactivate the container
exit_on_error csmport -e $CONTAINER_NAME

# return status
exit 0
```

13.3.9 Sample Monitor Script for a CSM Resource with a Segment Manager on It

Use the following sample monitor script to complete the fields for your CSM cluster resource on OES 11:

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

# define the IP address
RESOURCE_IP=10.10.10.44

# define the file system types
MOUNT_FS1=ext3
MOUNT_FS2=xfs

#define the container name
CONTAINER_NAME=csm44

#define the volume devices
VOLUME_DEV1=csm44p1
VOLUME_DEV2=csm44p2

#define the devices
MOUNT_DEV1=/dev/mapper/$VOLUME_DEV1
MOUNT_DEV2=/dev/mapper/$VOLUME_DEV2

#define the mount points
MOUNT_POINT1=/mnt/lxvolext3
MOUNT_POINT2=/mnt/lxvolxfs

#check the IP address
exit_on_error status_secondary_ipaddress $RESOURCE_IP

#check the volumes
exit_on_error status_fs $MOUNT_DEV1 $MOUNT_POINT1 $MOUNT_FS1
exit_on_error status_fs $MOUNT_DEV2 $MOUNT_POINT2 $MOUNT_FS2

#check the container
exit_on_error csmport -c $CONTAINER_NAME

# return status
exit 0
```

13.4 Configuring the Preferred Nodes and Cluster Settings for a CSM Cluster Resource

After you have modified the CSM cluster resource scripts, the CSM resource is ready to run only on OES 11 nodes. You must modify the preferred nodes for the cluster resource so that it fails over only to other OES 11 nodes. You should also disable the *Resource Follows Master* option until the conversion is finalized. The most preferred node is the first server in the *Assigned Nodes* list for the resource.

NOTE: Resources fail back only 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 verify the policy settings and modify the preferred nodes:

- 1 Ensure that you have offlined the CSM resources from OES 2 nodes and converted the load, unload, and monitor scripts as described in the following:
 - ♦ [Section 13.2, “Modifying the Scripts for a CSM Resource,” on page 230](#)
 - ♦ [Section 13.3, “Modifying the Scripts for a CSM Resource with a Segment Manager On It,” on page 234](#)
- 2 Add one or more OES 11 nodes to the cluster, but do not bring the CSM resources online at this time.
- 3 In iManager, select *Clusters > Cluster Manager*, then select the cluster.
- 4 Click the cluster resource name to open its properties pages.
- 5 In a mixed-mode cluster, deselect the *Resource Follows Master* check box to prevent the resource from inadvertently trying to fail over to a non-OES 11 node.

After the conversion to OES 11 is finalized (no longer a mixed-mode cluster), you can enable this option if you want the resource to always run on the master node. If the master node in the cluster fails, the resource fails over to whichever node becomes the master.
- 6 (Optional) Select the *Ignore Quorum* check box if you don't want the cluster-wide timeout period and node number limit enforced.

The quorum default values were set when you installed 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 Click *Next*.
- 9 Assign nodes to the resource as described in [Section 9.8, “Configuring Preferred Nodes for a Resource,” on page 140](#), then click *Finished*.

IMPORTANT: Ensure that the preferred nodes list contains only OES 11 nodes.

- 10 Configure resource priorities for load order as described in [Section 9.9, “Configuring Resource Priorities for Load Order,” on page 142](#).
- 11 (Optional) Assign the resource to a Resource Mutual Exclusion Group as described in [Section 9.10, “Configuring Resource Mutual Exclusion Groups,” on page 143](#).
- 12 Continue with [Section 13.5, “Verifying the Load and Unload Scripts,” on page 245](#).

13.5 Verifying the Load and Unload Scripts

After you have modified the resource, ensure that the load and unload scripts are working.

- 1 In iManager, select *Clusters > Cluster Manager*, then select the cluster.
- 2 Select the check box next to the new CSM cluster resource, then click *Online*.

Ensure that the resource successfully enters the *Running* state. If the resource goes comatose instead, it is probably because you made an error when typing values in the script definitions. Take the resource offline, go to the resource's *Properties > Scripts* page to review and modify its scripts as needed to correct errors, then try again to online the resource.

- 3 (Optional) Continue with [Section 13.6, "Creating a Virtual Server Object for a CSM Cluster Resource,"](#) on page 245.

13.6 Creating a Virtual Server Object for a CSM Cluster Resource

You can use an existing Virtual Server Object that you created before you moved the CSM resource to OES 11. After a conversion is finalized, you can use the procedure in this section if you need to create a new Virtual Server object for the resource.

When you cluster-enable the Linux POSIX volume, a virtual server object is not created automatically as it is when you cluster-enable an NSS pool. You can assign a virtual server name to the cluster resource by using the `/opt/novell/ncs/bin/ncs_ncpserv.py` script to create an `NCS:NCP Server` object for it. Having a virtual server object allows the Linux POSIX volume cluster resource to be viewed in the *Browse* panel in iManager. You can also bind the virtual server name to the IP address to allow users to access the resource via an assigned name in addition to its IP address

IMPORTANT: The `NCS:NCP Server` object does not give users NCP access to the data on a Linux POSIX volume. An NCP volume is required to do that. To create NCP volumes on the shared Linux POSIX volume, see ["Creating a Shared NCP Volume on the Linux POSIX Cluster Resource"](#) in the *OES 11: NCP Server for Linux Administration Guide*.

The virtual server name is stored in Novell eDirectory in an `NCS:NCP Server` object under the Cluster object where you created the resource. You must add a line to the load and unload scripts that identifies the name of this virtual server and a line that binds or unbinds the name to the IP address of the Linux POSIX cluster resource.

- ♦ [Section 13.6.1, "Creating the Virtual Server Object,"](#) on page 246
- ♦ [Section 13.6.2, "Modifying the Load Script,"](#) on page 246
- ♦ [Section 13.6.3, "Modifying the Unload Script,"](#) on page 247
- ♦ [Section 13.6.4, "Activating the Script Changes,"](#) on page 247
- ♦ [Section 13.6.5, "Verifying the Virtual Server,"](#) on page 247

13.6.1 Creating the Virtual Server Object

You use the `/opt/novell/ncs/bin/ncs_ncpserv.py` script to create a virtual server object (NCS:NCP Server) in eDirectory for the Linux POSIX volume cluster resource. If the resource does not have NCP volumes on it, you do not use the `-v` option. For information about the `ncs_ncpserv.py` script, see [Section A.6, “Virtual NCS NCP Server Object \(ncs_ncpserv.py\) Script,” on page 286](#).

- 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 lx_volumename -i resource_ip_address
```

Replace the `lx_volumename` and `resource_ip_address` with the information for your particular solution.

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, to create the NCS:NCP Server object for the `lxvol44` cluster resource where the IP address is `10.10.10.44` and the cluster context is `ou=clusters,ou=city,o=mycompany`, enter

```
./ncs_ncpserv.py -c lxvol44 -i 10.10.10.44
```

The confirmation message is displayed:

```
NCP Server 'cn=cluster1_lxvol44_server,ou=clusters,ou=city,o=mycompany'
created.
```

- 4 Continue with [Section 13.6.2, “Modifying the Load Script,” on page 246](#)

13.6.2 Modifying the Load Script

After you have created an NCS:NCP Server object, you must modify the load script so that it binds the NCS:NCP Server object to the resource.

- 1 In iManager, select *Clusters > Cluster Options*, then select the cluster.
- 2 Click the name link of the Linux POSIX cluster resource to open its Cluster Resource Properties page.
- 3 Click the *Scripts* tab to open the *Load* script.
- 4 In the definition area, add the following lines to define the virtual NCP server name:

```
# define NCP server name
NCP_SERVER=cluster1_lxvol44_server
```

Replace the NCP server name with the name for your virtual NCP server.

- 5 Under the mount line, add a line to bind the NCP server name to the resource IP address:

```
# bind the NCP volume
exit_on_error ncpcon bind --ncpservname=$NCP_SERVER --ipaddress=$RESOURCE_IP
```

- 6 Click *Apply* to save your changes.

The script changes are not active until the next time the cluster resource is taken offline, and then brought online. Do not active the script changes at this time.

- 7 Continue with [Section 13.6.3, “Modifying the Unload Script,” on page 247](#).

13.6.3 Modifying the Unload Script

After you have created an NCS:NCP Server object, you must modify the unload script so that it unbinds the NCS:NCP Server object from the resource.

- 1 In iManager, select *Clusters > Cluster Options*, then select the cluster.
- 2 Click the name link of the Linux POSIX cluster resource to open its Cluster Resource Properties page.
- 3 Click the *Scripts* tab, then click *Unload* to open the *Unload* script.
- 4 In the definition area, add the following lines to define the virtual NCP server name:

```
# define NCP server name
NCP_SERVER=cluster1_lxvol44_server
```

Replace the NCP server name with the name for your virtual NCP server. Use the same value for variable that you did in the load script.

- 5 Under the definition, add a line to unbind the NCP server name from the resource IP address:

```
# unbind the NCP server name
ignore_error ncpcon unbind --ncpservname=$NCP_SERVER --
ipaddress=$RESOURCE_IP
```

- 6 Click *Apply* to save your changes.
The script changes are not active until the next time the cluster resource is taken offline, and then brought online.
- 7 Continue with [Section 13.6.4, “Activating the Script Changes,” on page 247](#).

13.6.4 Activating the Script Changes

The script changes are not active until the next time the cluster resource is taken offline, and then brought online.

- 1 In iManager, select *Clusters > Cluster Manager*, then select the cluster.
- 2 Select the check box next to the Linux POSIX cluster resource, then click *Offline*.
Wait until the resource is reports an Offline status before continuing.
- 3 Select the check box next to the Linux POSIX cluster resource, then click *Online*.
Wait until the resource is reports an Online status before continuing.
- 4 Continue with [Section 13.6.5, “Verifying the Virtual Server,” on page 247](#).

13.6.5 Verifying the Virtual Server

Verify that an NCS:NCP Server object appears in the *Browse* panel in iManager.

- 1 In the iManager toolbar, click the *View Objects* icon.
- 2 In the left panel, click *Browse*.
- 3 Browse to the Cluster container to see the virtual server object for the cluster resource, such as `cluster1_lxvol44_server`.

13.7 Deleting a Cluster Resource

To avoid data corruption, ensure that you offline the cluster resource before attempting to delete it, especially if the resource uses shared storage.

WARNING: Deleting a cluster resource while it is online can result in deletion errors, and the data associated with the shared storage is not recoverable.

14 Configuring Novell Cluster Services in a Xen Virtualization Environment

Novell Cluster Services is installed and configured in a virtualization environment by using the same methods and processes as those used on a physical Novell Open Enterprise Server (OES) 11 server. It can be used in the host and guest virtual environments.

You can install Novell Cluster Services on an OES 11 Xen virtualization host server (Dom0). You can create cluster resources that contain the virtual machine information by using the Xen and XenLive resource templates. You can fail over or cluster migrate these virtual machine cluster resources to different physical nodes in your cluster. Only the Xen and XenLive templates can be used in the OES 11 Xen host server environment.

For a virtual machine, you can use any virtualization hypervisor (such as Xen, KVM, VMware, and Hyper-V) that supports the 64-bit SUSE Linux Enterprise Server 11 SP1 or later operating system, which is the platform used for OES 11 services. All templates except Xen and XenLive are valid on guest servers (DomU). When Novell Cluster Services is installed on an OES 11 guest server environment, no additional changes or special configuration are required. You can create clusters consisting of all virtual nodes, or use a combination of virtual and physical nodes.

Although many different cluster virtualization scenarios are possible, only those outlined in the sections below have been tested:

- ♦ [Section 14.1, “Prerequisites for Xen Host Server Environments,” on page 249](#)
- ♦ [Section 14.2, “Virtual Machines as Cluster Resources,” on page 250](#)
- ♦ [Section 14.3, “Virtual Machines as Cluster Nodes,” on page 258](#)
- ♦ [Section 14.4, “Virtual Cluster Nodes in Separate Clusters,” on page 259](#)
- ♦ [Section 14.5, “Mixed Physical and Virtual Node Clusters,” on page 259](#)
- ♦ [Section 14.6, “Additional Information,” on page 261](#)

14.1 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>) on Novell Support (<http://www.novell.com/support>).

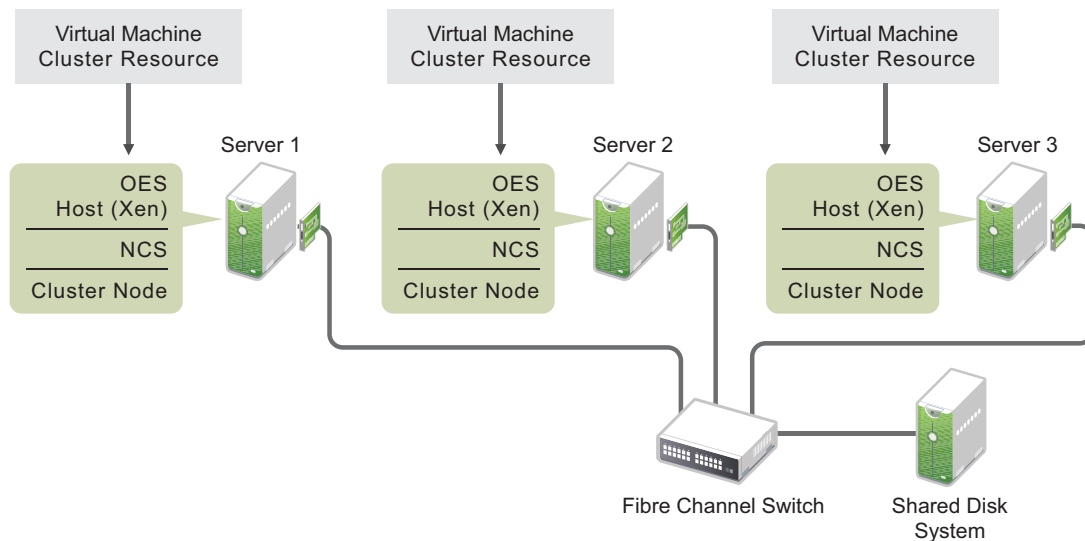
14.2 Virtual Machines as Cluster Resources

In this scenario, you have an OES 11 cluster configured on physical machines. OES 11 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 11 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 14-1 depicts how this setup might look. Novell Cluster Services (NCS) is installed and running on the virtual machine (VM) host server.

Figure 14-1 Virtual Machines as Cluster Resources



The following sections describe how to create a cluster resource and its cluster scripts for each virtual machine:

- [Section 14.2.1, “Creating a Xen Virtual Machine Cluster Resource,” on page 251](#)
- [Section 14.2.2, “Configuring Virtual Machine Load, Unload, and Monitor Scripts,” on page 252](#)
- [Section 14.2.3, “Setting Up Live Migration,” on page 257](#)

14.2.1 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.

Ensure 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.suse.com/documentation/sles11/book_xen/data/book_xen.html) (http://www.suse.com/documentation/sles11/book_xen/data/book_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 user name 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 252.

14.2.2 Configuring Virtual Machine Load, Unload, and Monitor Scripts

The Xen resource templates configure the virtual machine resource by automatically creating load, unload, and monitor scripts, setting failover and failback modes, and assigning the virtual machine as a resource to all nodes in the cluster.

The load, unload, and monitor scripts for virtual machine cluster resources do not need to be modified if all the following are true:

- ♦ The resource name is the same as the virtual machine name.
- ♦ The configuration file name is the same as the virtual machine name.
- ♦ The mount point directory name is the same as the virtual machine name.
- ♦ You are using the Reiser file system.

If you are not modifying the scripts, continue the setup by configuring the resource policies and the resource server assignments. For information, see [Section 9.7, “Configuring the Start, Failover, and Failback Modes for Cluster Resources,” on page 139](#) and [Section 9.8, “Configuring Preferred Nodes for a Resource,” on page 140](#).

If you are modifying the scripts, continue with the following sections:

- ♦ [“Configuring the Load Script” on page 252](#)
- ♦ [“Configuring the Unload Script” on page 254](#)
- ♦ [“Configuring the Monitor Script” on page 255](#)

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 252](#)
- ♦ [“Sample XenLive_Template Load Script” on page 253](#)

- 2 Click *Next* and continue with [“Configuring the Unload Script” on page 254](#).

Sample Xen_Template Load Script

The Xen_Template load script appears similar to the following example:

```

#!/bin/bash
. /opt/novell/ncs/lib/ncsfuns
export OCF_ROOT=/usr/lib/ocf/

#define domU name
OCF_RESOURCE_INSTANCE=xen_vm_name

#define shared volume name
VOLUME_NAME=${OCF_RESOURCE_INSTANCE}

#define the volume group name
CONTAINER_NAME=name
# define the device
MOUNT_DEV=/dev/${CONTAINER_NAME}/${VOLUME_NAME}

# filesystem settings
export OCF_RESKEY_device=$MOUNT_DEV
export OCF_RESKEY_directory=/mnt/${OCF_RESOURCE_INSTANCE}
export OCF_RESKEY_fstype=ext3
#export OCF_RESKEY_options=

# service settings
export OCF_RESKEY_xmfile=${OCF_RESKEY_directory}/${OCF_RESOURCE_INSTANCE}

#activate the volume group
exit_on_error vgchange -a ey $VOLGROUP_NAME

# mount the file system
exit_on_error ocf_start Filesystem

# start the service
exit_on_error ocf_start Xen

# return status
exit 0

```

Sample XenLive_Template Load Script

The XenLive_Template load script appears similar to the following example:

```

#!/bin/bash
. /opt/novell/ncs/lib/ncsfuns
export OCF_ROOT=/usr/lib/ocf/
OCF_RESOURCE_INSTANCE=xen_vm_name

# filesystem settings
export OCF_RESKEY_directory=/mnt/${OCF_RESOURCE_INSTANCE}

# service settings
export OCF_RESKEY_xmfile=${OCF_RESKEY_directory}/${OCF_RESOURCE_INSTANCE}

# start the service
if [ -n "$NCS_TOFROM" ]
then
    exit_on_error ocf_migrate_from Xen
else
    exit_on_error ocf_start Xen
fi

# return status
exit 0

```

Configuring the Unload Script

The virtual machine resource unload script page should now be displayed. The unload script contains commands to stop the virtual machine. You can customize some commands for your specific configuration.

- 1 View and, if necessary, edit the lines in the script for your specific directory structure, mount point, configuration files, and file system type (in the Xen_Template).

Use the same values that you specified in the load script.

See the following examples of the default Xen_Template and XenLive_Template unload scripts:

- ♦ [“Sample Xen_Template Unload Script” on page 254](#)
- ♦ [“Sample XenLive_Template Unload Script” on page 255](#)

- 2 Click *Next*, then continue the setup by configuring the resource policies and the resource server assignments.

For information, see [Section 9.7, “Configuring the Start, Failover, and Failback Modes for Cluster Resources,” on page 139](#) and [Section 9.8, “Configuring Preferred Nodes for a Resource,” on page 140](#).

- 3 If you want to enable monitoring for the resource, continue with [“Configuring the Monitor Script” on page 255](#).

Sample Xen_Template Unload Script

The Xen_Template unload script appears similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfncs
export OCF_ROOT=/usr/lib/ocf/

#define domU name
OCF_RESOURCE_INSTANCE=xen_vm_name

#define shared volume name
VOLUME_NAME=$OCF_RESOURCE_INSTANCE

#define the volume group name
CONTAINER_NAME=name
# define the device
MOUNT_DEV=/dev/$CONTAINER_NAME/$VOLUME_NAME

# filesystem settings
export OCF_RESKEY_device=$MOUNT_DEV
export OCF_RESKEY_directory=/mnt/$OCF_RESOURCE_INSTANCE
export OCF_RESKEY_fstype=ext3
#export OCF_RESKEY_options=

# service settings
export OCF_RESKEY_xmfile=$OCF_RESKEY_directory/$OCF_RESOURCE_INSTANCE

# stop the service
exit_on_error ocf_stop Xen

# umount the file system
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error ocf_stop Filesystem

#deactivate the volume group
exit_on_error vgchange -a n $VOLGROUP_NAME

# return status
exit 0
```

Sample XenLive_Template Unload Script

The XenLive_Template unload script appears similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuns
export OCF_ROOT=/usr/lib/ocf/
OCF_RESOURCE_INSTANCE=xen_vm_name

# filesystem settings
export OCF_RESKEY_directory=/mnt/${OCF_RESOURCE_INSTANCE}

# service settings
export OCF_RESKEY_xmfile=${OCF_RESKEY_directory}/${OCF_RESOURCE_INSTANCE}
export OCF_RESKEY_CRM_meta_migrate_target=${NCS_TOFROM}

RC=0
# stop the service
if [ -n "$NCS_TOFROM" ]
then
    RC=`ocf_migrate_to Xen`
    if [ $RC -ne 0 ]
    then
        ignore_error ocf_stop Xen
    fi
else
    ignore_error ocf_stop Xen
fi

# return status
exit $RC
```

Configuring the Monitor Script

The Xen_Template and XenLive_Template each include a resource 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 255](#)
- ♦ [“Viewing or Modifying the Monitor Script” on page 256](#)
- ♦ [“Sample Xen_Template Monitor Script” on page 256](#)
- ♦ [“Sample XenLive_Template Monitor Script” on page 257](#)

Enabling Resource Monitoring

To enable resource monitoring for a virtual machine cluster resource:

- 1 In iManager, select *Clusters > 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, to migrate to another server, or to reboot the hosting node (without synchronizing or unmounting the disks) if a failure action initiates. The reboot option is normally used only for a mission-critical cluster resource that must remain available.

If the failure action initiates and you chose the option to migrate the resource to another server, the resource migrates to the next server in its *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.

If the failure action initiates and you chose the option to reboot the hosting node without synchronizing or unmounting the disks, each of the resources on the hosting node will fail over to the next server in its *Assigned Nodes* list because of the reboot. This is a hard reboot, not a graceful one.

With resource monitoring, the *Failover*, *Failback*, and *Start* modes have no effect on where the resource migrates. This means that a resource that has been migrated by the resource monitoring failure action does not migrate back to the node it migrated from unless you manually migrate it back.

Viewing or Modifying the Monitor Script

To view or customize the monitor script for the virtual machine's cluster resource:

- 1 In iManager, select *Clusters > 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 256](#)
- ♦ ["Sample XenLive_Template Monitor Script" on page 257](#)

- 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 255](#). Cluster Services marks the monitor process as failed right after the defined timeout expires, but it must wait for the process to conclude before it can start other resource operations.

Sample Xen_Template Monitor Script

The *Xen_Template* monitor script appears similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
export OCF_ROOT=/usr/lib/ocf/

#define domU name
OCF_RESOURCE_INSTANCE=xen_vm_name

#define shared volume name
VOLUME_NAME=${OCF_RESOURCE_INSTANCE}

#define the volume group name
CONTAINER_NAME=name
# define the device
MOUNT_DEV=/dev/${CONTAINER_NAME}/${VOLUME_NAME}

# filesystem settings
export OCF_RESKEY_device=${MOUNT_DEV}
export OCF_RESKEY_directory=/mnt/${OCF_RESOURCE_INSTANCE}
export OCF_RESKEY_fstype=ext3
#export OCF_RESKEY_options=

# service settings
export OCF_RESKEY_xmfile=${OCF_RESKEY_directory}/${OCF_RESOURCE_INSTANCE}

#check the logical volume
exit_on_error status_lv $MOUNT_DEV

# status of the file system
exit_on_error ocf_status Filesystem

# status of the service
exit_on_error ocf_status Xen

# return status
exit 0
```

Sample XenLive_Template Monitor Script

The XenLive_Template monitor script appears similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
export OCF_ROOT=/usr/lib/ocf/
OCF_RESOURCE_INSTANCE=xen_vm_name

# filesystem settings
export OCF_RESKEY_directory=/mnt/${OCF_RESOURCE_INSTANCE}

# service settings
export OCF_RESKEY_xmfile=${OCF_RESKEY_directory}/${OCF_RESOURCE_INSTANCE}

# status of the service
exit_on_error ocf_status Xen

# return status
exit 0
```

14.2.3 Setting Up Live Migration

Live migrations use the XenLive template. You can manually copy the virtual machine configuration files to the same path on each node of the cluster, or you can set up an OCFS2 file system for the configuration files. Do one of the following:

- ♦ Manually copy the configuration file for the virtual machine to the same directory (the path must be the same) on each cluster node where the virtual machine will run.

- Configure the OCFS2 file system on a shared disk system and copy the virtual machine configuration file to a directory on the file system. You also must ensure that all cluster nodes where the virtual machine will run have access to the OCFS2 file system on the shared disk system.

Ensure that your OCFS2 file system is working properly before you attempt to use it with Novell Cluster Services.

An overview of OCFS2 is available in “Oracle Cluster File System 2” (http://www.novell.com/documentation/sle_ha/book_sleha/data/cha_ha_ocfs2.html) in the *SUSE Linux Enterprise Server 11 SP1 High Availability Administration Guide* (http://www.novell.com/documentation/sle_ha/book_sleha/data/book_sleha.html). 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>).

14.3 Virtual Machines as Cluster Nodes

In this scenario, you have OES 11 (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.

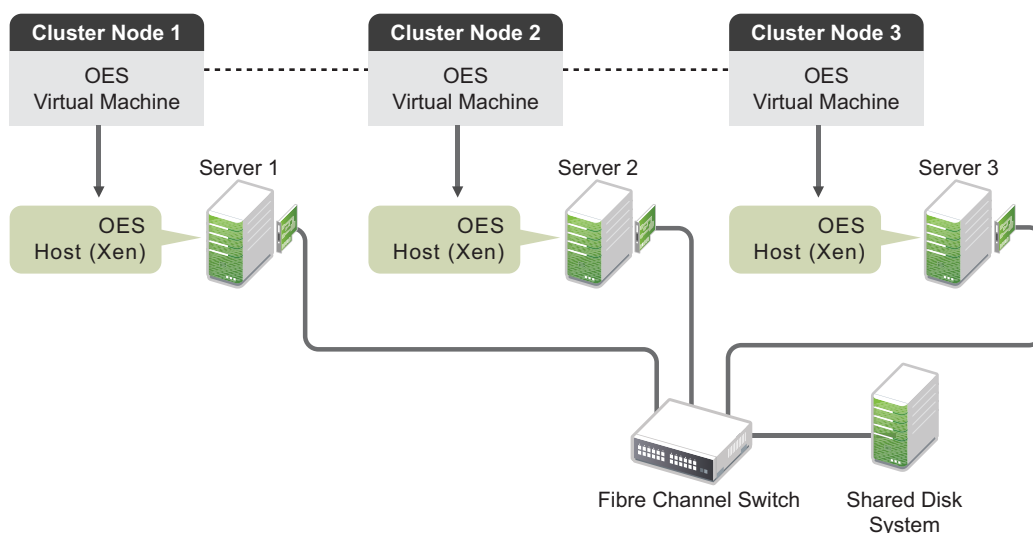
IMPORTANT: All virtual cluster nodes in the same cluster should be either Linux or NetWare. Do not mix Linux and NetWare cluster nodes in the same cluster.

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 14-2 depicts using virtual machines as cluster nodes.

Figure 14-2 Virtual Machines as Cluster Nodes



14.4 Virtual Cluster Nodes in Separate Clusters

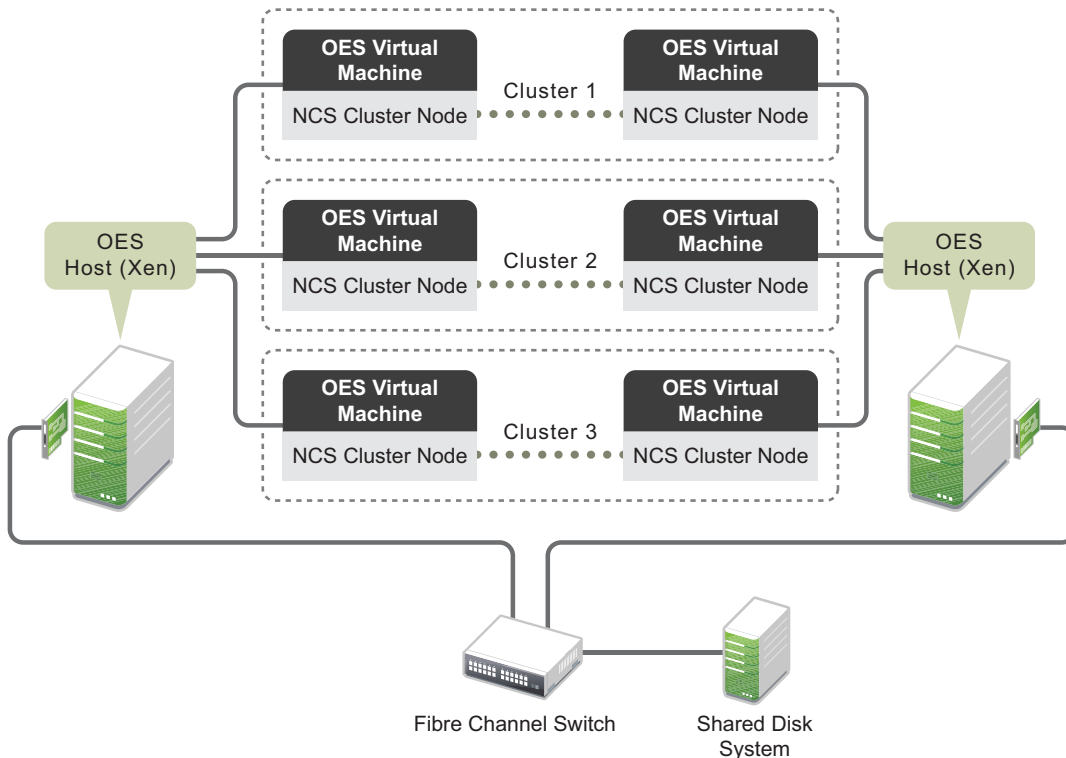
In this scenario, you have OES 11 (Xen) installed and configured on each node (physical machine). You then create multiple NetWare or Linux virtual machines on each physical machine and install and configure Novell Cluster Services on each virtual machine. During the Novell Cluster Services installation, 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 14-3 depicts using virtual cluster nodes in separate clusters.

Figure 14-3 Virtual Cluster Nodes in Separate Clusters



14.5 Mixed Physical and Virtual Node Clusters

This is a temporary scenario that is used for upgrading cluster nodes or converting clusters from physical to virtual cluster nodes.

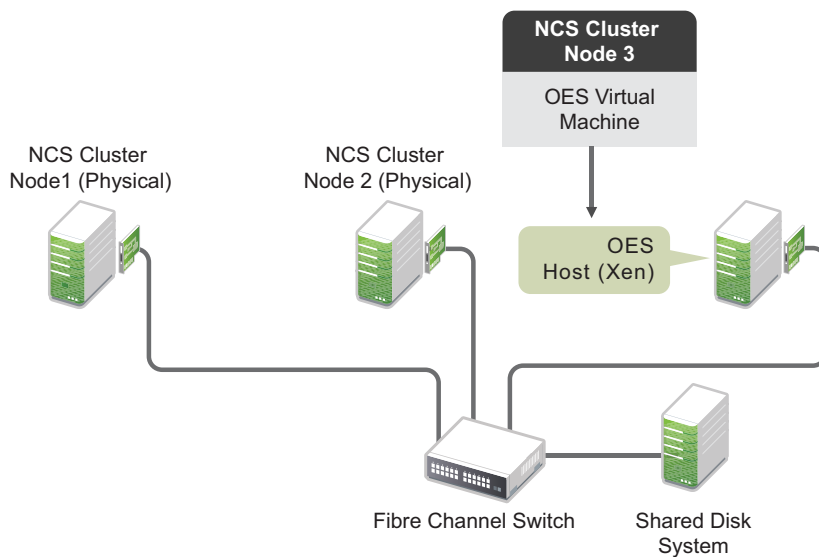
This can be done through several different methods. One method is to add a virtual NetWare cluster node to an existing physical NetWare cluster. To do this, you install an OES 11 (Xen) server (physical machine). You then create a NetWare virtual machine on the physical machine and install and

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 14-4 depicts how this setup might look.

Figure 14-4 *Mixed Physical and Virtual Node Cluster*

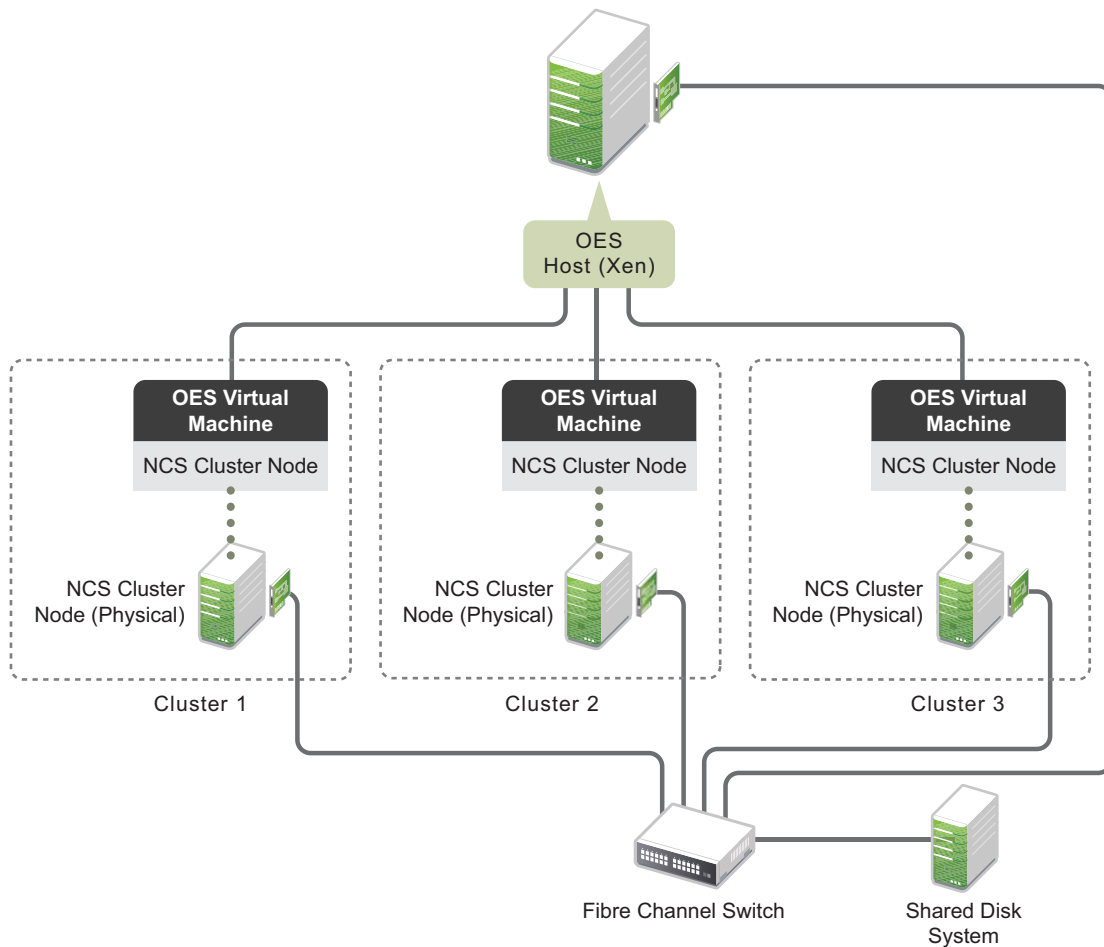


Another method is to install Novell Cluster Services on physical NetWare nodes and create a separate cluster for each node. You then install an OES 11 (Xen) server (physical machine) and create NetWare virtual machines and install Novell Cluster Services on each virtual machine. You can then add one 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 14-5 depicts how this setup might look.

Figure 14-5 *Separate Mixed Physical and Virtual Node Clusters*



14.6 Additional Information

To get started with virtualization, see “[Introduction to Xen Virtualization](http://www.suse.com/documentation/sles11/book_xen/data/cha_xen_basics.html)” (http://www.suse.com/documentation/sles11/book_xen/data/cha_xen_basics.html) in the *Virtualization with Xen* (http://www.suse.com/documentation/sles11/book_xen/data/book_xen.html) guide.

For information on setting up virtualized NetWare, see “[Installing, Upgrading, or Updating OES on a VM](#)” in the *OES 11: Installation Guide* guide.

For information on setting up virtualized OES 11, see “[Installing, Upgrading, or Updating OES on a VM](#)” in the *OES 11: Installation Guide* guide.

15 Troubleshooting Novell Cluster Services

This section describes known issues and frequently asked questions for managing Novell Cluster Services.

- ♦ [Section 15.1, “File Location Information,” on page 264](#)
- ♦ [Section 15.2, “Version Issues,” on page 264](#)
- ♦ [Section 15.3, “Diagnosing a Poison Pill Condition,” on page 264](#)
- ♦ [Section 15.4, “Diagnosing Cluster Problems,” on page 264](#)
- ♦ [Section 15.5, “Why didn’t my resource go to the node I specified?,” on page 265](#)
- ♦ [Section 15.6, “Cluster Search Times Out \(Bad XML Error\),” on page 266](#)
- ♦ [Section 15.7, “iManager Error 500 Occurs while Connecting to a Cluster,” on page 266](#)
- ♦ [Section 15.8, “A Device Name Is Required to Create a Cluster Partition,” on page 266](#)
- ♦ [Section 15.9, “Cluster Resource Goes Comatose Immediately After Migration or Failover,” on page 266](#)
- ♦ [Section 15.10, “smdr.novell Is Not Registered with SLP for a New Cluster Resource,” on page 267](#)
- ♦ [Section 15.11, “NSS Takes Up to 10 Minutes to Load When the Server Is Rebooted \(Linux\),” on page 267](#)
- ♦ [Section 15.12, “Could Not Delete This Resource Data_Server \(Error 499\),” on page 267](#)
- ♦ [Section 15.13, “Problem Authenticating to Remote Servers during Cluster Configuration,” on page 267](#)
- ♦ [Section 15.14, “Problem Connecting to an iSCSI Target,” on page 268](#)
- ♦ [Section 15.15, “Problem Deleting a Cluster Resource or Clustered Pool,” on page 268](#)
- ♦ [Section 15.16, “Can’t Find the Prevent Cascading Failover Option,” on page 268](#)
- ♦ [Section 15.17, “Is there a way to uninstall Novell Cluster Services from a server?,” on page 268](#)
- ♦ [Section 15.18, “Running supportconfig Concurrently on Multiple Nodes Causes Cluster Problems,” on page 268](#)

15.1 File Location Information

Information about Novell Cluster Services scripts and configuration can be found in the following locations:

Information	File Location
Load scripts	/var/opt/novell/ncs
Unload scripts	
Cluster Services configuration files	
Load and unload script output files (.out)	/var/opt/novell/log/ncs
Most resource load and unload issues can be solved with the information in the .out files.	
Cluster Services scripts, such as ldncs and unldncs	/opt/novell/ncs/bin
Resource templates	/opt/novell/ncs/templates

15.2 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 289](#).

15.3 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

15.4 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?

15.5 Why didn't my resource go to the node I specified?

In a failover situation, the cluster looks for the best solution for the node. It goes to the most preferred node if it can. If that node is not available, it attempts to go to the next node in the resource's preferred nodes list. It considers if the node is available and if it has an Resource Mutual Exclusion rules that must be obeyed for resources that are already mounted on the target node. If all of the nodes in the resource's preferred nodes list are not available or have RME conflicts, the resource goes comatose instead of failing over to one of the nodes in its preferred nodes list.

In a cluster migration situation, the resource is currently online and you specify a target node. The resource cannot go to the specified node if any of the following conditions exist, and it stays online on the original node. You can view the target node's log to understand why the resource was not available.

- ♦ The specified node is not in the resource's preferred nodes list.
- ♦ The specified node is in the resource's preferred nodes list, but the node is currently running another resource that is a Resource Mutual Exclusion conflict for the resource you are attempting to migrate.
- ♦ The specified node is in the resource's preferred nodes list, but the node is not available.

15.6 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.

15.7 iManager Error 500 Occurs while Connecting to a Cluster

Novell iManager might hang or time-out while trying to connect to a server with the Cluster plug-in or the Storage plug-in. Error 500 is reported. This error can occur if your login user name for iManager is not a user that is enabled with Linux User Management (LUM) on the target server.

To avoid this problem, ensure that the user name is LUM enabled on the target server. The Cluster plug-in connects to the master node in the cluster (that is, the node hosting the Master IP Resource).

For more information, see *Technical Information Document: iManager error 500 while trying to managing cluster or storage (TID 7010295)* in the [Novell Knowledgebase \(http://www.novell.com/support/\)](http://www.novell.com/support/).

15.8 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.

15.9 Cluster Resource Goes Comatose Immediately After Migration or Failover

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

For information, see "[SLP](#)" on page 45.

15.10 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
```

15.11 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”](#) in the *OES 11: NSS File System Administration Guide for Linux*.

15.12 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
```

15.13 Problem Authenticating to Remote Servers during Cluster Configuration

During the 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.

15.14 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 reboot the Linux initiator server. This is required for each Linux initiator server that will access the iSCSI target.

For instructions on configuring an OES 11 server as an iSCSI initiator and connecting to an iSCSI target, go to “Mass Storage over IP Networks--iSCSI” (http://www.suse.com/documentation/sles11/stor_admin/data/cha_inst_system_iscsi.html) in the *SUSE Linux Enterprise Server 11 Storage Administration Guide*. (http://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html)

15.15 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.

15.16 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 `ldnccs.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.

15.17 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 [Section 8.14, “Deleting a Cluster Node from a Cluster,”](#) on page 113.

15.18 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.

16 Security Considerations

This section describes security issues and recommendations for Novell Cluster Services for Novell Open Enterprise Server (OES) 11. 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 16.1, “Cluster Administration Rights,” on page 269](#)
- ♦ [Section 16.2, “Ports,” on page 269](#)
- ♦ [Section 16.3, “Email Alerts,” on page 270](#)
- ♦ [Section 16.4, “Log Files,” on page 270](#)

16.1 Cluster Administration Rights

For information about the rights needed to install and manage Novell Cluster Services, see the following sections:

- ♦ [“Novell eDirectory 8.8.6” on page 42](#)
- ♦ [“Novell Domain Services for Windows” on page 53](#)
- ♦ [“SFCB and CIMOM” on page 48](#)
- ♦ [“SLP” on page 45](#)
- ♦ [“eDirectory Tree” on page 42](#)
- ♦ [“Cluster Installation Administrator” on page 37](#)
- ♦ [“Cluster Administrator or Administrator-Equivalent User” on page 39](#)

16.2 Ports

For each cluster, you can specify the port used for cluster communication. The default cluster port number is 7023, and is automatically assigned when the cluster is created. You might need to modify this value if there is a port conflict. For information, see [Section 7.6, “Modifying the Cluster IP Address and Port Properties,” on page 97](#).

You must specify the eDirectory port when you install clusters. Port 636 is the default port number for eDirectory communications in the tree.

If you are using a firewall, the port must be opened for CIMOM communications between the cluster and iManager. Port 5989 is the default setting for secure HTTP (HTTPS) communications.

16.3 Email Alerts

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

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

For information, see [Section 7.4, “Configuring Cluster Event Email Notification,”](#) on page 95.

16.4 Log Files

NCS normally writes log messages to `/var/log/messages`. You can view events in the *NCS Event Log* in the Clusters plug-in to Novell iManager.

NCS installation writes log messages to `/var/opt/novell/install/ncslog`.

NCS writes resource runtime messages to files under `/var/opt/novell/log/ncs/`. Output log files are named with the resource name, load or unload to indicate which script, and the `.out` extension. For example: `/var/opt/novell/log/ncs/myresource.unload.out` reports the log for the myresource unload script messages.

A Console Commands for Novell Cluster Services

Novell Cluster Services provides several console commands to help you perform certain cluster-related tasks.

- ♦ [Section A.1, “Cluster Management Commands,” on page 271](#)
- ♦ [Section A.2, “Business Continuity Clustering Commands,” on page 276](#)
- ♦ [Section A.3, “extend_schema Command,” on page 276](#)
- ♦ [Section A.4, “SBD Utility,” on page 277](#)
- ♦ [Section A.5, “Cluster Segment Manager Import/Export \(csmport\) Utility,” on page 280](#)
- ♦ [Section A.6, “Virtual NCS NCP Server Object \(ncs_ncpserv.py\) Script,” on page 286](#)
- ♦ [Section A.7, “Configuration Update \(ncs-configd.py\) Script,” on page 288](#)

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>	<p>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.</p> <p>Specify the resource name in the command and use the YES or NO switch to accept or reject an alert on a resource to control fail over, fail back, or manual start.</p> <p>Example</p> <pre>cluster alert res1 yes</pre>

Cluster Console Command	Description
CONVERT preview <i>[resource]</i> CONVERT preview CONVERT commit	<p>Finalizes the cluster conversion from NetWare to Linux after all nodes in a mixed cluster have been converted to Linux. The <code>CLUSTER CONVERT</code> command can be executed only on Linux cluster nodes.</p> <p>Specify a resource name with the Preview option to view the resource load and unload script changes prior to finalizing the conversion. If you do not provide a resource name, it displays the information for all resources.</p> <p>Use the Commit switch without specifying a resource to finalize the conversion for all cluster resources.</p> <p>Examples</p> <pre>cluster convert preview res1 cluster convert preview cluster convert commit</pre>
DOWN	<p>Removes all cluster nodes from the cluster. This command has the same effect as executing the <code>CLUSTER LEAVE</code> command on every server in the cluster.</p> <p>You are prompted to confirm the cluster down.</p>
EXEC "path_to_script"	<p>Executes the specified script on all nodes in the cluster.</p> <p>Example</p> <pre>cluster exec "command"</pre>
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 email 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> <p>Example</p> <pre>cluster info protocol</pre>
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 and running 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>

Cluster Console Command	Description
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> <p>Running the command without the {ON OFF} parameter reports the maintenance status of the cluster (that is, whether maintenance mode is on or off).</p> <p>Example</p> <pre>cluster maintenance on</pre>
MIGRATE {resource} {node}	<p>Migrates the specified resource from the node where it is currently running to the node that you specify in the command. The node that you specify must be running in the cluster and also be in the resource's <i>Assigned Nodes</i> list.</p> <p>Example</p> <pre>cluster migrate res1 node1</pre>
MONITOR {resource} <start stop status>	<p>Manually starts, stops, or checks the status of resource monitoring. The resource must be running on the node where you issue the command. Resource monitoring must be enabled for the resource in order to use this command.</p> <p>Example</p> <pre>cluster monitor sh_pool_res_01 status</pre>
OFFLINE {resource}	<p>Unloads the specified resource from the node where it is currently running.</p> <p>Example</p> <pre>cluster offline res1</pre>
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. The node that you specify must be running in the cluster and also be in the resource's <i>Assigned Nodes</i> list.</p> <p>Example</p> <pre>cluster online res1 cluster online res1 node1</pre>
POOLS	<p>Lists the NSS pools on the shared disk system that are accessible by Novell Cluster Services.</p>
RENAME <old_resource_name> <new_resource_name>	<p>Renames a pool cluster resource. This command must be issued from the master node. The resource must be in offline state to be renamed. The new name must not exist prior to the renaming.</p> <p>Renaming the resource does not modify the resource's virtual server name (NCS:NCP Server object). You can delete and re-create the NCS:NCS Server object with a new name by using iManager.</p> <p>Example</p> <pre>cluster rename POOL1_SERVER customized_name22</pre>

Cluster Console Command	Description
RESOURCES	Lists all resources that currently exist in the cluster. The resources do not need to be online or running.
RESTART [seconds]	<p>Restarts Novell Cluster Services software on all servers in the cluster. The <code>cluster leave</code> process begins immediately. Specify a value of 60 or more seconds as the time to wait before the <code>cluster join</code> begins.</p> <p>The default setting is 60 seconds. The default setting is used if a value is not provided, if the value provided is smaller than 60 seconds, or if the input is invalid.</p> <p>You are prompted to confirm the cluster restart.</p> <p>Example</p> <pre>cluster restart 60</pre>
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.suse.com/documentation/sles11/stor_admin/data/scandev.html) in the <i>SLES 11 SP1: Storage Administration Guide</i> (http://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html).</p>

Cluster Console Command	Description
SET {Parameter} {Value}	<p>Sets cluster parameters individually for the cluster. See Chapter 7, "Configuring Cluster Policies and Priorities," on page 91 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 master 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 IP port number of the cluster master.</p> <p>QUORUMWAIT sets the amount of time in seconds that the cluster waits before resources start to load.</p> <p>QUORUM sets the number of nodes that must be running in the cluster before resources will start to load.</p> <p>HEARTBEAT sets the amount of time in seconds between transmits for all nodes in the cluster except the master.</p> <p>TOLERANCE sets the amount of time in seconds that the master node gives all other nodes in the cluster to signal that they are alive.</p> <p>MASTERWATCHDOG sets the amount of time in seconds between transmits for the master node in the cluster.</p> <p>SLAVEWATCHDOG sets the amount of time in seconds that the slave nodes give the master node in the cluster to signal that it is alive.</p> <p>MAXRETRANSMITS sets the maximum number of times transmits will be attempted between the master node and slave nodes.</p> <p>ENABLEEMAIL enables and disables email notification. You can set the value to OFF to disable email notification, or either CRITICAL or VERBOSE to enable email notification.</p> <p>EMAILADDRESSES lets you specify the email addresses used for email notification. The addresses should be separated by spaces. Using this parameter without specifying any addresses clears existing addresses that have been set previously.</p> <p>EMAILOPTIONS sets the email notification options. Specify XML as the value to receive email notification in XML format. Not specifying any value with this parameter turns notification in XML format off.</p> <p>RESOURCEPRIORITY sets the resource priority for the cluster.</p> <p>Example</p> <pre>cluster set ipaddress 10.1.1.1</pre>
STATS {Display, Clear}	<p>The Display parameter reports the node number, node name, and heartbeat information to the console screen.</p> <p>The Clear parameter resets the various stats counters. This command is useful for troubleshooting network problems related to the cluster heartbeat protocol.</p> <p>Example</p> <pre>cluster stats display</pre>
STATUS {resource}	<p>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.</p>
VIEW	<p>Displays the node name, cluster epoch number, master node name, and a list of nodes that are currently members of the cluster.</p>

A.2 Business Continuity Clustering Commands

If Novell Business Continuity Clustering (BCC) is installed and running on the cluster, you can use the following additional `cluster` commands to manage the BCC cluster environment.

```
connections
credentials
disable
enable
nsmi
resetresources
```

For information about using these commands, see “Console Commands for BCC” (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

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.

- ♦ [Section A.3.1, “Syntax,” on page 276](#)
- ♦ [Section A.3.2, “Options,” on page 277](#)
- ♦ [Section A.3.3, “Example,” on page 277](#)

A.3.1 Syntax

```
/opt/novell/oes-install/util/extend_schema --port port_num admin_username  
admin_password server_ip_address schema_file
```

To extend the schema, the tree administrator user modifies the following schema files in the given order:

```
/opt/novell/ncs/schema/ncs.ldif  
  
/opt/novell/ncs/schema/ncpsrvr.prelldif
```

A.3.2 Options

Replace the parameters with the credentials to access and location of the eDirectory schema files.

Parameter	Description	Example
port_num	The port number you assigned for eDirectory communications in the tree where you plan to install clusters. The default port is 636.	636
admin_username	The typeful fully distinguished user name of the administrator who has the eDirectory rights needed to extend the schema.	cn=admin,o=example
admin_password	The password of the administrator user.	pas5W0rd
server_ip_address	The IP address of the eDirectory server that contains the schema files.	10.10.10.1

A.3.3 Example

For example, enter the following commands in the order shown, using the values for your particular solution:

```
/opt/novell/oes-install/util/extend_schema --port 636 cn=admin,o=example pas5W0rd  
10.1.1.1 /opt/novell/ncs/schema/ncs.ldif
```

```
/opt/novell/oes-install/util/extend_schema --port 636 cn=admin,o=example pas5W0rd  
10.1.1.1 /opt/novell/ncs/schema/ncpsrvr.preldef
```

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 the Novell Storage Services (NSS) Management Utility (`nssmu(8)`) or the Storage plug-in to iManager. The NSS utility called `ncsinit(8)` is available for initializing a device and setting it to a shared state.

- ♦ [Section A.4.1, “Syntax,” on page 278](#)
- ♦ [Section A.4.2, “Options,” on page 278](#)
- ♦ [Section A.4.3, “Return Value,” on page 279](#)
- ♦ [Section A.4.4, “Examples,” on page 279](#)

A.4.1 Syntax

```
sbdutil [-c|-f|-i|-v] [-s] [-d device] [-d device] [-p partition] [-n cluster_name]
sbdutil -c -d device [-s [size]] [-n cluster_name]
sbdutil -f [-s] [-n cluster_name]
sbdutil -i -p partition [-s] [-n cluster_name]
sbdutil -v [-p partition] [-s] [-n cluster_name]
```

Log in to a node in the cluster where you want to create the SBD partition, then enter the command at a terminal console as the `root` user or any other user in `admin` or `ncsgroup`. If the command succeeds, the partition name is printed. See [Section A.4.3, “Return Value,” on page 279](#) for more information.

A.4.2 Options

-c

Create an SBD partition. This option requires at least one device to be specified. You can create a mirrored SBD by supplying more than one device with multiple instances of the `-d` option.

IMPORTANT: Do not create an SBD partition for a cluster that already has an SBD partition. If you need to re-create the SBD for a cluster, delete its existing SBD first.

To delete an SBD:

1. Enter `cluster down` at the server console of one cluster server.
This causes all cluster servers to leave the cluster.
2. Delete the SBD partition.
You can use `nssmu` or other utilities to delete the SBD partition.

-f

Find the SBD partition.

-i

Initialize the SBD partition. If there is information in the `sbdutil -v` view from old nodes that have been removed from the cluster, you can run `sbdutil -i` to initialize the view. This can be done without any issues while the cluster is running.

-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. Specify only the base (leaf) names (such as `sdb` or `mpathd`) with the `-d` option.

-p partition

Use this partition instead of searching for one.

-n cluster_name

Use the specified cluster name instead of getting the name from `cluster.xml`. If this option is not specified, the SBD partition is named by default with the cluster name that is found in the `cluster.xml` file on the node.

-s

Assume the device is a shared disk system instead of checking `cluster.xml`. An optional partition size (in MB) can also be specified when creating a partition (-c). The default size is 8 MB. Some of the allocated space is used for storing metadata such as the partition table and MBR (master boot record) for the partition.

Specify the size as -1 to use all free space on the device. This option allows Novell Cluster Services to use a whole disk/LUN (or LUNs) that you set aside for SBD.

A.4.3 Return Value

If the command succeeds, it prints the partition name to the screen and returns "0". Otherwise, no partition name is printed out, and a non-zero error code is returned.

You can use "echo \$?" at the end of the command to check the return code:

```
/opt/novell/ncs/bin/sbdutil -c -n <cluster_name> -d <device> -d <device> -s <size>;  
echo $?
```

Examples

The following example demonstrates a successful command where the `mycluster1.sbd` partition is created:

```
/opt/novell/ncs/bin/sbdutil -c -n mycluster1 -d MD_EMC_001G_A -d MD_EMC_001G_B -s -  
1; echo $?  
/dev/nss/mycluster1.sbd  
0
```

The following example demonstrates a failed command where no partition is created, and a non-zero error code is returned:

```
/opt/novell/ncs/bin/sbdutil -c -n mycluster1 -d MD_EMC_001G_A -d MD_EMC_001G_B -s  
2048; echo $?  
2
```

A.4.4 Examples

This program normally runs as root on a node in the cluster that will be using the SBD.

sbdutil -f

This tells you whether an SBD partition exists and identifies the device on the SAN where the SBD partition is located.

sbdutil -c -n mycluster1 -d sdb -s -1

Creates an SBD partition for the cluster named `mycluster1` on the LUN `/dev/sdb`. The size is set to use all of the free space on the device.

sbdutil -c -n mycluster1 -d sdb -d sdc

Creates a mirrored SBD partition for the cluster named `mycluster1` on the LUN `/dev/sdb` with the mirror on `/dev/sdc`.

A.5 Cluster Segment Manager Import/Export (csmport) Utility

The Novell Cluster Services Cluster Segment Manager (CSM) Import/Export utility allows you to use a Linux POSIX volume that was cluster-enabled with a CSM container. This conversion is required because the Enterprise Volume Segment Manager (EVMS) was deprecated on SUSE Linux Enterprise Server 11, and the CSM is no longer available on OES 11 servers.

The utility can be used to scan for CSM containers and check if they can be imported for use in the cluster. This is the default mode when you run it without any options, or with only the `-v` option. By default, the utility searches block devices (`/dev/sd*`) for containers. Information about each device is displayed. Details about the CSM container are present if the container exists and can be imported for use in the cluster.

The utility also allows you to do the following:

- ♦ Search for a specified CSM container and check if it can be imported for use in the cluster.
- ♦ Import a specified CSM container for exclusive use in the cluster, and generate a corresponding Device Mapper object for it as `/dev/mapper/csm_name`.
- ♦ Export a specified CSM container by removing its Device Mapper object from the `/dev/mapper` directory. This allows the container to be imported by other nodes.

You can specify devices to search for any of the options. Wildcards such as the question mark (?) and asterisk (*) can be used when specifying devices. A question mark replaces a character in that position. An asterisk replaces any number of characters in that position.

- ♦ [Section A.5.1, “Syntax,” on page 280](#)
- ♦ [Section A.5.2, “Options,” on page 280](#)
- ♦ [Section A.5.3, “Examples,” on page 282](#)

A.5.1 Syntax

Open a console on a cluster node, then issue the command as the root user.

```
/opt/novell/ncs/bin/csmport [-v|-s] [-h] [-e devmapper_object | -i container_name |  
-c container_name] [device...]
```

A.5.2 Options

No options

Use the command with no options to search the block devices (`/dev/sd*`) for containers.

```
/opt/novell/ncs/bin/csmport
```

[devices]

Use the command with no options and specify the devices to search in order to specify which devices to search for CSM containers. By default, the utility searches block devices (`/dev/sd*`) for containers.

```
/opt/novell/ncs/bin/csmport [device...]
```

When specifying devices, wildcards (? and *) can be used.

Examples

The following command searches the multipath device `/dev/mapper/MD_EMC_005G` for CSM containers.

```
/opt/novell/ncs/bin/csmport -v /dev/mapper/MD_EMC_005G
```

The following command uses three question mark wildcards to search multipath devices with any three characters in those positions in the device name:

```
/opt/novell/ncs/bin/csmport -v /dev/mapper/MD_EMC_???G
```

The following command uses an asterisk wildcard to search multipath devices with any number of characters in that position in the device name:

```
/opt/novell/ncs/bin/csmport -v /dev/mapper/MD_EMC_*
```

-h

Prints the help page for this command.

-c <container_name> [device...]

Searches the `/dev/sd*` devices or specified devices for the named CSM container.

```
/opt/novell/ncs/bin/csmport -c container_name
```

-i <container_name> [device...]

Use this option to import a CSM container on a node for exclusive use in the cluster. The utility achieves this by searching for the container and if found, extracting the volume inside the container and exposing it exclusively at `/dev/mapper/container_name`.

The utility uses the container name instead of the Linux POSIX volume name because some CSM containers might contain compatibility volumes, which do not have names. The `/dev/mapper/container_name` object can be mounted at the old mount point that was used for the shared Linux POSIX volume in that container. The mount point can have the same name or different name than the Device Mapper object.

```
/opt/novell/ncs/bin/csmport -i container_name
```

Example

Suppose that you created a Linux POSIX cluster resource with the following settings on an OES 2 server:

```
RESOURCE_IP:          10.10.10.44
CONTAINER_NAME:       csm44
EVMS_VOLUME_NAME:     lxvol44
MOUNT_DEV:            /dev/evms/$CONTAINER_NAME/$EVMS_VOLUME_NAME
MOUNT_POINT:          /mnt/users
```

On an OES 11 server, you can use the following command to generate the Device Mapper object at `/dev/mapper/csm44`:

```
/opt/novell/ncs/bin/csmport -i csm44
```

The Device Mapper object can be mounted at the old mount point.

The resulting OES 11 settings are:

```
RESOURCE_IP:          10.10.10.44
CONTAINER_NAME:       csm44
MOUNT_DEV:            /dev/mapper/$CONTAINER_NAME
MOUNT_POINT:          /mnt/users
```

-e <devmapper_object> [device...]

Use this option to export the Device Mapper object from the node. It removes the object from the `/dev/mapper` directory. Basically, the option undoes what the `-i` option did, so that the container can be imported by other nodes.

```
/opt/novell/ncs/bin/csmport -e devmapper_object
```

Example

The following command exports the Device Mapper object `/dev/mapper/csm44`.

```
/opt/novell/ncs/bin/csmport -e csm44
```

-v

Verbose mode can be used with any of the options.

-s

Silent mode can be used with any of the options.

A.5.3 Examples

This section provides examples of typical tasks you perform when using the `csmport` command.

- ♦ [“Scanning for CSM Containers” on page 282](#)
- ♦ [“Importing a CSM Container for Exclusive Use on a Node in a Cluster” on page 285](#)
- ♦ [“Exporting a CSM Container” on page 286](#)
- ♦ [“Searching for a Container on a Specified Device” on page 286](#)
- ♦ [“Searching for Containers on Multiple Devices with Wildcards” on page 286](#)
- ♦ [“Importing a CSM Container on a Specified Device” on page 286](#)

Scanning for CSM Containers

Scan for CSM containers and check if they can be imported for use in the cluster.

```
/opt/novell/ncs/bin/csmport -v
```

In the following sample output, only device `/dev/sdi` has a CSM container `CSM5GN0EVMSV01` and can be imported to `/dev/mapper/CSM5GN0EVMSV01` for its exclusive use on the cluster node.

Scan for CSM containers

```
Search device '/dev/sdn' for CSM containers.
Retrieve device information.
Device '/dev/sdn': major=8, minor=208.
Refresh device '/dev/sdn'.
Open device '/dev/sdn' for reading.
Check CSM metadata on device '/dev/sdn'.
/dev/sdn is not a CSM container.
```

```
Search device '/dev/sdm' for CSM containers.
Retrieve device information.
Device '/dev/sdm': major=8, minor=192.
Refresh device '/dev/sdm'.
Open device '/dev/sdm' for reading.
Check CSM metadata on device '/dev/sdm'.
/dev/sdm is not a CSM container.
```

```

Search device '/dev/sdh' for CSM containers.
Retrieve device information.
Device '/dev/sdh': major=8, minor=112.
Refresh device '/dev/sdh'.
Open device '/dev/sdh' for reading.
Check CSM metadata on device '/dev/sdh'.
/dev/sdh has a CSM container.
Cluster ID doesn't match. On disk = 37037513, in memory = 135657986.
Container name = 'AJTestContainer01' on disk.
Container is accessible (1).

Search device '/dev/sdc' for CSM containers.
Retrieve device information.
Device '/dev/sdc': major=8, minor=32.
Refresh device '/dev/sdc'.
Open device '/dev/sdc' for reading.
Check CSM metadata on device '/dev/sdc'.
/dev/sdc is not a CSM container.

Search device '/dev/sdl' for CSM containers.
Retrieve device information.
Device '/dev/sdl': major=8, minor=176.
Refresh device '/dev/sdl'.
Open device '/dev/sdl' for reading.
Check CSM metadata on device '/dev/sdl'.
/dev/sdl is not a CSM container.

Search device '/dev/sdi' for CSM containers.
Retrieve device information.
Device '/dev/sdi': major=8, minor=128.
Refresh device '/dev/sdi'.
Open device '/dev/sdi' for reading.
Check CSM metadata on device '/dev/sdi'.
/dev/sdi has a CSM container.
Container name = 'CSM5GNoEVMSVol' on disk.
Container is accessible (4).
No EVMS volume.
New name: /dev/mapper/CSM5GNoEVMSVol
Size: 10485757 sectors.
Device '/dev/sdi' has a CSM container 'CSM5GNoEVMSVol' and can be converted to
/dev/mapper/CSM5GNoEVMSVol.

Search device '/dev/sdk' for CSM containers.
Retrieve device information.
Device '/dev/sdk': major=8, minor=160.
Refresh device '/dev/sdk'.
Open device '/dev/sdk' for reading.
Check CSM metadata on device '/dev/sdk'.
/dev/sdk is not a CSM container.

Search device '/dev/sdg' for CSM containers.
Retrieve device information.
Device '/dev/sdg': major=8, minor=96.
Refresh device '/dev/sdg'.
Open device '/dev/sdg' for reading.
Check CSM metadata on device '/dev/sdg'.
/dev/sdg is not a CSM container.

Search device '/dev/sdb' for CSM containers.
Retrieve device information.
Device '/dev/sdb': major=8, minor=16.
Refresh device '/dev/sdb'.
Open device '/dev/sdb' for reading.
Check CSM metadata on device '/dev/sdb'.
/dev/sdb is not a CSM container.

Search device '/dev/sdt' for CSM containers.
Retrieve device information.
Device '/dev/sdt': major=65, minor=48.
Refresh device '/dev/sdt'.

```

```

Open device '/dev/sdt' for reading.
Failed to read the 1st sector of device '/dev/sdt': Input/output error (5).

Search device '/dev/sdx' for CSM containers.
Retrieve device information.
Device '/dev/sdx': major=65, minor=112.
Refresh device '/dev/sdx'.
Open device '/dev/sdx' for reading.
Failed to read the 1st sector of device '/dev/sdx': Input/output error (5).

Search device '/dev/sdj' for CSM containers.
Retrieve device information.
Device '/dev/sdj': major=8, minor=144.
Refresh device '/dev/sdj'.
Open device '/dev/sdj' for reading.
Check CSM metadata on device '/dev/sdj'.
/dev/sdj has a CSM container.
Cluster ID doesn't match. On disk = 37037513, in memory = 135657986.
Container name = 'CMS4GCon' on disk.
Container is accessible (4).

Search device '/dev/sdy' for CSM containers.
Retrieve device information.
Device '/dev/sdy': major=65, minor=128.
Refresh device '/dev/sdy'.
Open device '/dev/sdy' for reading.
Failed to read the 1st sector of device '/dev/sdy': Input/output error (5).

Search device '/dev/sdu' for CSM containers.
Retrieve device information.
Device '/dev/sdu': major=65, minor=64.
Refresh device '/dev/sdu'.
Open device '/dev/sdu' for reading.
Failed to read the 1st sector of device '/dev/sdu': Input/output error (5).

Search device '/dev/sdd' for CSM containers.
Retrieve device information.
Device '/dev/sdd': major=8, minor=48.
Refresh device '/dev/sdd'.
Open device '/dev/sdd' for reading.
Failed to read the 1st sector of device '/dev/sdd': Input/output error (5).

Search device '/dev/sdq' for CSM containers.
Retrieve device information.
Device '/dev/sdq': major=65, minor=0.
Refresh device '/dev/sdq'.
Open device '/dev/sdq' for reading.
Check CSM metadata on device '/dev/sdq'.
/dev/sdq is not a CSM container.

Search device '/dev/sds' for CSM containers.
Retrieve device information.
Device '/dev/sds': major=65, minor=32.
Refresh device '/dev/sds'.
Open device '/dev/sds' for reading.
Failed to read the 1st sector of device '/dev/sds': Input/output error (5).

Search device '/dev/sdv' for CSM containers.
Retrieve device information.
Device '/dev/sdv': major=65, minor=80.
Refresh device '/dev/sdv'.
Open device '/dev/sdv' for reading.
Failed to read the 1st sector of device '/dev/sdv': Input/output error (5).

Search device '/dev/sdaa' for CSM containers.
Retrieve device information.
Device '/dev/sdaa': major=65, minor=160.
Refresh device '/dev/sdaa'.
Open device '/dev/sdaa' for reading.
Failed to read the 1st sector of device '/dev/sdaa': Input/output error (5).

```

```

Search device '/dev/sde' for CSM containers.
  Retrieve device information.
  Device '/dev/sde': major=8, minor=64.
  Refresh device '/dev/sde'.
  Open device '/dev/sde' for reading.
  Check CSM metadata on device '/dev/sde'.
  /dev/sde is not a CSM container.

Search device '/dev/sdz' for CSM containers.
  Retrieve device information.
  Device '/dev/sdz': major=65, minor=144.
  Refresh device '/dev/sdz'.
  Open device '/dev/sdz' for reading.
  Failed to read the 1st sector of device '/dev/sdz': Input/output error (5).

Search device '/dev/sdw' for CSM containers.
  Retrieve device information.
  Device '/dev/sdw': major=65, minor=96.
  Refresh device '/dev/sdw'.
  Open device '/dev/sdw' for reading.
  Failed to read the 1st sector of device '/dev/sdw': Input/output error (5).

Search device '/dev/sdr' for CSM containers.
  Retrieve device information.
  Device '/dev/sdr': major=65, minor=16.
  Refresh device '/dev/sdr'.
  Open device '/dev/sdr' for reading.
  Failed to read the 1st sector of device '/dev/sdr': Input/output error (5).

Search device '/dev/sdp' for CSM containers.
  Retrieve device information.
  Device '/dev/sdp': major=8, minor=240.
  Refresh device '/dev/sdp'.
  Open device '/dev/sdp' for reading.
  Failed to read the 1st sector of device '/dev/sdp': Input/output error (5).

Search device '/dev/sdo' for CSM containers.
  Retrieve device information.
  Device '/dev/sdo': major=8, minor=224.
  Refresh device '/dev/sdo'.
  Open device '/dev/sdo' for reading.
  Failed to read the 1st sector of device '/dev/sdo': Input/output error (5).

Search device '/dev/sdf' for CSM containers.
  Retrieve device information.
  Device '/dev/sdf': major=8, minor=80.
  Refresh device '/dev/sdf'.
  Open device '/dev/sdf' for reading.
  Check CSM metadata on device '/dev/sdf'.
  /dev/sdf is not a CSM container.

Search device '/dev/sda' for CSM containers.
  Retrieve device information.
  Device '/dev/sda': major=8, minor=0.
  Refresh device '/dev/sda'.
  Open device '/dev/sda' for reading.
  Check CSM metadata on device '/dev/sda'.
  /dev/sda is not a CSM container.

```

Importing a CSM Container for Exclusive Use on a Node in a Cluster

Use the `-i` option to import the CSM container named `CSM5GNoEVMSV01` for exclusive use on a node in a cluster.

```
/opt/novell/ncs/bin/csmport -i CSM5GNoEVMSV01
```

This command generates the Device Mapper object `/dev/mapper/CSM5GNoEVMSVol`. You can map this object to the desired mount point.

Exporting a CSM Container

Use the `-e` option to export the Device Mapper object `/dev/mapper/CSM5GNoEVMSVol`.

```
/opt/novell/ncs/bin/csmport -e CSM5GNoEVMSVol
```

This removes the object and allows the container to be imported by other nodes.

Searching for a Container on a Specified Device

Specify a device to be searched in order to speed up the search for a container or containers. For example, to search for CSM containers on a multipath device `/dev/mapper/MD_EMC_005G`, enter

```
/opt/novell/ncs/bin/csmport /dev/mapper/MD_EMC_005G
```

Searching for Containers on Multiple Devices with Wildcards

Specify multiple devices to be searched and use wildcards to include multiple devices of that type.

```
/opt/novell/ncs/bin/csmport -v /dev/mapper/MD_EMC_???G /dev/disk/by-id/dm-uuid-mpath-*
```

Importing a CSM Container on a Specified Device

When importing the `CSM5GNoEVMSVol` container, specify the multipath device `/dev/mapper/MD_EMC_005G` in order to speed up the search for the container.

```
/opt/novell/ncs/bin/csmport -i CSM5GNoEVMSVol /dev/mapper/MD_EMC_005G
```

A.6 Virtual NCS NCP Server Object (ncs_ncpserv.py) Script

The `/opt/novell/ncs/bin/ncs_ncpserv.py` script creates a virtual NCP Server object (NCS:NCP Server) in eDirectory, and associates it with none, one, or multiple NCP volumes that you specify. It automatically renames the NCP Volume objects to use the cluster name instead of the server name where the NCP volume was created. NCP clients access files on the Linux POSIX volume via the virtual server name.

- ♦ [Section A.6.1, “Syntax,” on page 286](#)
- ♦ [Section A.6.2, “Examples,” on page 287](#)

A.6.1 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 lx_volumename -i resource_ip_address [-v <ncp_volumename /  
"ncp_volumename1:ncp_volumename2:..."> ]
```

Replace the `lx_volumename`, `resource_ip_address`, and `ncp_volumename` with the information for your particular solution.

If the -v option is not specified, all of the NCP volumes that currently exist on the Linux POSIX cluster resource are bound to the IP address.

If you enter multiple volume names, use colons to delimit the names and put quotation marks around the list of names. Volume names can be listed by the volume name (MY_NNCP_VOL06) or by the volume distinguished name (cn=CLUS_02_MY_NNCP_VOL06,o=novell), or any combination of the two methods.

A.6.2 Examples

In the following examples, the resource IP address is 10.10.10.44, the cluster name is cluster1 and the cluster context is ou=clusters,o=mycompany.

Example 1

To specify a single NCP volume named USERS on the lxvol44 cluster resource, enter

```
./ncs_ncpserv.py -c lxvol44 -i 10.10.10.44 -v USERS
```

The following confirmation message is displayed:

```
NCP Server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' created.

Object 'cn=servername_USERS,ou=clusters,o=mycompany' renamed to
'cn=cluster1_USERS,ou=clusters,o=mycompany'.
The volume name you need to use in the scripts is: USERS
NCP server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' and volume
'cn=cluster1_USERS,ou=clusters,o=mycompany' are linked with each other.
```

Example 2

To specify multiple NCP volumes on the lxvol44 cluster resource, enter

```
./ncs_ncpserv.py -c lxvol44 -i 10.10.10.44 -v
"USERS:MY_NCP_VOL06:cn=servername_MY_NCP_VOL07,ou=clusters,o=novell"
```

The following confirmation message is displayed:

```
NCP Server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' created.

Object 'cn=servername_USERS,ou=clusters,o=mycompany' renamed to
'cn=cluster1_USERS,ou=clusters,o=mycompany'.
The volume name you need to use in the scripts is: USERS
NCP server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' and volume
'cn=cluster1_USERS,ou=clusters,o=mycompany' are linked with each other.

Object 'cn=servername_MY_NCP_VOL06,ou=clusters,o=mycompany' renamed to
'cn=cluster1_MY_NCP_VOL06,ou=clusters,o=mycompany'.
The volume name you need to use in the scripts is: MY_NCP_VOL06
NCP server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' and volume
'cn=cluster1_MY_NCP_VOL06,ou=clusters,o=mycompany' are linked with each other.

Object 'cn=servername_MY_NCP_VOL07,ou=clusters,o=mycompany' renamed to
'cn=cluster1_MY_NCP_VOL07,ou=clusters,o=mycompany'.
The volume name you need to use in the scripts is: MY_NCP_VOL07
NCP server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' and volume
'cn=cluster1_MY_NCP_VOL07,ou=clusters,o=mycompany' are linked with each other.
```

A.7 Configuration Update (ncs-configd.py) Script

You can use the `/opt/novell/ncs/bin/ncs-configd.py` script to push updates to all nodes in a cluster.

Syntax

As the root user, enter

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

Novell Cluster Services can be running or not running when you issue the command.

B Files for Novell Cluster Services

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 File Name and Path	Purpose
<code>/etc/init.d/novell-ncs</code>	LSB Compliant Service
<code>/etc/opt/novell/ncs/clstrlib.conf</code>	Cluster configuration file
<code>/etc/opt/novell/ncs/nodename</code>	Configuration file - the node's name
<code>/opt/novell/ncs/bin/ClusterCli.pl</code>	Cluster CLI engine
<code>/opt/novell/ncs/bin/ClusterCliSnapinInterface.pm</code>	Cluster CLI engine
<code>/opt/novell/ncs/bin/ClusterCliUtils.pm</code>	Cluster CLI engine
<code>/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Alert.pm</code>	Cluster CLI command
<code>/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Convert.pm</code>	Cluster CLI command
<code>/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Down.pm</code>	Cluster CLI command
<code>/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Exec.pm</code>	Cluster CLI command
<code>/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Info.pm</code>	Cluster CLI command
<code>/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Join.pm</code>	Cluster CLI command
<code>/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Leave.pm</code>	Cluster CLI command
<code>/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Maintenance.pm</code>	Cluster CLI command
<code>/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Migrate.pm</code>	Cluster CLI command
<code>/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Monitor.pm</code>	Cluster CLI command

NCS File Name and Path	Purpose
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Offline.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Online.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Pools.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Rename.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Resources.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Restart.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Set.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Stats.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Status.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_View.pm	Cluster CLI command
/opt/novell/ncs/bin/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 email daemon
/opt/novell/ncs/bin/ncs-resourced.py	Daemon used to run load and unload scripts.
/opt/novell/ncs/bin/csmport	CSM Import/Export utility
/opt/novell/ncs/bin/ncstempl.py	Used to install or update cluster resource templates in the Cluster container
/opt/novell/ncs/bin/sbdutil	SBD partition utility
/opt/novell/ncs/bin/uldncs	Unloads NCS; used by the Cluster Stop command.
/opt/novell/ncs/lib/ncs-1.0.0.so	NCS snap-in
/opt/novell/ncs/lib/ncsfuns	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

NCS File Name and Path	Purpose
/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
/usr/share/man/man8/cluster.8.gz	CLUSTER man page
/usr/share/man/man8/csmport.8.gz	CSMPORT man page
/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.
/lib/modules/kernel_dir/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/kernel_dir/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/kernel_dir/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/kernel_dir/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/kernel_dir/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/kernel_dir/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/kernel_dir/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/kernel_dir/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/kernel_dir/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/kernel_dir/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.

C Comparing Novell Cluster Services for Linux and NetWare

[Table C-1](#) compares the features and capabilities of Novell Cluster Services on Novell Open Enterprise Server (OES) 11 to Novell Cluster Services on NetWare 6.5 SP8.

Table C-1 Comparison of Novell Cluster Services for OES 11 and NetWare 6.5 SP8

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
Operating system	OES 11	NetWare 6.5 SP8
Two-node cluster with OES 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	Not yet available for OES 11 BCC 1.2.2 for OES 2 SP3	BCC 1.1 SP2 for NetWare 6.5 SP8
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 5.6, "Configuring Additional Administrators," on page 82 .	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.
OES Common Proxy User	Supported	Not applicable
NCS_Management Group	Yes	Not applicable
Preferred list for LDAP servers	Yes; in the <code>/etc/opt/novell/ncs/clstrlib.conf</code> file	No
Directory-based cluster configuration	Yes; common schema for NetWare and Linux	Yes; 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>The user that installs Novell Cluster Services must have schema extension rights.</p> <p>The schema extension can be performed separately from the Novell Cluster Services install by a user with schema extension rights. For information, see Section 5.2, “Extending the eDirectory Schema to Add Cluster Objects,” on page 60.</p> <p>Afterwards, any administrator with sufficient rights can install Novell Cluster Services. For information, see Section 5.3, “Assigning Install Rights for Container Administrators (or Non-Administrator Users),” on page 62.</p>	The user who installs Novell Cluster Services must have schema extension rights.
Forward migration for Novell Cluster Services	<p>OES 2 to OES 11. Down cluster and rolling cluster upgrade are supported.</p> <p>Special handling is required for Linux POSIX file system cluster resources. For information, see Chapter 6, “Upgrading Clusters from OES 2 to OES 11,” on page 85.</p>	<p>NetWare 6.5 SP8 to OES 11. Down cluster and rolling cluster upgrade are supported.</p> <p>NetWare 6.5 SP6 to NetWare 6.5 SP8. 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 SP8 to OES 11</p> <p>NetWare 6.5 SP6 to OES 11</p> <p>NetWare 6.0 with latest service packs and patches to OES 11. Uses the same process as NetWare 6.5 to OES 11.</p>	Not applicable
mixed-mode 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>

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
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	During the install. After the install and before adding a second node to the cluster by using the sbdutil and specifying two devices. Mirroring an existing SBD by using NLVM.	During the install on the first server in the cluster, or by using the sbdutil after the install.
Shared disks	Fibre Channel SAN LUNs iSCSI SAN LUNs SCSI disks (shared external drive arrays)	Fibre Channel SAN LUNs iSCSI SAN LUNs SCSI disks (shared external drive arrays)
Cluster-aware shared devices	Share devices by using NLVM, NSSMU, or the Storage plug-in to iManager.	Share devices by using NSSMU or the Storage plug-in to iManager.
Requires Novell Storage Services (NSS)	Required to use NLVM and NSSMU to manage devices, and to create NSS pools as cluster resources.	NSS is the default file system on NetWare. Novell Cluster Services is not supported on NetWare traditional volumes.
Volume manager for NSS pools	NetWare partitions on devices managed by Linux Logical Volume Manager	NetWare Segment Manager
Volume manager for Linux POSIX file systems	Linux Logical Volume Manager	NetWare Segment Manager
Requires NCP (NetWare Core Protocol)	NCP Server is required for resources that use NCP, such as NSS pools, NCP volumes, and Dynamic Storage Technology shadow volume pairs.	NCP is the default file access protocol on NetWare.
NSS pools as cluster resources	Yes For information, see Chapter 11, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 155. Shareable for Clustering Multiple-Server Activation Prevention (MSAP) Cluster volume broker; Linux kernel module handles NSS pool events.	Yes For information, see “Setting Up Cluster Resources for Novell Cluster Services” in the NW6.5 SP8: Novell Cluster Services 1.8.5 Administration Guide. Shareable for Clustering Multiple-Server Activation Prevention (MSAP) Cluster volume broker

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
Linux POSIX file systems as cluster resources	<p>Yes</p> <p>For information, see Chapter 12, “Configuring and Managing Cluster Resources for LVM Volume Groups,” on page 195</p> <p>For information about managing resources migrated from OES 2 to OES 11, see Chapter 13, “Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers,” on page 229.</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 11: 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” in the <i>OES 11: 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 14.2, “Virtual Machines as Cluster Resources,” on page 250.</p>	Not applicable
iManager 2.7.4	Yes	Yes
Clusters plug-in and Storage Management plug-in for iManager	Yes	Yes
Cluster-enabling NSS pools by using the Storage plug-in for iManager	Yes	Yes
Cluster-enabling NSS pools by using the NSS Management Utility (NSSMU)	Yes	Yes
Cluster-enabling NSS pools by using NLVM commands	Yes	Not applicable
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.

Feature or Capability	Cluster Services for Linux	Cluster Services for 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 “Planning the Conversion of Load and Unload Scripts” in the <i>OES 11: Novell Cluster Services NetWare to Linux Conversion Guide</i> .	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; with cross-protocol file locking	Yes; with cross-protocol file locking
Novell CIFS support for accessing files on shared NSS pools	Yes; with cross-protocol file locking	Yes; with cross-protocol file locking
Linux Samba/CIFS support for accessing files on shared NSS pools on Linux	Yes; with cross-protocol file locking Requires users to be Linux-enabled with Linux User Management. Requires Universal Password.	Not applicable
Linux Samba/CIFS support for accessing files on shared Linux POSIX file systems	Yes Requires users to be 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>

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
Multipath I/O	<p>Device Mapper - Multipath I/O, or third-party MPIO solutions</p> <p>For information, see “Managing Multipath I/O for Devices” (http://www.suse.com/documentation/sles11/stor_admin/data/multipathing.html) in the <i>SLES 11 SP1: Storage Administration Guide</i>. (http://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html)</p>	<p>Media Manager Multipath I/O, or third-party MPIO solutions</p> <p>For information, see “Managing Multipath I/O to Devices (NetWare)” in the <i>NW 6.5 SP8: NSS File System Administration Guide</i>.</p>
Monitoring script	Yes	No
Resource Mutual Exclusion Groups	Yes	No
Monitoring script	Yes	No

D Comparing Resources Support for Linux and NetWare

[Table D-1](#) compares clustering support for Novell Open Enterprise Server (OES) 11 services when using Novell Cluster Services on OES 11 and on NetWare 6.5 SP8.

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 the [OES 11: Novell Cluster Services NetWare to Linux Conversion Guide](#).

Table D-1 Comparison of Clustering Support for OES 11 Services on Linux and NetWare

Service	NetWare 6.5 SP8	OES 11	For conversion information, see
AFP (Apple Filing Protocol)	Yes See “ Setting Up for Macintosh ” in the NW 6.5 SP8: AFP, CIFS, and NFS (NFAP) Administration Guide .	Yes See “ Configuring AFP with Novell Cluster Services for an NSS File System ” in the OES 11: Novell AFP Administration Guide .	“ Novell AFP ” in the OES 11: Novell Cluster Services NetWare to Linux Conversion Guide
Apache Web Server	Yes See “ Apache with Novell Cluster Services ” in the NW 6.5 SP8: Novell Cluster Services 1.8.5 Resource Configuration Guide .	Yes; use the standard Apache Web Server for Linux.	“ Apache Web Server ” in the OES 11: Novell Cluster Services NetWare to Linux Conversion Guide
Archive and Version Services (Novell)	Yes Uses a MySQL database. See “ Installing and Configuring New Archive Servers ” in the NW 6.5 SP8: Novell Archive and Version Services 2.1 Administration Guide .	Yes Uses a PostgreSQL database, which must reside on a Linux POSIX file system. See “ Configuring Archive and Version Service for Novell Cluster Services ” in the OES 11: Novell Archive and Version Services 2.1 Administration Guide .	“ Novell Archive and Version Services ” in the OES 11: Novell Cluster Services NetWare to Linux Conversion Guide

Service	NetWare 6.5 SP8	OES 11	For conversion information, see
CIFS (Windows File Services)	Yes; Novell CIFS See “ Setting Up for Windows ” in the <i>NW 6.5 SP8: AFP, CIFS, and NFS (NFAP) Administration Guide</i> .	Yes; Novell CIFS See “ Configuring CIFS with Novell Cluster Services for an NSS File System ” in the <i>OES 11: Novell CIFS for Linux Administration Guide</i> .	“ Novell CIFS ” in the <i>OES 11: Novell Cluster Services NetWare to Linux Conversion Guide</i>
DFS (Novell Distributed File Services) Volume location database (VLDB)	Yes See “ Clustering Novell Distributed File Services ” in the <i>NW 6.5 SP8: Novell Distributed File Services Administration Guide</i> .	Yes See “ Clustering Novell Distributed File Services ” in the <i>OES 11 SP1: Novell Distributed File Services Administration Guide for Linux</i> .	“ Novell Distributed File Services VLDB ” in the <i>OES 11: Novell Cluster Services NetWare to Linux Conversion Guide</i>
DHCP	Yes See “ Clustering in NetWare 6.5 ” in the <i>NW 6.5 SP8: Novell DNS/DHCP Services Administration Guide</i> .	Yes DHCP for Linux supports using a shared Linux POSIX file system or a shared NSS pool for the cluster resource. See “ Configuring DHCP with Novell Cluster Services for the Linux File System ” in the <i>OES 11: Novell DNS/DHCP Services for Linux Administration Guide</i> .	“ DHCP Server ” in the <i>OES 11: Novell Cluster Services NetWare to Linux Conversion Guide</i>
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 11: Novell DNS/DHCP Services for Linux Administration Guide</i> .	“ DNS Server ” in the <i>OES 11: Novell Cluster Services NetWare to Linux Conversion Guide</i>
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 11 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.

Service	NetWare 6.5 SP8	OES 11	For conversion information, see
eDirectory Certificate Server	Yes See “Server Certificate Objects and Clustering” in the Novell Certificate Server 3.3.1 Administration Guide (http://www.novell.com/documentation/crt33/).	Yes See “ Server Certificate Objects and Clustering ” in the Novell Certificate Server 3.3.4 Administration Guide .	“eDirectory Server Certificates” in the OES 11: Novell Cluster Services NetWare to Linux Conversion Guide
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 NW6.5 SP8: Novell Cluster Services 1.8.5 Resource Configuration Guide .	Not available on Linux.	This install option was discontinued beginning in NetWare 6.5 SP3. See also MySQL .
FTP Server	Yes See “ Cluster-Enabling NetWare FTP Server ” in the NW 6.5 SP8: Novell FTP Administration Guide .	No; use the Linux Pure-FTPd solution. See “ Cluster Enabling Pure-FTPd in an OES 11 Environment ” in the OES 11: Planning and Implementation Guide .	Not applicable
Identity Manager 3.6.1 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. Requires Identity Manager 3.6.1.	
iPrint	Yes See the NW 6.5 SP8: iPrint Administration Guide .	Yes See “ Configuring iPrint with Novell Cluster Services ” in the OES 11: iPrint Linux Administration Guide .	“Novell iPrint” in the OES 11: Novell Cluster Services NetWare to Linux Conversion Guide
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” in the OES 11: Novell Cluster Services NetWare to Linux Conversion Guide
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.

Service	NetWare 6.5 SP8	OES 11	For conversion information, see
NetStorage	Yes See “ Configuring NetStorage with Novell Cluster Services ” in the <i>NW 6.5 SP8: NetStorage Administration Guide</i> .	Yes See “ Configuring NetStorage with Novell Cluster Services ” in the <i>OES 11: NetStorage Administration Guide for Linux</i> .	No known issues.
NFS	Yes See “ Cluster-Enabling Native File Access for UNIX ” in the <i>NW 6.5 SP8: AFP, CIFS, and NFS (NFAP) Administration Guide</i> .	No; use the standard Linux solution for NFS.	
Novell iFolder 2.1.x	Yes	Not applicable	iFolder 2.1x is replaced by Novell iFolder 3.9 .
Novell iFolder 3.9	Not applicable	Yes See “ Clustering iFolder Servers with Novell Cluster Services for Linux ” in the <i>Novell iFolder 3.9 Administration Guide</i> .	“ Migrating iFolder 2.x ” in the <i>OES 11: Migration Tool 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 11: Novell QuickFinder Server 5.0 Administration Guide</i> .	“ QuickFinder Server ” in the <i>OES 11: Novell Cluster Services NetWare to Linux Conversion Guide</i>
Storage, DST shadow volume pairs	Not applicable	Yes See “ Configuring DST Shadow Volumes with Novell Cluster Services ” in the <i>OES 11: 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.

Service	NetWare 6.5 SP8	OES 11	For conversion information, see
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 11: 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	
Storage, NSS pools and volumes	Yes See “Setting Up Cluster Resources for Novell Cluster Services” in the <i>NW6.5 SP8: Novell Cluster Services 1.8.5 Administration Guide</i> .	Yes See Chapter 11, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 155. For a feature comparison, see “Cross-Platform Issues for NSS” in the <i>OES 11: NSS File System Administration Guide for Linux</i> .	“Novell Storage Services Pools” in the <i>OES 11: Novell Cluster Services NetWare to Linux Conversion Guide</i>
Tomcat	Yes See “Configuring Tomcat and Novell Cluster Services” in the <i>NW6.5 SP8: Novell Cluster Services 1.8.5 Resource Configuration Guide</i> .	Yes; native to Linux Use a similar procedure to the one outlined for Tomcat on NetWare, but use the Linux locations and files.	You cannot convert the NetWare Tomcat configuration for a Linux server.
Xen guest servers as nodes in a cluster	Virtualized NetWare nodes can be used in NetWare clusters. Nodes can be any combination of virtual and physical servers.	Virtualized OES 11 nodes can be used in OES 11 clusters. Nodes can be any combination of virtual and physical servers.	See Chapter 14, “Configuring Novell Cluster Services in a Xen Virtualization Environment,” on page 249.
Xen virtual machines on the host server	Not applicable	Yes; use the Xen and XenLive templates.	See Section 14.2, “Virtual Machines as Cluster Resources,” on page 250.

E Documentation Updates

This section contains information about documentation content changes made to the *Novell Cluster Services Administration Guide* since the initial 2.0 release for Novell Open Enterprise Server 11.

This document was updated on the following dates:

- ♦ [Section E.1, “August 6, 2012,” on page 305](#)
- ♦ [Section E.2, “July 20, 2012,” on page 306](#)
- ♦ [Section E.3, “June 27, 2012,” on page 306](#)
- ♦ [Section E.4, “May 31, 2012,” on page 307](#)
- ♦ [Section E.5, “April 30, 2012,” on page 307](#)
- ♦ [Section E.6, “January 31, 2012,” on page 308](#)

E.1 August 6, 2012

Updates were made to the following section. The changes are explained below.

- ♦ [Section E.1.1, “Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers,” on page 305](#)

E.1.1 Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers

Location	Change
Section 13.3.5, “Configuring the OES 11 NLVM Configuration File,” on page 237	The path for the NLVM configuration file is <code>/etc/opt/novell/nss/nlvm.conf</code> .

E.2 July 20, 2012

Updates were made to the following section. The changes are explained below.

- ♦ [Section E.2.1, “Configuring Cluster Policies and Protocols,” on page 306](#)

E.2.1 Configuring Cluster Policies and Protocols

Location	Change
Table 7-2, “Cluster Protocols,” on page 93	Modify the Slave Watchdog and Master Watchdog parameters only when supervised by Novell Technical Support.
Section 7.3, “Configuring Cluster Protocol Properties,” on page 94	Modify the Slave Watchdog and Master Watchdog parameters only when supervised by Novell Technical Support.

E.3 June 27, 2012

Updates were made to the following sections. The changes are explained below.

- ♦ [Section E.3.1, “Configuring and Managing Cluster Resources,” on page 306](#)
- ♦ [Section E.3.2, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 306](#)

E.3.1 Configuring and Managing Cluster Resources

Location	Change
Section 9.1.2, “Using Parameters with Space Characters in a Cluster Script,” on page 128	This section is new.

E.3.2 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
“Expanding the Size of a Clustered Pool” on page 182	Corrected the procedures for an online pool expansion.

E.4 May 31, 2012

Updates were made to the following section. The changes are explained below.

- ♦ [Section E.4.1, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 307](#)

E.4.1 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
“NCP Server for Linux” on page 158	NCP Server does not support cross-protocol locks across a cluster migration or failover of the resource.
“Adding NFS Export for a Clustered Pool Resource” on page 175	This section is new.

E.5 April 30, 2012

Updates were made to the following sections. The changes are explained below.

- ♦ [Section E.5.1, “Configuring Cluster Policies and Priorities,” on page 307](#)
- ♦ [Section E.5.2, “Files for Novell Cluster Services,” on page 307](#)
- ♦ [Section E.5.3, “Managing Clusters,” on page 308](#)
- ♦ [Section E.5.4, “Planning for Novell Cluster Services,” on page 308](#)
- ♦ [Section E.5.5, “Troubleshooting Novell Cluster Services,” on page 308](#)

E.5.1 Configuring Cluster Policies and Priorities

Location	Change
Section 7.4, “Configuring Cluster Event Email Notification,” on page 95	Because both Postfix and the GWIA default to using port 25, you must configure the GWIA to bind exclusively to its resource’s secondary IP address in order to avoid a port conflict between Postfix and the GWIA. .

E.5.2 Files for Novell Cluster Services

Location	Change
Table B-1 on page 289	References to the hang-check timer were removed. It was obsoleted in OES 2 SP2 Patches. The NCS tolerance is now used to determine whether NCS has been starved of CPU. You can use directive <code>kernel.panic</code> in file <code>/etc/sysctl.conf</code> file to control the behaviors after NCS has determined to kill the node.

E.5.3 Managing Clusters

Location	Change
Table 8-2, “Cluster Resource States,” on page 102	Updated the Alert entry to identify alert types: Start alert, Failover alert, and Failback alert.
Section 8.16, “Configuring STONITH (Feature Preview),” on page 122	This section is new.

E.5.4 Planning for Novell Cluster Services

Location	Change
Section 4.3, “Volume ID Requirements,” on page 40	Volume IDs for BCC-enabled clustered volumes must be unique across all nodes in every peer cluster.
Section 4.10.3, “Modifying the Polling Interval, No Path Retry, and Failback Settings in the multipath.conf File,” on page 57	Added information about the polling_interval setting.
Section 4.6.8, “SFCB and CIMOM,” on page 48	The user name that you use to log in to Novell iManager when you manage a cluster and the BCC cluster must be a Novell eDirectory user name that has been LUM-enabled.

E.5.5 Troubleshooting Novell Cluster Services

Location	Change
“iManager Error 500 Occurs while Connecting to a Cluster” on page 266	This section is new.

E.6 January 31, 2012

Updates were made to the following sections. The changes are explained below.

- ♦ [Section E.6.1, “Configuring Cluster Resources for Shared NSS Pools and Volumes,” on page 309](#)
- ♦ [Section E.6.2, “Planning for Novell Cluster Services,” on page 309](#)
- ♦ [Section E.6.3, “Upgrading Clusters from OES 2 to OES 11,” on page 309](#)

E.6.1 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
“NCP Server for Linux” on page 158	NCP Server does not support cross-protocol locks across a cluster migration or failover of the resource.

E.6.2 Planning for Novell Cluster Services

Location	Change
“NCP Server for Linux” on page 50	NCP Server does not support cross-protocol locks across a cluster migration or failover of the resource.

E.6.3 Upgrading Clusters from OES 2 to OES 11

Location	Change
“Supported Upgrade Paths from OES 2 to OES 11” on page 87	This section is new.

