

Novell Cluster Services™ 2.1 for Linux Administration Guide

Open Enterprise Server 11 SP1

January 2014

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 © 2007–2014 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	15
1 Overview of Novell Cluster Services	17
1.1 Why Should I Use Clusters?	17
1.2 Benefits of Novell Cluster Services	17
1.3 Product Features	18
1.4 Clustering for High Availability	18
1.5 Shared Disk Scenarios	20
1.5.1 Using Fibre Channel Storage Systems	21
1.5.2 Using iSCSI Storage Systems	22
1.5.3 Using Shared SCSI Storage Systems	23
1.6 Terminology	23
1.6.1 The Cluster	23
1.6.2 Cluster Resources	24
1.6.3 Failover Planning	26
2 What's New or Changed for Novell Cluster Services	29
2.1 What's New (OES 11 SP1 September 2013 Patches)	29
2.2 What's New (OES 11 SP1 May 2013 Patches)	29
2.3 What's New (OES 11 May 2013 Patches)	30
2.4 What's New (OES 11 SP1 April 2013 Patches)	30
2.5 What's New (OES 11 April 2013 Patches)	31
2.6 What's New (OES 11 SP1 January 2013 Patches)	31
2.7 What's New (OES 11 January 2013 Patches)	32
2.8 What's New (OES 11 SP1 November 2012 Patches)	33
2.9 What's New (OES 11 September 2012 Patches)	33
2.10 What's New (OES 11 SP1)	33
2.11 What's New (OES 11)	37
3 Planning for a Cluster	41
3.1 Determining Your Design Criteria	41
3.2 Using Cluster Best Practices	41
3.3 Planning the LAN Connectivity	42
3.3.1 VLAN	42
3.3.2 Channel Bonding	43
3.3.3 Spanning Tree Protocol	43
3.3.4 IP Addresses	43
3.3.5 Name Resolution	43
3.4 Planning the Shared Storage Connectivity	44
3.5 Planning the Shared Storage Solution	44
3.6 Planning the eDirectory Deployment	44
3.6.1 Object Location	44
3.6.2 Cluster OU Context	45
3.6.3 Cluster OU Partitioning and Replication	45
3.7 Planning for Shared Storage as Cluster Resources	45
3.8 Planning for OES Services as Cluster Resources	46

4	Planning for Novell Cluster Services	49
4.1	Cluster Administration Requirements	49
4.1.1	Cluster Installation Administrator	49
4.1.2	NCS Proxy User	50
4.1.3	Cluster Administrator or Administrator-Equivalent User	52
4.2	IP Address Requirements	52
4.3	Volume ID Requirements	52
4.4	Hardware Requirements	53
4.5	Virtualization Environments	53
4.6	Software Requirements for Cluster Services	54
4.6.1	Novell Open Enterprise Server 11 SP1	54
4.6.2	Novell Cluster Services	54
4.6.3	Novell eDirectory 8.8.x	54
4.6.4	SLP	58
4.6.5	Novell iManager	59
4.6.6	Clusters Plug-in for iManager	60
4.6.7	Storage-Related Plug-Ins for iManager	61
4.6.8	SFCB and CIMOM	61
4.6.9	CASA 1.7	62
4.6.10	Web Browser	63
4.7	Software Requirements for Cluster Resources	63
4.7.1	NCP Server for Linux	63
4.7.2	Novell Storage Services File System for Linux	64
4.7.3	LVM Volume Groups and Linux POSIX File Systems	65
4.7.4	NCP Volumes on Linux POSIX File Systems	65
4.7.5	Dynamic Storage Technology Shadow Volume Pairs	65
4.7.6	NCP File Access	65
4.7.7	Novell AFP	65
4.7.8	Novell CIFS	66
4.7.9	Novell Samba	66
4.7.10	Novell Domain Services for Windows	66
4.8	Shared Disk Configuration Requirements	66
4.8.1	Shared Devices	67
4.8.2	SBD Partitions	67
4.8.3	Shared iSCSI Devices	69
4.8.4	Shared RAID Devices	69
4.9	SAN Rules for LUN Masking	69
4.10	Multipath I/O Configuration Requirements	70
4.10.1	Path Failover Settings for Device Mapper Multipath	70
4.10.2	Modifying the Port Down Retry Setting in the modprobe.conf.local File	70
4.10.3	Modifying the Polling Interval, No Path Retry, and Failback Settings in the multipath.conf File	71
4.10.4	Modifying the Port Down Retry and Link Down Retry Settings for an HBA BIOS	72
5	Installing and Configuring Novell Cluster Services on OES 11 SP1	73
5.1	Novell Cluster Services Licensing	73
5.2	Extending the eDirectory Schema to Add Cluster Objects	74
5.2.1	Prerequisites for Extending the Schema	74
5.2.2	Extending the Schema	75
5.3	Assigning Install Rights for Container Administrators or Non-Administrator Users	76
5.4	Installing Novell Cluster Services	77
5.4.1	Before You Install Novell Cluster Services	77
5.4.2	Installing Novell Cluster Services during an OES 11 SP1 Installation	78
5.4.3	Installing Novell Cluster Services on an Existing OES 11 SP1 Server	79
5.5	Configuring Novell Cluster Services	80
5.5.1	Initializing and Sharing a Device to Use for the SBD Partition	81

5.5.2	Opening the Novell Open Enterprise Server Configuration Page	84
5.5.3	Using Different LDAP Credentials for the Cluster Configuration	85
5.5.4	Accessing the Novell Cluster Services Configuration Page in YaST	85
5.5.5	Configuring a New Cluster	86
5.5.6	Adding a Node to an Existing Cluster	91
5.6	Adding a Node That Was Previously in the Cluster	95
5.7	Configuring Additional Administrators	96
5.8	Installing or Updating the Clusters Plug-in for iManager	96
5.8.1	Prerequisites for Installing or Updating the Clusters Plug-In	97
5.8.2	Installing the Clusters Plug-In	97
5.8.3	Uninstalling and Reinstalling the Clusters Plug-In	97
5.8.4	Updating Role-Based Services for the Clusters Plug-In for OES 11 SP1	98
5.9	Patching Novell Cluster Services	98
5.10	What's Next	99
6	Upgrading OES 11 Clusters	101
6.1	Requirements and Guidelines for Upgrading Clusters from OES 11 to OES 11 SP1	101
6.2	Adding OES 11 SP1 Nodes to an OES 11 Cluster (Rolling Cluster Upgrade)	102
6.3	Upgrade Issues for OES 11 SP1	104
6.3.1	Updating the Clusters Plug-in and Role-Based Services in Novell iManager	104
6.3.2	DHCP Cluster Resource: Path Change for dhcpd.pid	104
7	Upgrading Clusters from OES 2 SP3 to OES 11x	105
7.1	What's New and Changed for Clustered Storage	105
7.1.1	NSS Pools	105
7.1.2	Linux POSIX File Systems on an EVMS CSM (Compatibility Only)	106
7.1.3	Linux POSIX File Systems on an LVM Volume Group	107
7.2	Supported Upgrade Paths from OES 2 SP3 to OES 11x	107
7.3	Requirements and Guidelines for Upgrading Clusters from OES 2 SP3	108
7.4	Adding OES 11x Nodes to an OES 2 SP3 Cluster (Rolling Cluster Upgrade)	108
7.5	Updating the Clusters Plug-in for Novell iManager	112
7.6	Upgrading a Cluster with DST Resources from OES 2 SP3 to OES 11x	112
8	Configuring Cluster Policies, Protocols, and Properties	113
8.1	Understanding Cluster Settings	113
8.1.1	Cluster Policies	114
8.1.2	Cluster Priorities	115
8.1.3	Cluster Protocols	115
8.1.4	RME Groups	116
8.1.5	BCC	116
8.1.6	Cascading Failover Prevention	116
8.1.7	Cluster Node Reboot Behavior	116
8.1.8	STONITH	116
8.1.9	Cluster Information	116
8.2	Setting Up a Personalized List of Clusters to Manage	117
8.2.1	Adding Clusters to Your My Clusters List	117
8.2.2	Viewing Information about Clusters in Your My Clusters List	118
8.2.3	Personalizing the My Clusters Display	119
8.2.4	Removing Clusters from a My Clusters List	120
8.3	Configuring Quorum Membership and Timeout Policies	121
8.3.1	Quorum Membership (Number of Nodes)	122
8.3.2	Quorum Timeout	122
8.4	Configuring Cluster Protocols	122
8.4.1	Heartbeat	123

8.4.2	Tolerance	123
8.4.3	Master Watchdog	124
8.4.4	Slave Watchdog	124
8.4.5	Maximum Retransmits	124
8.5	Configuring Cluster Event Email Notification	124
8.5.1	Configuring an Email Port	124
8.5.2	Configuring the Notification Policies	125
8.6	Configuring Cascade Failover Prevention	126
8.6.1	Understanding the Cascade Failover Prevention Quarantine	126
8.6.2	Releasing a Resource from Quarantine	128
8.6.3	Enabling or Disabling Cascade Failover Prevention	128
8.7	Configuring the Cluster Node Reboot Behavior	129
8.8	Configuring STONITH	130
8.8.1	Sample Script for HP iLO	131
8.8.2	Sample Script for Web-Based Power Switches	131
8.9	Viewing or Modifying the NCS Proxy User Assignment in the NCS_Management_Group	132
8.10	Viewing or Modifying the Cluster IP Address or Port	133
8.11	Moving a Cluster, or Changing the Node IP Addresses, LDAP Servers, or Administrator Credentials for a Cluster	134
8.11.1	Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster	134
8.11.2	Moving a Cluster or Changing IP Addresses of Cluster Nodes and Resources	136
8.12	What's Next	138

9 Managing Clusters 139

9.1	Starting and Stopping Novell Cluster Services	140
9.1.1	Starting Novell Cluster Services	140
9.1.2	Stopping Novell Cluster Services	140
9.1.3	Enabling and Disabling the Automatic Start of Novell Cluster Services	141
9.2	Taking the Cluster Down	142
9.3	Monitoring Cluster and Resource States	142
9.4	Responding to a Start, Failover, or Failback Alert	145
9.5	Viewing the Cluster Event Log in iManager	146
9.6	Viewing the Cluster Event Log at the Command Line	149
9.7	Generating a Cluster Configuration Report	150
9.8	Cluster Migrating Resources to Different Nodes	151
9.9	Onlining and Offlining (Loading and Unloading) Cluster Resources from a Cluster Node	152
9.9.1	Guidelines for Onlining and Offlining Resources	153
9.9.2	Using iManager to Online and Offline Resources	153
9.9.3	Using the Cluster Online Command	153
9.9.4	Using the Cluster Offline Command	153
9.9.5	Using the Cluster Migrate Command	154
9.10	Removing (Leaving) a Node from the Cluster	154
9.10.1	Removing One Node at a Time	154
9.10.2	Removing All Nodes at the Same Time (Taking the Cluster Down)	154
9.11	Joining a Node to the Cluster (Rejoining a Cluster)	155
9.12	Shutting Down Novell Cluster Services When Servicing Shared Storage	155
9.13	Enabling or Disabling Cluster Maintenance Mode	156
9.14	Viewing the Cluster Node Properties	156
9.15	Deleting a Cluster Node from a Cluster	157
9.16	Creating or Deleting Cluster SBD Partitions	157
9.16.1	Requirements and Guidelines for Creating an SBD Partition	158
9.16.2	Before You Create a Cluster SBD Partition	160
9.16.3	Creating a Non-Mirrored Cluster SBD Partition with SBDUTIL	160
9.16.4	Mirroring an Existing SBD Partition with NSSMU	164
9.16.5	Creating a Mirrored Cluster SBD Partition with SBDUTIL	167
9.16.6	Removing a Segment from a Mirrored Cluster SBD Partition	170

9.16.7	Deleting a Non-Mirrored Cluster SBD Partition	172
9.16.8	Deleting a Mirrored Cluster SBD	173
9.16.9	Additional Information about SBD Partitions	174
9.17	Customizing Cluster Services Management	174

10 Configuring and Managing Cluster Resources 177

10.1	Planning for Cluster Resources	177
10.1.1	Naming Conventions for Cluster Resources	178
10.1.2	Using Parameter Values with Spaces in a Cluster Script	178
10.1.3	Using Double Quotation Marks in a Cluster Script	178
10.1.4	Script Length Limits	179
10.1.5	Planning Cluster Maintenance	179
10.1.6	Number of Resources	179
10.1.7	Linux POSIX File System Types	179
10.2	Setting Up a Personalized List of Resources to Manage	179
10.2.1	Adding Resources to a My Resources List	180
10.2.2	Viewing Information about Resources in a My Resources List	180
10.2.3	Personalizing the My Resources Display	182
10.2.4	Removing Resources from a My Resources List	183
10.2.5	Updating the My Resources List after Renaming or Deleting a Resource	184
10.3	Using Cluster Resource Templates	184
10.3.1	Default Resource Templates	184
10.3.2	Viewing or Modifying a Resource Template in iManager	185
10.3.3	Creating a Resource Template	187
10.3.4	Synchronizing Locally Modified Resource Templates with Templates in Novell eDirectory	189
10.4	Creating Cluster Resources	189
10.5	Configuring a Load Script for a Cluster Resource	190
10.6	Configuring an Unload Script for a Cluster Resource	191
10.7	Enabling Monitoring and Configuring the Monitor Script	192
10.7.1	Understanding Resource Monitoring	192
10.7.2	Configuring Resource Monitoring	194
10.7.3	Example Monitor Scripts	196
10.7.4	Monitoring Services that Are Critical to Clustering	196
10.8	Configuring the Start, Failover, and Failback Modes for Cluster Resources	197
10.8.1	Understanding Cluster Resource Modes	197
10.8.2	Setting the Start, Failover, and Failback Modes for a Resource	197
10.9	Configuring Preferred Nodes and Failover Order for a Resource	198
10.10	Configuring Resource Priorities for Load Order	200
10.11	Configuring Resource Mutual Exclusion Groups	201
10.11.1	Understanding Resource Mutual Exclusion	202
10.11.2	Setting Up RME Groups	204
10.11.3	Viewing RME Group Settings	204
10.12	Controlling Resource Monitoring	205
10.13	Changing the IP Address of a Cluster Resource	206
10.13.1	Pool Cluster Resource	206
10.13.2	Non-Pool Cluster Resources	207
10.14	Renaming a Cluster Resource	208
10.15	Deleting Cluster Resources, or Disabling Clustering for a Pool, LVM Volume Group, or Service	209
10.16	Additional Information for Creating Cluster Resources	212
10.16.1	Creating Storage Cluster Resources	212
10.16.2	Creating Service Cluster Resources	213
10.16.3	Creating Virtual Machine Cluster Resources	213

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

12.1	Requirements for Creating Pool Cluster Resources	218
12.1.1	Novell Cluster Services	218
12.1.2	Resource IP Address	218
12.1.3	Novell eDirectory	218
12.1.4	Novell Storage Services	218
12.1.5	Shared Storage	220
12.1.6	NCP Server for Linux	220
12.1.7	Novell CIFS for Linux	221
12.1.8	Novell AFP for Linux	221
12.1.9	Novell Samba	222
12.1.10	Domain Services for Windows	222
12.2	Guidelines for Using Pool Cluster Resources	222
12.2.1	Guidelines for Shared Pools	222
12.2.2	Guidelines for Shared Volumes	223
12.3	Initializing and Sharing a SAN Device	223
12.3.1	Initializing and Sharing a Device with iManager	224
12.3.2	Initializing and Sharing a Device with NSSMU	225
12.3.3	Initializing and Sharing a Device with NLVM	225
12.4	Creating Cluster-Enabled Pools and Volumes	227
12.4.1	Creating a Cluster-Enabled Pool and Volume with iManager	227
12.4.2	Creating a Cluster-Enabled Pool and Volume with NSSMU	230
12.5	Cluster-Enabling an Existing NSS Pool and Its Volumes	233
12.6	Configuring a Load Script for the Shared NSS Pool	240
12.6.1	Sample NSS Load Script	240
12.6.2	Specifying a Customized Mount Point for an NSS Volume	241
12.6.3	Specifying Customized Mount Options	241
12.6.4	Loading Novell AFP as an Advertising Protocol	242
12.6.5	Loading Novell CIFS as an Advertising Protocol	242
12.7	Configuring an Unload Script for the Shared NSS Pool	242
12.7.1	Sample NSS Unload Script	243
12.7.2	Using Volume Dismount Commands	243
12.7.3	Unloading Novell AFP as an Advertising Protocol	243
12.7.4	Unloading Novell CIFS as an Advertising Protocol	244
12.8	Configuring a Monitor Script for the Shared NSS Pool	244
12.8.1	Sample NSS Monitor Script	244
12.8.2	Monitoring Novell AFP as an Advertising Protocol	245
12.8.3	Monitoring Novell CIFS as an Advertising Protocol	245
12.9	Adding Advertising Protocols for NSS Pool Cluster Resources	245
12.10	Adding NFS Export for a Clustered Pool Resource	247
12.10.1	Sample NFS Load Script	247
12.10.2	Sample NFS Unload Script	248
12.11	Mirroring and Cluster-Enabling Shared NSS Pools and Volumes	248
12.11.1	Understanding NSS Mirroring	248
12.11.2	Requirements for NSS Mirroring	250
12.11.3	Creating and Mirroring NSS Pools on Shared Storage	250
12.11.4	Verifying the NSS Mirror Status in the Cluster	252
12.12	Renaming a Clustered NSS Pool	252
12.12.1	Renaming a Shared Pool with iManager	253
12.12.2	Renaming a Shared Pool with NSSMU	255
12.12.3	Renaming a Shared Pool and Its Resource Objects	258
12.12.4	Renaming a Shared Pool in a DST Cluster Resource	262
12.13	Renaming a Clustered NSS Volume	264
12.13.1	Renaming a Shared NSS Volume with iManager	265

12.13.2	Renaming a Shared NSS Volume with NSSMU	269
12.13.3	Renaming a Shared NSS Volume in a DST Cluster Resource	273
12.14	Renaming the Mount Point Path for a Clustered NSS Volume	275
12.15	Changing the Name Space for a Clustered NSS Volume	281
12.16	Adding a Volume to a Clustered Pool	282
12.17	Changing an Assigned Volume ID	283
12.18	Expanding the Size of a Clustered Pool	284
12.18.1	Planning to Expand a Clustered Pool	284
12.18.2	Increasing the Size of a Clustered Pool by Extending the Size of Its LUN	285
12.18.3	Increasing the Size of a Clustered Pool by Adding Space from a Different Shared LUN	291
12.18.4	Verifying the Expanded Pool Cluster Resource on All Nodes	295
12.19	Disabling Clustering for a Pool	296
12.20	Deleting and Re-Creating a Pool Cluster Resource (Disabling and Re-Enabling Clustering for a Pool)	298
12.21	Deleting a Clustered Pool	301
12.21.1	Deleting a Clustered Pool on the Master Node	302
12.21.2	Deleting a Clustered Pool on a Non-Master Node	303

13 Configuring and Managing Cluster Resources for Shared LVM Volume Groups 305

13.1	Requirements for Creating LVM Cluster Resources	305
13.1.1	Novell Cluster Services	306
13.1.2	Linux Logical Volume Manager 2 (LVM2)	306
13.1.3	Clustered Logical Volume Manager Daemon (CLVMD)	306
13.1.4	Resource IP Address	306
13.1.5	Shared Storage Devices	306
13.1.6	All Nodes Must Be Present	307
13.1.7	Working in Mixed Node OES Clusters	307
13.1.8	NCP File Access with Novell NCP Server	307
13.1.9	SMB/CIFS File Access with Novell Samba	308
13.1.10	Linux File Access Protocols	308
13.2	Initializing a SAN Device	308
13.2.1	Using NSSMU to Initialize a Device	309
13.2.2	Using NLVM Commands to Initialize a Device	311
13.3	Configuring an LVM Volume Group Cluster Resource with NSS Management Tools	312
13.3.1	Sample Values	312
13.3.2	Creating an LVM Volume Group Cluster Resource with NSSMU	314
13.3.3	Creating an LVM Volume Group Cluster Resource with NLVM Commands	318
13.3.4	Configuring the LVM Cluster Resource Settings	321
13.3.5	Viewing or Modifying the LVM Resource Scripts	324
13.3.6	Sample LVM Resource Scripts	326
13.4	Configuring an LVM Cluster Resource with LVM Commands and the Generic File System Template	330
13.4.1	Creating a Shared LVM Volume with LVM Commands	330
13.4.2	Creating a Generic File System Cluster Resource for an LVM Volume Group	335
13.4.3	Sample Generic LVM Resource Scripts	341
13.5	Creating a Virtual Server Object for an LVM Volume Group Cluster Resource	343
13.5.1	Creating an NCP Virtual Server for the LVM Resource	343
13.5.2	Adding NCP Virtual Server Commands to the Resource Scripts	344
13.5.3	Sample LVM Resource Scripts with an NCP Virtual Server	346
13.6	Enabling NCP File Access for a Clustered LVM Volume	348
13.6.1	Adding NCP Volume Commands to the Resource Scripts	348
13.6.2	Creating a Shared NCP Volume Object	351
13.6.3	Sample LVM Resource Scripts Enabled for NCP File Access	353
13.7	Modifying the LVM Resource Scripts for Other Services	355
13.8	Renaming the Mount Point Path for a Clustered LVM Volume	356

13.9	Disabling Clustering for an LVM Volume Group and Logical Volume	357
13.10	Deleting a Clustered LVM Volume Group and Logical Volume	362
13.11	Deleting a Clustered LVM Volume (Created in NSSMU or NLVM)	366
13.11.1	Deleting a Cluster-Enabled LVM Volume on the Master Node	366
13.11.2	Deleting a Cluster-Enabled LVM Volume on a Non-Master Node	367
13.12	Linux LVM Management Tools	368
13.12.1	Using the YaST Expert Partitioner	369
13.12.2	Using LVM Commands	369
14	Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers	371
14.1	Requirements for Using CSM Cluster Resources	371
14.2	Identifying CSM-Based Resources	372
14.3	Deporting the CSM Containers with Segment Managers	373
14.4	Modifying the Scripts for CSM Resources without a Segment Manager	374
14.4.1	Offlining the CSM Cluster Resources without Segment Managers	375
14.4.2	Configuring Scripts for a CSM Cluster Resource without a Segment Manager	375
14.4.3	Sample Load Script for a CSM Resource without a Segment Manager	377
14.4.4	Sample Unload Script for CSM Resource without a Segment Manager	377
14.4.5	Sample Monitor Script for a CSM Resource without a Segment Manager	378
14.5	Modifying the Scripts for CSM Resources with a Segment Manager	378
14.5.1	Offlining the CSM Cluster Resources with a Segment Manager	380
14.5.2	Configuring the Scripts for a CSM Cluster Resource with a Segment Manager	380
14.5.3	Sample Load Script for a CSM Resource with a Segment Manager	383
14.5.4	Sample Unload Script for a CSM Resource with a Segment Manager	384
14.5.5	Sample Monitor Script for a CSM Resource with a Segment Manager on It	385
14.6	Configuring and Adding OES 11 SP1 Nodes to the OES 2 SP3 Cluster	385
14.7	Configuring the Preferred Nodes and Cluster Settings for a CSM Cluster Resource	386
14.8	Verifying the Load and Unload Scripts	387
14.9	Creating an NCP Virtual Server Object for a CSM Cluster Resource	387
14.9.1	Creating the Virtual Server Object	388
14.9.2	Modifying the Load Script	389
14.9.3	Modifying the Unload Script	389
14.9.4	Activating the Script Changes	390
14.9.5	Verifying the Virtual Server	390
14.10	Deleting a Cluster Resource	390
15	Configuring Novell Cluster Services in a Virtualization Environment	391
15.1	Prerequisites for Xen Host Server Environments	392
15.2	Virtual Machines as Cluster Resources	392
15.2.1	Creating a Xen Virtual Machine Cluster Resource	393
15.2.2	Configuring Virtual Machine Load, Unload, and Monitor Scripts	394
15.2.3	Setting Up Live Migration	399
15.3	Virtual Machines as Cluster Nodes	400
15.4	Virtual Cluster Nodes in Separate Clusters	401
15.5	Mixed Physical and Virtual Node Clusters	401
15.6	Additional Information	403
16	Troubleshooting Novell Cluster Services	405
16.1	File Location Information	406
16.2	Version Issues	406
16.3	Diagnosing a Poison Pill Condition	406
16.4	Diagnosing Cluster Problems	407

16.5	Why didn't my resource go to the node I specified?	407
16.6	iManager: Error 500 Occurs while Connecting to a Cluster	408
16.7	iManager: Error 500 Occurs after Installing or Updating the Clusters Plug-In for iManager 2.7.5	408
16.8	iManager: Problems with Clusters Plug-In After Update for Novell iManager 2.7.5	408
16.9	A Device Name Is Required to Create a Cluster Partition	409
16.10	Cluster Resource Goes Comatose Immediately After Migration or Failover	409
16.11	Cannot Authenticate to Remote Servers during Cluster Configuration	409
16.12	Cannot Connect to an iSCSI Target.	409
16.13	Errors when Deleting a Cluster Resource or Clustered Pool	409
16.14	Is there a way to uninstall Novell Cluster Services from a server?	410
16.15	Running supportconfig Concurrently on Multiple Nodes Causes Cluster Problems	410
16.16	Error Bad XML: Cluster Search Times Out	410
16.17	Error NFS Stale Handle Reported on Host when Cluster Failover Occurs	410
16.18	Error 20897 - This node is not a cluster member.	411
16.19	Error 499 - Could not delete this resource Data_Server	411
16.20	Error 603 - smdr.novell is not registered with SLP for a new cluster resource.	411
16.21	Migration Authentication Error Occurs when Connecting to a New Cluster Resource on the Target Server	412
16.22	Communications Errors	412

17 Security Considerations 415

17.1	Cluster Administration Rights	415
17.2	Ports	415
17.3	Email Alerts	416
17.4	Log Files	416
17.5	Configuration Files	416

A Console Commands for Novell Cluster Services 417

A.1	Cluster Management Commands	417
A.2	Business Continuity Clustering Commands	423
A.3	extend_schema Command	423
	A.3.1 Syntax	423
	A.3.2 Options	424
	A.3.3 Example	424
A.4	SBD Utility	424
	A.4.1 Syntax	425
	A.4.2 Options	425
	A.4.3 Return Value	426
	A.4.4 Examples	426
A.5	Cluster Segment Manager Import/Export (csmport) Utility	427
	A.5.1 Syntax	427
	A.5.2 Options	427
	A.5.3 Examples	429
A.6	Configuration Update (ncs-configd.py) Script	433
A.7	Virtual NCS NCP Server Object (ncs_ncpserv.py) Script	433
	A.7.1 Syntax	434
	A.7.2 Options	434
	A.7.3 Examples	434
A.8	NPCON Mount Command for NCP and NSS Volumes	435
	A.8.1 Syntax	435
	A.8.2 NCP Volume Examples	435
	A.8.3 NSS Volume Examples	436
	A.8.4 Other Examples	436

B	Files for Novell Cluster Services	437
C	Electing a Master Node	441
C.1	Master-Election Algorithm for OES 11 and Earlier	441
C.2	New Master-Election Algorithm for OES 11 SP1	442
D	Comparing Novell Cluster Services for Linux and NetWare	445
E	Comparing Resources Support for Linux and NetWare	451
F	Clusters Plug-In Changes for Novell iManager 2.7.5	457
G	Documentation Updates	463
G.1	January 6, 2014	464
G.1.1	Upgrading OES 11 Clusters	464
G.2	December 6, 2013	464
G.2.1	Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes	464
G.2.2	Troubleshooting Novell Cluster Resources	464
G.3	September 2013 Scheduled Maintenance	464
G.3.1	Configuring and Managing Cluster Resources	465
G.3.2	Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes	465
G.3.3	Managing Clusters	465
G.3.4	Planning for Novell Cluster Services	466
G.3.5	What's New or Changed in Novell Cluster Services	466
G.4	September 21, 2013.	466
G.4.1	Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes	466
G.5	September 6, 2013.	467
G.5.1	Configuring and Managing Cluster Resources	467
G.5.2	Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes	467
G.6	August 16, 2013	468
G.6.1	Planning for Novell Cluster Services	468
G.7	August 6, 2013	468
G.7.1	Configuring Cluster Policies and Priorities.	468
G.8	July 30, 2013	468
G.8.1	Configuring and Managing Cluster Resources	469
G.8.2	Configuring and Managing Cluster Resources for Shared LVM Volume Groups	469
G.8.3	Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes	469
G.8.4	Configuring Cluster Policies and Priorities.	469
G.9	July 11, 2013	470
G.9.1	Planning for Novell Cluster Services	470
G.9.2	Upgrading Clusters from OES 2 SP3 to OES 11x.	470
G.10	May 2013 Scheduled Maintenance	470
G.10.1	Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes	471
G.10.2	Console Commands for Novell Cluster Services.	471
G.10.3	Managing Clusters.	471
G.10.4	Security Considerations.	471
G.10.5	Planning for Novell Cluster Services	472
G.10.6	What's New or Changed in Novell Cluster Services	472
G.11	May 21, 2013	472
G.11.1	Configuring and Managing Cluster Resources	472
G.11.2	Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes	473
G.11.3	What's New or Changed in Novell Cluster Services	473
G.12	May 7, 2013	473
G.12.1	Troubleshooting Novell Cluster Services.	473

G.13	April 2013 Scheduled Maintenance	473
G.13.1	Managing Clusters	473
G.13.2	Troubleshooting Novell Cluster Services	474
G.13.3	What's New or Changed in Novell Cluster Services	474
G.14	April 18, 2013	474
G.14.1	Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes	474
G.14.2	Troubleshooting Novell Cluster Services	475
G.14.3	Upgrading Clusters from OES 2 SP3 to OES 11x	475
G.14.4	What's New or Changed in Novell Cluster Services	475
G.15	April 12, 2013	475
G.15.1	Configuring and Managing Cluster Resources for Shared LVM Volume Groups	476
G.15.2	Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes	476
G.16	January 2013 Scheduled Maintenance	477
G.16.1	Clusters Plug-In Changes for Novell iManager 2.7.5	477
G.16.2	Configuring and Managing Cluster Resources	477
G.16.3	Upgrading Clusters from OES 2 SP3 to OES 11 SP1	477
G.16.4	What's New or Changed in Novell Cluster Services	477
G.17	November 2012 Scheduled Maintenance	478
G.17.1	Configuring Cluster Policies, Protocols, and Properties	478
G.17.2	Configuring and Managing Cluster Resources	478
G.17.3	Configuring and Managing Cluster Resources for Shared LVM Volume Groups	478
G.17.4	Managing Clusters	479
G.17.5	Upgrading OES 11 Clusters	479
G.17.6	What's New or Changed in Novell Cluster Services	479
G.18	September 2012 Scheduled Maintenance	479
G.18.1	Configuring and Managing Cluster Resources	479
G.18.2	Planning for Novell Cluster Services	480
G.18.3	Troubleshooting Novell Cluster Services	480
G.18.4	What's New or Changed in Novell Cluster Services	480
G.19	September 12, 2012	480
G.19.1	Configuring and Managing Cluster Resources	481
G.19.2	Configuring Cluster Policies, Protocols, and Priorities	481
G.19.3	Installing and Configuring Novell Cluster Services	481
G.19.4	Troubleshooting Novell Cluster Services	481
G.20	August 28, 2012 (OES 11 SP1)	481
G.20.1	Console Commands for Novell Cluster Services	482
G.20.2	Comparing Novell Cluster Services for Linux and NetWare	482
G.20.3	Configuring and Managing Cluster Resources	482
G.20.4	Configuring and Managing Cluster Resources for Shared Linux LVM Volume Groups	483
G.20.5	Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes	483
G.20.6	Configuring Cluster Policies, Protocols, and Priorities	484
G.20.7	Configuring Novell Cluster Services in a Virtualization Environment	484
G.20.8	Electing a Master Node	484
G.20.9	Files for Novell Cluster Services	485
G.20.10	Installing and Configuring Novell Cluster Services on OES 11	485
G.20.11	Managing Clusters	485
G.20.12	Planning for Novell Cluster Services	485
G.20.13	Security Considerations	486
G.20.14	Troubleshooting Novell Cluster Services	486
G.20.15	Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers	486
G.20.16	Upgrading Clusters from OES 2 SP3 to OES 11 SP1	486
G.20.17	Upgrading OES 11 Clusters	487
G.20.18	What's New or Changed in Novell Cluster Services	487

About This Guide

This guide describes how to install, upgrade, configure, and manage Novell Cluster Services 2.1 for Novell Open Enterprise Server (OES) 11 Support Pack 1 (SP1). It is divided into the following sections:

- ♦ Chapter 1, “Overview of Novell Cluster Services,” on page 17
- ♦ Chapter 2, “What’s New or Changed for Novell Cluster Services,” on page 29
- ♦ Chapter 3, “Planning for a Cluster,” on page 41
- ♦ Chapter 4, “Planning for Novell Cluster Services,” on page 49
- ♦ Chapter 5, “Installing and Configuring Novell Cluster Services on OES 11 SP1,” on page 73
- ♦ Chapter 6, “Upgrading OES 11 Clusters,” on page 101
- ♦ Chapter 7, “Upgrading Clusters from OES 2 SP3 to OES 11x,” on page 105
- ♦ Chapter 8, “Configuring Cluster Policies, Protocols, and Properties,” on page 113
- ♦ Chapter 9, “Managing Clusters,” on page 139
- ♦ Chapter 10, “Configuring and Managing Cluster Resources,” on page 177
- ♦ Chapter 11, “Creating and Managing Service Cluster Resources,” on page 215
- ♦ Chapter 12, “Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes,” on page 217
- ♦ Chapter 13, “Configuring and Managing Cluster Resources for Shared LVM Volume Groups,” on page 305
- ♦ Chapter 14, “Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers,” on page 371
- ♦ Chapter 15, “Configuring Novell Cluster Services in a Virtualization Environment,” on page 391
- ♦ Chapter 16, “Troubleshooting Novell Cluster Services,” on page 405
- ♦ Chapter 17, “Security Considerations,” on page 415
- ♦ Appendix A, “Console Commands for Novell Cluster Services,” on page 417
- ♦ Appendix B, “Files for Novell Cluster Services,” on page 437
- ♦ Appendix C, “Electing a Master Node,” on page 441
- ♦ Appendix D, “Comparing Novell Cluster Services for Linux and NetWare,” on page 445
- ♦ Appendix E, “Comparing Resources Support for Linux and NetWare,” on page 451
- ♦ Appendix F, “Clusters Plug-In Changes for Novell iManager 2.7.5,” on page 457
- ♦ Appendix G, “Documentation Updates,” on page 463

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 SP1: Novell Cluster Services 2.1 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 17](#)
- ♦ [Section 1.2, “Benefits of Novell Cluster Services,” on page 17](#)
- ♦ [Section 1.3, “Product Features,” on page 18](#)
- ♦ [Section 1.4, “Clustering for High Availability,” on page 18](#)
- ♦ [Section 1.5, “Shared Disk Scenarios,” on page 20](#)
- ♦ [Section 1.6, “Terminology,” on page 23](#)

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 SP1 servers. You can configure up to 32 OES 11 SP1 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 stand-alone servers:

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

1.3 Product Features

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:

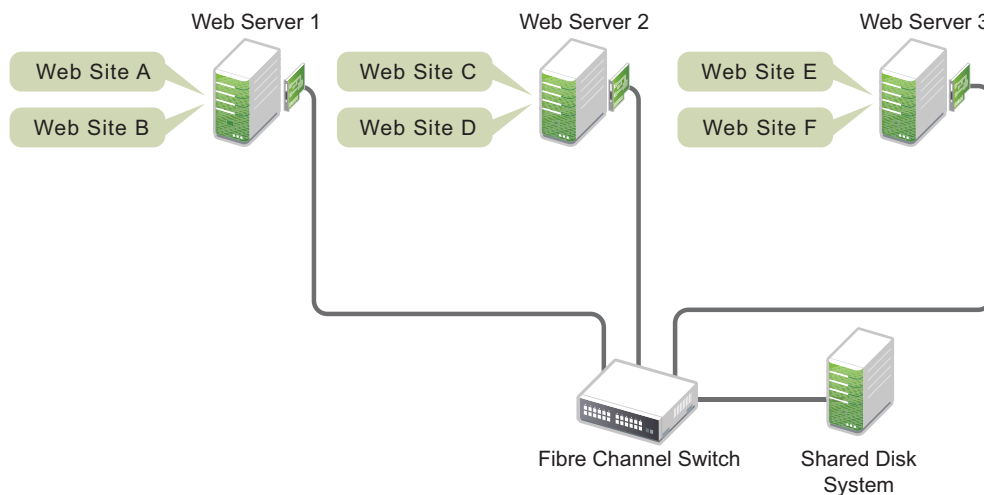
- ♦ 2 to 32 OES 11 SP1 servers, each containing at least one local disk device.
- ♦ Novell Cluster Services software running on each Linux server in the cluster.
- ♦ A shared disk subsystem connected to all servers in the cluster (optional, but recommended for most configurations).

- ♦ 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.

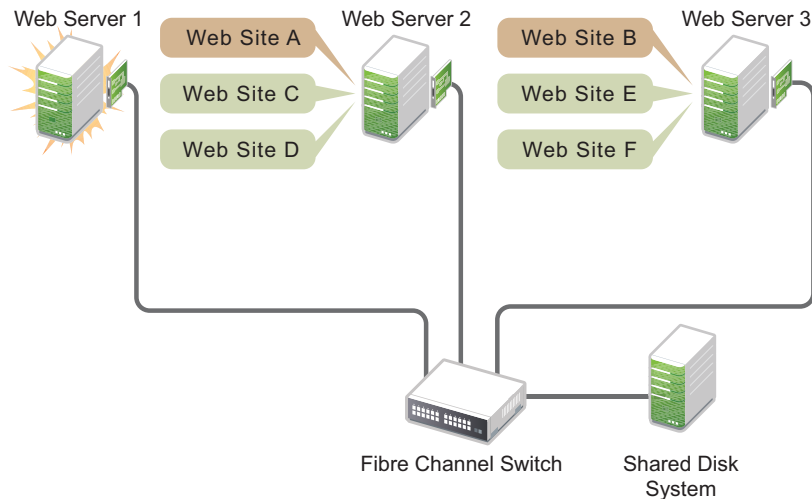
Figure 1-1 *Three-Server Cluster*



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 automatically fail back (that is, they are moved back to Web Server 1) if failback is configured for those resources, and Web Server operation returns to the way it was before Web Server 1 failed.

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 manually move Web Site A or Web Site B from Web Server 1 to either of the other servers in the cluster. You might want to do this to upgrade or perform scheduled maintenance on Web Server 1, or to increase performance or accessibility of the 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 21](#)
- ♦ [Section 1.5.2, “Using iSCSI Storage Systems,” on page 22](#)
- ♦ [Section 1.5.3, “Using Shared SCSI Storage Systems,” on page 23](#)

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 23](#)
- ♦ [Section 1.6.2, “Cluster Resources,” on page 24](#)
- ♦ [Section 1.6.3, “Failover Planning,” on page 26](#)

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 24](#)
- ♦ [“Server IP Address” on page 24](#)
- ♦ [“Master Node” on page 24](#)
- ♦ [“Slave Node” on page 24](#)
- ♦ [“Split-Brain Detector \(SBD\)” on page 24](#)
- ♦ [“Shared Storage” on page 24](#)

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 Novell eDirectory. If the master node fails, Cluster Services migrates the cluster IP address to another server in the cluster, and that server becomes the master node. For information about how a new master is determined, see [Appendix C, “Electing a Master Node,” on page 441](#).

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 one node at a time.

- ♦ [“Resource IP Address” on page 25](#)
- ♦ [“NCS Virtual Server” on page 25](#)
- ♦ [“Resource Templates” on page 25](#)
- ♦ [“Service Cluster Resource” on page 25](#)
- ♦ [“Pool Cluster Resource” on page 25](#)
- ♦ [“Linux POSIX Volume Cluster Resource” on page 26](#)

- ♦ [“NCP Volume Cluster Resource” on page 26](#)
- ♦ [“DST Volume Cluster Resource” on page 26](#)
- ♦ [“Cluster Resource Scripts” on page 26](#)

Resource IP Address

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

NCS Virtual Server

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

Resource Templates

A resource template contains the default load, unload, and monitor scripts and default settings for service or file system cluster resources. Resource templates are available for the following OES 11 SP1 Services and file systems:

- ♦ Novell Archive and Version Services
- ♦ 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 10.3, “Using Cluster Resource Templates,” on page 184](#).

Service Cluster Resource

An application or OES 11 SP1 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 13.5, “Creating a Virtual Server Object for an LVM Volume Group Cluster Resource,”](#) on page 343.

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 13.5, “Creating a Virtual Server Object for an LVM Volume Group Cluster Resource,”](#) on page 343.

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 27
- ♦ [“Quorum”](#) on page 27
- ♦ [“Preferred Nodes”](#) on page 27
- ♦ [“Resource Priority”](#) on page 27
- ♦ [“Resource Mutual Exclusion Groups”](#) on page 27
- ♦ [“Failover”](#) on page 27
- ♦ [“Fan-Out Failover”](#) on page 27
- ♦ [“Failback”](#) on page 27
- ♦ [“Cluster Migrate”](#) on page 27
- ♦ [“Leave a Cluster”](#) on page 27
- ♦ [“Join a Cluster”](#) on page 28

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 cluster migrated to other nodes in their preferred nodes list.

Join a Cluster

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

2 What's New or Changed for Novell Cluster Services

This section describes enhancements and changes in Novell Cluster Services since the initial release of version 2.0 in Novell Open Enterprise Server (OES) 11.

- [Section 2.1, “What’s New \(OES 11 SP1 September 2013 Patches\),” on page 29](#)
- [Section 2.2, “What’s New \(OES 11 SP1 May 2013 Patches\),” on page 29](#)
- [Section 2.3, “What’s New \(OES 11 May 2013 Patches\),” on page 30](#)
- [Section 2.4, “What’s New \(OES 11 SP1 April 2013 Patches\),” on page 30](#)
- [Section 2.5, “What’s New \(OES 11 April 2013 Patches\),” on page 31](#)
- [Section 2.6, “What’s New \(OES 11 SP1 January 2013 Patches\),” on page 31](#)
- [Section 2.7, “What’s New \(OES 11 January 2013 Patches\),” on page 32](#)
- [Section 2.8, “What’s New \(OES 11 SP1 November 2012 Patches\),” on page 33](#)
- [Section 2.9, “What’s New \(OES 11 September 2012 Patches\),” on page 33](#)
- [Section 2.10, “What’s New \(OES 11 SP1\),” on page 33](#)
- [Section 2.11, “What’s New \(OES 11\),” on page 37](#)

2.1 What's New (OES 11 SP1 September 2013 Patches)

In addition to bug fixes, Novell Cluster Services provides the following enhancement and behavior changes in the September 2013 Scheduled Maintenance for OES 11 SP1:

Unsharing a Device without Stopping NCS

NSS tools have been modified so that it is not necessary to stop NCS in order to disable the Shareable for Clustering setting. The device should not contain partitions that are used by shared pools, clustered pools, or SBDs. This change is in keeping with how unsharing a device works on OES 2 SP3 and earlier.

Support for VLAN on NIC Bonding

Novell Cluster Services added support for VLAN on NIC bonding. No modifications to scripts are required. You can use `ethx` or `vlanx` interfaces in any combination for nodes in a cluster.

2.2 What's New (OES 11 SP1 May 2013 Patches)

In addition to bug fixes, Novell Cluster Services provides the following enhancement and behavior changes in the May 2013 Scheduled Maintenance for OES 11 SP1:

Updated the DHCP PID File Location in the DHCP_Template

In the DHCP_Template for DHCP cluster resources, the PID file location was changed to `/var/lib/dhcp/var/run/dhcpd.pid`. This is the location used by DHCP since OES 11 SP1. The change applies automatically to any newly created DHCP cluster resources. For information about configuring DHCP cluster resources, see “Configuring DHCP with Novell Cluster Services for the Linux File System” in the *OES 11 SP1: Novell DNS/DHCP Services for Linux Administration Guide*.

SLP Refresh Interval for Cluster Resource Virtual NCP Servers

NCP Server was modified to refresh its OpenSLP registration of cluster resource virtual NCP servers based on the setting for the eDirectory advertise-life-time (`n4u.nds.advertise-life-time`) parameter. The `n4u.nds.advertise-life-time` parameter is set by default to 3600 seconds (1 hour) and has a valid range of 1 to 65535 seconds. Previously, NCP Server re-registered the virtual NCP servers with SLP every 30 minutes regardless of the eDirectory advertise-life-time setting. For information about setting the eDirectory advertise-life-time parameter in a cluster, see “SLP” (http://www.novell.com/documentation/oes11/clus_admin_lx/data/ncs_swreqs.html#slp) in the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

2.3 What’s New (OES 11 May 2013 Patches)

In addition to bug fixes, Novell Cluster Services provides the following enhancement and behavior changes in the May 2013 Scheduled Maintenance for OES 11:

SLP Refresh Interval for Cluster Resource Virtual NCP Servers

NCP Server was modified to refresh its OpenSLP registration of cluster resource virtual NCP servers based on the setting for the eDirectory advertise-life-time (`n4u.nds.advertise-life-time`) parameter. The `n4u.nds.advertise-life-time` parameter is set by default to 3600 seconds (1 hour) and has a valid range of 1 to 65535 seconds. Previously, NCP Server re-registered the virtual NCP servers with SLP every 30 minutes regardless of the eDirectory advertise-life-time setting. For information about setting the eDirectory advertise-life-time parameter in a cluster, see “SLP” (http://www.novell.com/documentation/oes11/clus_admin_lx/data/ncs_swreqs.html#slp) in the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

2.4 What’s New (OES 11 SP1 April 2013 Patches)

In addition to bug fixes, Novell Cluster Services provides the following enhancement in the April 2013 Scheduled Maintenance for OES 11 SP1:

Changing the Nodes or Failover Order in the Preferred Nodes List

When you change the nodes or failover order in a resource’s *Preferred Nodes* list, the change takes effect immediately when you click *OK* or *Apply* if the resource is running. Otherwise, it takes effect the next time the resource is brought online. Previously, the change did not occur until the cluster resource was taken offline and brought online.

2.5 What's New (OES 11 April 2013 Patches)

Upgrade to eDirectory 8.8.7

An upgrade to Novell eDirectory 8.8 SP7 is available in the April 2013 Scheduled Maintenance for OES 11. For information about the eDirectory upgrade, see [TID 7011599 \(http://www.novell.com/support/kb/doc.php?id=7011599\)](http://www.novell.com/support/kb/doc.php?id=7011599) in the Novell Knowledgebase.

There will be no further eDirectory 8.8 SP6 patches for the OES platform. Previous patches for Novell eDirectory 8.8 SP6 are available on [Novell Patch Finder \(http://download.novell.com/patch/finder/#familyId=112&productId=29503\)](http://download.novell.com/patch/finder/#familyId=112&productId=29503).

For information about configuring and managing eDirectory servers, see the [Novell eDirectory 8.8 SP7 documentation Web site \(http://www.netiq.com/documentation/edir887\)](http://www.netiq.com/documentation/edir887).

2.6 What's New (OES 11 SP1 January 2013 Patches)

Upgrade to Novell iManager 2.7.6

The January 2013 Scheduled Maintenance for OES 11 SP1 includes a channel upgrade from Novell iManager 2.7.5 to Novell iManager 2.7.6.

Novell iManager 2.7.6 provides the following enhancements:

- Microsoft Internet Explorer 10 certification in the desktop user interface view on Windows 8 (excluding Windows 8 RT) and Windows Server 2012.
- Apple Safari 6.0 certification on Mac OS X Mountain Lion (version 10.8).
- iManager Workstation certification on Windows 8 Enterprise Edition (32-bit and 64-bit).
- Manager 2.7.6 support for Tomcat 7.0.32, and Java 1.7.0_04 versions.

iManager documentation links in this guide have been updated to reflect this change.

iManager 2.7.6 and earlier documentation is available at <https://www.netiq.com/documentation/imanager27/>. For earlier iManager versions, see ["Previous Releases" \(https://www.netiq.com/documentation/imanager27/#prev\)](https://www.netiq.com/documentation/imanager27/#prev).

Novell Client Support for Windows 8 and Server 2012

The January 2013 Scheduled Maintenance for OES 11 SP1 announces the availability of Novell Client 2 SP3 for Windows with support for:

- Windows 8 (32-bit and 64-bit) excluding Windows 8 RT
- Windows Server 2012 (64-bit)

Novell Client 2 documentation links in this guide have been updated to reflect the release of SP3.

Novell Client 2 SP3 for Windows documentation is available on the [Web \(http://www.novell.com/documentation/windows_client/\)](http://www.novell.com/documentation/windows_client/). Documentation for earlier versions is available under [Previous Releases \(http://www.novell.com/documentation/windows_client/#previous\)](http://www.novell.com/documentation/windows_client/#previous).

2.7 What's New (OES 11 January 2013 Patches)

Upgrade to Novell iManager 2.7.6

The January 2013 Scheduled Maintenance for OES 11 includes a channel upgrade from Novell iManager 2.7.5 to Novell iManager 2.7.6.

Novell iManager 2.7.6 provides the following enhancements:

- Microsoft Internet Explorer 10 certification in the desktop user interface view on Windows 8 (excluding Windows 8 RT) and Windows Server 2012.
- Apple Safari 6.0 certification on Mac OS X Mountain Lion (version 10.8).
- iManager Workstation certification on Windows 8 Enterprise Edition (32-bit and 64-bit).
- Manager 2.7.6 support for Tomcat 7.0.32. and Java 1.7.0_04 versions.

iManager documentation links in this guide have been updated to reflect this change.

iManager 2.7.6 documentation is available at <https://www.netiq.com/documentation/imanager27/>. For earlier iManager versions, see “Previous Releases” (<https://www.netiq.com/documentation/imanager27/#prev>).

Novell Client Support for Windows 8 and Server 2012

The January 2013 Scheduled Maintenance for OES 11 announces the availability of Novell Client 2 SP3 for Windows with support for:

- Windows 8 (32-bit and 64-bit) excluding Windows 8 RT
- Windows Server 2012 (64-bit)

Novell Client 2 documentation links in this guide have been updated to reflect the release of SP3.

Novell Client 2 SP3 for Windows documentation is available on the [Web \(http://www.novell.com/documentation/windows_client/\)](http://www.novell.com/documentation/windows_client/). Documentation for earlier versions is available under [Previous Releases \(http://www.novell.com/documentation/windows_client/#previous\)](http://www.novell.com/documentation/windows_client/#previous).

Novell Cluster Services

In addition to bug fixes, Novell Cluster Services provides the following enhancement in the January 2013 Scheduled Maintenance for OES 11:

Clusters Plug-in for iManager 2.7.5 and Later

The Clusters plug-in for Novell iManager 2.7.5 or later was introduced in the OES 11 SP1 release. It supports the management of OES and NetWare clusters and resources. The availability of cluster management features depends on the version of Novell Cluster Services and server platform that are installed on the cluster being managed. A comparison of the old and new interface is available in “What's New (January 2013 Patches)” (http://www.novell.com/documentation/oes2/clus_admin_lx/data/ncs_new_jan2013.html) in the *OES 2 SP3: Novell Cluster Services 1.8.8 Administration Guide for Linux* (http://www.novell.com/documentation/oes2/clus_admin_lx/data/h4hgu4hs.html).

2.8 What's New (OES 11 SP1 November 2012 Patches)

In addition to bug fixes, Novell Cluster Services provides the following enhancement and behavior changes in the November 2012 Scheduled Maintenance for OES 11 SP1:

Enforcing Name Restrictions for NCP Enabled LVM Volumes

For NCP enabled LVM volume names, the NSSMU utility and `nlvm create linux` command have been modified to enforce the length restriction of 14 characters and to use uppercase letters for the name in Novell Cluster Services cluster resource scripts for a shared volume.

2.9 What's New (OES 11 September 2012 Patches)

In addition to bug fixes, Novell Cluster Services provides the following enhancement and behavior change in the September 2012 Scheduled Maintenance for OES 11:

Master Election Process

Novell Cluster Services introduces some intelligence in the master election process when the master leaves a cluster (voluntarily or involuntarily). The same master is elected under both the old and new algorithms, but the conclusion is reached sooner, especially for larger clusters. The new algorithm substantially reduces the time needed for master election in some cases.

This algorithm was first introduced in OES 11 SP1. For information about the old and new algorithms, see “Electing a Master Node” (http://www.novell.com/documentation/oes11/clus_admin_lx/data/ncs_master_election.html) in the *OES 11 SP1: Novell Cluster Services 2.1 NetWare to Linux Conversion Guide* (http://www.novell.com/documentation/oes11/clus_conversion_lx/data/bookinfo.html).

2.10 What's New (OES 11 SP1)

Novell Cluster Services 2.1 supports OES 11 SP1 services and file systems running on 64-bit SUSE Linux Enterprise Server (SLES) 11 SP2. In addition to bug fixes, Novell Cluster Services provides the following enhancements and behavior changes in the OES 11 SP1 release:

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 *Online Resource after Creation* option on the Cluster Pool Information page is available and deselected by default. If the pool is deactive when you cluster-enable it, you can select this option to automatically bring the resource online, which activates the pool on a preferred node and mounts the pool's volumes.

Previously, the option was deselected and dimmed, the resource was automatically placed in an offline state, and no action was taken for the state of the pool and its volumes.

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 2.1 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

Default CIFS Virtual Server Name

When you enable Novell CIFS as an advertising protocol for a cluster-enabled pool, the default name is based on the NCP virtual server object name. If you change the NCP virtual server name in NSSMU or in the Storage plug-in for iManager, the CIFS virtual server name is updated accordingly. For information, see “Novell CIFS for Linux” (http://www.novell.com/documentation/oes11/clus_admin_lx/data/bff13ks.html#bgh0tjp) in the *OES 11: Novell Cluster Services 2.1 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

Enabling NCP for a Clustered Linux POSIX Volume Group

When you create a cluster-enabled Linux POSIX volume group with NSSMU or with the `nlvm create linux volume` command, you can use the `ncp` option to enable users to access the Linux POSIX file system on the volume with the NetWare Core Protocol (NCP). Previously, the NCP volume (share) was created manually after the resource was created and brought online.

For information, see “Configuring an LVM Volume Group Cluster Resource with NSS Management Tools” (http://www.novell.com/documentation/oes11/clus_admin_lx/data/bxv3rde.html) in the *OES 11: Novell Cluster Services 2.1 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

Btrfs File System

The SUSE Linux Enterprise Server 11 SP2 operating system supports the Btrfs file system. The `btrfsprogs` package must be installed.

Novell Cluster Services supports using the Btrfs file system when you create a clustered Linux LVM volume group by using the NSS Management Utility (NSSMU) or Novell Linux Volume Manager (NLVM) commands. If you use Linux commands to create the LVM volume group, you can use the `mkfs.btrfs(8)` command to add the Btrfs file system to the volume. You specify `btrfs` as the file system type for the LVM volume group cluster resource. You can enable NCP file access for the Btrfs volume.

For information, see “Configuring an LVM Volume Group Cluster Resource with NSS Management Tools” (http://www.novell.com/documentation/oes11/clus_admin_lx/data/bxv3rde.html) in the *OES 11: Novell Cluster Services 2.1 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

Cluster Migrate Command

The `cluster migrate` command now offers the `-list`, `-most`, and `-next` options. For information, see “MIGRATE” in the Cluster Management Commands table (http://www.novell.com/documentation/oes11/clus_admin_lx/data/bfktr09.html) of the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

Clusters Plug-In for Novell iManager

The Clusters plug-in for Novell iManager 2.7.5 has been reorganized. In *Roles and Tasks* under *Clusters*, the *Cluster Manager*, *BCC Manager*, *Cluster Event Log*, and *Cluster Options* menu items have been replaced with two options: *My Clusters* and *My Resources*. Links on these pages take you to the familiar cluster options presented as tabs: *Cluster Manager*, *BCC Manager*, *Cluster Event Log*, and *Cluster Options*.

- ♦ **My Clusters:** The logged-in cluster administrator can set up a personalized list of clusters to manage. This allows an administrator to view at a glance the status of multiple clusters. An administrator can also customize the display to sort the entries, modify the columns, or filter the entries. The list of clusters and display preferences persist between the administrator's logins to iManager on the same server. For information, see "[Setting Up a Personalized List of Clusters to Manage](http://www.novell.com/documentation/oes11/clus_admin_lx/data/myclusters.html)" (http://www.novell.com/documentation/oes11/clus_admin_lx/data/myclusters.html) in the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).
- ♦ **My Resources:** The logged-in cluster administrator can set up a personalized list of cluster resources to manage. This allows an administrator to view at a glance the status of multiple cluster resources for multiple clusters. An administrator can also customize the display to sort the entries, modify the columns, or filter the entries. The list of cluster resources and display preferences persist between the administrator's logins to iManager on the same server. For information, see "[Setting Up a Personalized List of Resources to Manage](http://www.novell.com/documentation/oes11/clus_admin_lx/data/myresources.html)" (http://www.novell.com/documentation/oes11/clus_admin_lx/data/myresources.html) in the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

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

If you use Role-Based Services (RBS), upgrading the Clusters plug-in does not automatically update the RBS settings. The RBS Configuration page reports that the Clusters plug-in is out-of-date. The plug-in must be reinstalled on the RBS Configuration page in order to pick up the *My Clusters* and *My Resources* menu options. For information, see "[Updating Role-Based Services for the Clusters Plug-in for OES 11 SP1](http://www.novell.com/documentation/oes11/clus_admin_lx/data/inst_clusters.html#inst_clusters_rb)" (http://www.novell.com/documentation/oes11/clus_admin_lx/data/inst_clusters.html#inst_clusters_rb) in the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

STONITH

The STONITH (shoot-the-other-node-in-the-head) capability allows Novell Cluster Services to kill a suspect node by using remote power control. Unlike a poison pill, it does not require a response from the suspect node. STONITH is used after a poison pill is issued; it does not replace the poison pill. For information, see "[Configuring STONITH](http://www.novell.com/documentation/oes11/clus_admin_lx/data/stonith.html)" (http://www.novell.com/documentation/oes11/clus_admin_lx/data/stonith.html) in the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

Master Election Process

Novell Cluster Services for OES 11 SP1 introduces some intelligence in the master election process when the master leaves a cluster (voluntarily or involuntarily). The same master is elected under both the old and new algorithms, but the conclusion is reached sooner, especially for larger clusters. The new algorithm substantially reduces the time needed for master election in some cases.

For information about the old and new algorithms, see “Electing a Master Node” (http://www.novell.com/documentation/oes11/clus_admin_lx/data/ncs_master_election.html) in the *OES 11 SP1: Novell Cluster Services 2.1 NetWare to Linux Conversion Guide* (http://www.novell.com/documentation/oes11/clus_conversion_lx/data/bookinfo.html).

NetWare to Linux Cluster Conversion: File Access Protocol Commands for Monitor Scripts

When you convert NetWare 6.5 SP8 clusters to OES 11 SP1, if the AFP or CIFS protocols are configured as advertising protocols for an NSS pool cluster resource on the NetWare cluster, the corresponding monitoring commands are automatically added to the monitor script for the resource on OES 11 SP1. Monitoring is not available for NetWare clusters and is disabled while the cluster is in mixed mode. Monitoring can be enabled for the resource after the final cluster conversion is complete. For information, see “File Access Protocols for Monitor Scripts” (http://www.novell.com/documentation/oes11/clus_conversion_lx/data/bh9em7s.html#b115ip7x) in the *OES 11 SP1: Novell Cluster Services 2.1 NetWare to Linux Conversion Guide* (http://www.novell.com/documentation/oes11/clus_conversion_lx/data/bookinfo.html).

DHCP Cluster Resource: Path Change for dhcpd.pid

The DHCP process ID file location was changed from `/var/run/dhcpd.pid` to `/var/lib/dhcp/var/run/dhcpd.pid`. During a cluster upgrade from OES 11 to OES 11 SP1, you must modify the location of the `dhcpd.pid` file in the DHCP cluster resource’s unload script. After you modify the script, you should bring the resource online only on OES 11 SP1 nodes. This is necessary to ensure a graceful shutdown of the `dhcpd.pid` process when the DHCP resource fails over to a different node. For information, see “Path Change for dhcpd.pid” and “Changing the Path for dhcpd.pid” in the *OES 11 SP1: Novell DNS/DHCP Services for Linux Administration Guide*.

Change in the Command Sequence for Pool Cluster Resource Scripts

The default sequence of commands has changed for the load script and unload script of newly created NSS pool cluster resources. In the load script, the secondary IP address is added after the volumes are mounted. In the unload script, the secondary IP address is removed before the pool is deactivated.

The new sequence improves the handling of client reconnects because it leverages the clients’ auto-reconnect mechanisms. With the previous order, clients might attempt to connect to the IP address before the data is ready or after the data is unavailable, which results in an `Object Not Found` error. Some clients might automatically try again to connect, but some do not. Thus, it is more efficient to add the IP address after data is available and remove it before the data is not available.

Newly created pool cluster resources use the following sequence for commands in a load script. The order is reversed in an unload script.

NSS

[NFS] (If NFS is used, you manually add the command here.)

Secondary IP address

NCP

[AFP]

[CIFS]

Existing scripts are not modified during a cluster upgrade. We recommend that you modify the command sequence accordingly in your existing load scripts and unload scripts for pool cluster resources.

2.11 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 backward 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-node 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, Ext3, 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, Linux POSIX volume cluster resources use a Cluster Segment Manager (CSM) container on devices that are managed by EVMS. Because EVMS has been deprecated in OES 11, you must modify the scripts and cluster settings so that the CSM containers 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 the resource is configured to run on OES 11, you must configure its Preferred Nodes list to ensure that it can fail over only to OES 11 nodes in the mixed-node cluster. 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 host names 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).

Cascading Failover Prevention

Novell Cluster Services added the cascading failover prevention function that detects if a node has failed because of a bad cluster resource and prevents that bad resource from failing over to other servers in the cluster. This function was previously available on NetWare, but not on Linux.

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 49](#).

- ♦ [Section 3.1, “Determining Your Design Criteria,” on page 41](#)
- ♦ [Section 3.2, “Using Cluster Best Practices,” on page 41](#)
- ♦ [Section 3.3, “Planning the LAN Connectivity,” on page 42](#)
- ♦ [Section 3.4, “Planning the Shared Storage Connectivity,” on page 44](#)
- ♦ [Section 3.5, “Planning the Shared Storage Solution,” on page 44](#)
- ♦ [Section 3.6, “Planning the eDirectory Deployment,” on page 44](#)
- ♦ [Section 3.7, “Planning for Shared Storage as Cluster Resources,” on page 45](#)
- ♦ [Section 3.8, “Planning for OES Services as Cluster Resources,” on page 46](#)

3.1 Determining Your Design Criteria

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

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

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

3.2 Using Cluster Best Practices

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

- ♦ Ensure that eDirectory is stable before implementing a cluster.

- ♦ Ensure that you have full Read/Write replicas of the entire eDirectory tree co-located in the data center where you are setting up the cluster.
- ♦ Ensure that IP addresses are unique.
- ♦ Consistently apply IP address assignments for each cluster and its cluster resources.
- ♦ Make IP address changes for the cluster and cluster resources only by using the procedure described in [Section 8.11.2, “Moving a Cluster or Changing IP Addresses of Cluster Nodes and Resources,” on page 136](#).

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

- ♦ Ensure that Volume IDs used for a cluster resources are unique across all nodes.
Each cluster node automatically assigns volume ID 0 to volume `sys` and volume ID 1 to volume `_ADMIN`. Cluster-enabled volumes use high volume IDs, starting from 254 in descending order. Volume IDs can be assigned in the cluster load script. You can view the volume IDs assigned on a node by using the `ncpcon volumes` command.
The Novell Client uses the volume ID to access a volume.
- ♦ Consider each node’s configuration requirements for each of the services it is intended to host.
- ♦ Create failover matrixes for each cluster resource so that you know what service is supported and which nodes are the preferred nodes for failover.

3.3 Planning the LAN Connectivity

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

- ♦ [Section 3.3.1, “VLAN,” on page 42](#)
- ♦ [Section 3.3.2, “Channel Bonding,” on page 43](#)
- ♦ [Section 3.3.3, “Spanning Tree Protocol,” on page 43](#)
- ♦ [Section 3.3.4, “IP Addresses,” on page 43](#)
- ♦ [Section 3.3.5, “Name Resolution,” on page 43](#)

3.3.1 VLAN

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

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

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

3.3.2 Channel Bonding

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

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

3.3.3 Spanning Tree Protocol

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

The default settings for STP inhibit the heartbeat for over 30 seconds whenever there is a change in link status. Test your STP configuration with 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 58](#).

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 70](#)
 - ♦ [“Managing Multipath I/O for Devices” \(http://www.suse.com/documentation/sles11/stor_admin/data/multipathing.html\)](#) 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 69](#).

3.5 Planning the Shared Storage Solution

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

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

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

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

3.6 Planning the eDirectory Deployment

Your 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 44](#)
- ♦ [Section 3.6.2, “Cluster OU Context,” on page 45](#)
- ♦ [Section 3.6.3, “Cluster OU Partitioning and Replication,” on page 45](#)

3.6.1 Object Location

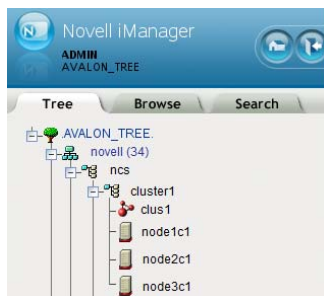
Cluster nodes and Cluster objects can exist in any container in the eDirectory tree. The Virtual Server object for the cluster and the objects for cluster resources are automatically created in the eDirectory context of the server where the cluster resource is created and cluster-enabled.

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 134.

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 12, "Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes," on page 217
Linux LVM volume groups and logical volumes	Chapter 13, "Configuring and Managing Cluster Resources for Shared LVM Volume Groups," on page 305
Linux POSIX volumes	Chapter 14, "Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers," on page 371
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 SP1: NCP Server for Linux Administration Guide</i>
Dynamic Storage Technology (DST) volumes (NSS volumes configured in a shadow volume pair)	"Configuring DST Shadow Volume Pairs with Novell Cluster Services" in the <i>OES 11 SP1: 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 HTTP Server" in the <i>OES 11 SP1: Novell Cluster Services 2.1 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 SP1: Novell AFP for Linux Administration Guide</i>
Archive and Version Services	"Configuring Archive and Version Service for Novell Cluster Services" in the <i>OES 11 SP1: Novell Archive and Version Services 2.1 Administration Guide</i>
Certificate Server (eDirectory Server Certificates)	"eDirectory Server Certificates" in the <i>OES 11 SP1: Novell Cluster Services 2.1 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 SP1: 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 SP1: 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 SP1: Novell DNS/DHCP Services for Linux Administration Guide</i>
DNS Server	"Configuring DNS with Novell Cluster Services" in the <i>OES 11 SP1: 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 3.9.1 Administration Guide</i>
iPrint	"Configuring iPrint with Novell Cluster Services" in the <i>OES 11 SP1: 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 SP1: Novell Cluster Services 2.1 NetWare to Linux Conversion Guide</i>
NetStorage	"Configuring NetStorage with Novell Cluster Services" in the <i>OES 11 SP1: NetStorage Administration Guide for Linux</i> .

Service	For information, see
PureFTPd (Novell FTP)	“Cluster Enabling Pure-FTPd in an OES 11 SP1 Environment” in the OES 11 SP1: Planning and Implementation Guide
QuickFinder 5.0.x (Server Synchronization Feature)	“QuickFinder Server” in the OES 11 SP1: Novell Cluster Services 2.1 NetWare to Linux Conversion Guide
Samba (Novell Samba)	“Configuring Samba for LVM Volume Groups and Novell Cluster Services” in the OES 11 SP1: 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 SP1 servers.

IMPORTANT: For information about designing your cluster and cluster environment, see [Chapter 3, “Planning for a Cluster,”](#) on page 41.

- ♦ [Section 4.1, “Cluster Administration Requirements,”](#) on page 49
- ♦ [Section 4.2, “IP Address Requirements,”](#) on page 52
- ♦ [Section 4.3, “Volume ID Requirements,”](#) on page 52
- ♦ [Section 4.4, “Hardware Requirements,”](#) on page 53
- ♦ [Section 4.5, “Virtualization Environments,”](#) on page 53
- ♦ [Section 4.6, “Software Requirements for Cluster Services,”](#) on page 54
- ♦ [Section 4.7, “Software Requirements for Cluster Resources,”](#) on page 63
- ♦ [Section 4.8, “Shared Disk Configuration Requirements,”](#) on page 66
- ♦ [Section 4.9, “SAN Rules for LUN Masking,”](#) on page 69
- ♦ [Section 4.10, “Multipath I/O Configuration Requirements,”](#) on page 70

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 49
- ♦ [Section 4.1.2, “NCS Proxy User,”](#) on page 50
- ♦ [Section 4.1.3, “Cluster Administrator or Administrator-Equivalent User,”](#) on page 52

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 50
- ♦ [“Container Administrator”](#) on page 50

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 74](#).

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 76](#).

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

The specified NCS Proxy User for the node is automatically assigned as a member in the `<cluster_name>_MGT_GRP.<context>` group that resides in the Cluster object container, such as `clus1_MGT_GRP.ncs.novell`. The group accommodates the server-specific NCS Proxy Users that you assign when you configure each node for the cluster. Members are added to the group as you configure each node for a cluster. Each member of the group has the necessary rights for configuring the cluster and cluster resources and for exchanging information with eDirectory.

For example, [Figure 4-1](#) shows that an OES Common Proxy User has been assigned as the NCS Proxy User for each node in a four-node cluster. The nodes are named *avalon*, *camelot*, *excalibur*, and *cornwall*. The context is *novell*.

Figure 4-1 Members of the NCS Management Group for a Cluster



IMPORTANT: You can modify this default administrator user name or password for the user name assigned as the NCS Proxy User after the install by following the procedure in [Section 8.11, “Moving a Cluster, or Changing the Node IP Addresses, LDAP Servers, or Administrator Credentials for a Cluster,”](#) on page 134.

Consider the following caveats for the three proxy user options:

- ♦ “OES Common Proxy User” on page 51
- ♦ “LDAP Admin User” on page 52
- ♦ “Another Administrator User” on page 52

OES Common Proxy User

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

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

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

For information about enabling or disabling the OES Common Proxy User, see the [OES 11 SP1: Installation Guide](#). For caveats and troubleshooting information for the OES Common Proxy user, see the [OES 11 SP1: Planning and Implementation Guide](#).

LDAP Admin User

If you specify the LDAP Admin user as the NCS Proxy User, you typically continue using this identity while you set up the cluster and cluster resources. After the cluster configuration is completed, you create another user identity to use for NCS Proxy User, and grant that user sufficient administrator rights as specified in [“Cluster Administrator or Administrator-Equivalent User” on page 52](#).

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 `<cluster_name>_MGT_GRP.<context>` group.

4.1.3 Cluster Administrator or Administrator-Equivalent User

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

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

4.2 IP Address Requirements

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

4.3 Volume ID Requirements

A volume ID is a value assigned to represent the volume when it is mounted by NCP Server on an OES server. The 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 the same volume IDs. Duplicate volume IDs can prevent resources from going online if the resource is BCC-migrated to a different cluster. When you BCC-enable a volume,

you must manually edit its load script to ensure that its volume ID is unique across all nodes in every peer cluster. You can use the `ncpcon volumes` command on each node in every peer cluster to identify the volume IDs in use by all mounted volumes. Compare the results for each server to identify the clustered volumes that have duplicate volume IDs assigned. Modify the load scripts to manually assign unique volume IDs.

4.4 Hardware Requirements

The following hardware requirements for installing 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 to Ethernet NICs, Novell Cluster Services supports VLAN on NIC bonding in OES 11 SP1 with the latest patches applied. No modifications to scripts are required. You can use `ethx` or `vlanx` interfaces in any combination for nodes in a cluster.

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 SP1: Installation Guide*.

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

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

4.5 Virtualization Environments

Xen and KVM virtualization software is included with SUSE Linux Enterprise Server. Novell Cluster Services supports using Xen or KVM 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 15, “Configuring Novell Cluster Services in a Virtualization Environment,”](#) on page 391.

Novell Cluster Services is supported to run on a 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. These virtualization templates are general, and can be adapted to work for other virtualization environments, such as KVM and VMware. For information about setting up Xen and XenLive cluster resources for a virtualization host server, see [Section 15.2, “Virtual Machines as Cluster Resources,”](#) on page 392.

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 SP1,” on page 54](#)
- [Section 4.6.2, “Novell Cluster Services,” on page 54](#)
- [Section 4.6.3, “Novell eDirectory 8.8.x,” on page 54](#)
- [Section 4.6.4, “SLP,” on page 58](#)
- [Section 4.6.5, “Novell iManager,” on page 59](#)
- [Section 4.6.6, “Clusters Plug-in for iManager,” on page 60](#)
- [Section 4.6.7, “Storage-Related Plug-Ins for iManager,” on page 61](#)
- [Section 4.6.8, “SFCB and CIMOM,” on page 61](#)
- [Section 4.6.9, “CASA 1.7,” on page 62](#)
- [Section 4.6.10, “Web Browser,” on page 63](#)

4.6.1 Novell Open Enterprise Server 11 SP1

Novell Cluster Services 2.1 for Linux supports Novell Open Enterprise Server 11 SP1 running on 64-bit SUSE Linux Enterprise Server 11 Service Pack 2 (SP2). Novell Cluster Services is one of the OES Services patterns.

We recommend having uniform nodes in the cluster. The same release version 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 SP3	Chapter 7, “Upgrading Clusters from OES 2 SP3 to OES 11x,” on page 105
NetWare 6.5 SP8	OES 11 SP1: Novell Cluster Services 2.1 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 SP1 servers. It is one of the OES Services patterns on the OES 11 SP1 Add-On disk.

4.6.3 Novell eDirectory 8.8.x

Novell eDirectory 8.8.7 (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.

For information about Novell eDirectory, see [Novell eDirectory 8.8 SP7 Administration Guide](#).

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

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

- ♦ [“eDirectory Tree” on page 55](#)
- ♦ [“eDirectory Context” on page 55](#)
- ♦ [“Cluster Object Container” on page 56](#)
- ♦ [“Cluster Objects Stored in eDirectory” on page 57](#)
- ♦ [“LDAP Server List” on page 58](#)

eDirectory Tree

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

eDirectory Context

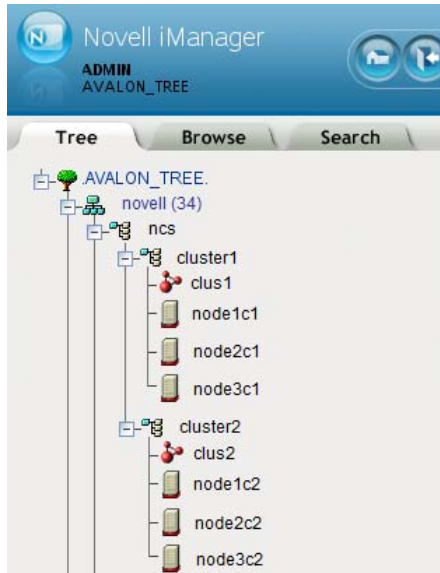
If you are creating a new cluster, the eDirectory context where the new Cluster object will reside must be an existing context. Specifying a new context during the Novell Cluster Services configuration does not create a new context.

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 context. Multiple Cluster objects can co-exist in the same eDirectory container. In iManager, use *Directory Administration > Create Object* to create a container for the cluster before you configure the cluster.

For example, [Figure 4-2](#) 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 4-2 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](#)” in the *Novell eDirectory 8.8 SP7 Administration Guide* for more information.

Renaming a pool involves changing information in the Pool object in eDirectory. If Server objects for the cluster nodes are in different containers, you must ensure that the shared pool is active on a cluster node that has its NCP server object in the same context as the Pool object of the pool you are going to rename. For information about renaming a shared pool, see [Section 12.12, “Renaming a Clustered NSS Pool,”](#) on page 252.

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 4-3 shows the 13 default eDirectory objects that are created in the Cluster container as viewed from the Tree view in iManager:

Figure 4-3 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 Novell 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 134.

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 SP1 installation. You can enable and configure SLP on the eDirectory Configuration - NTP & SLP page. For information, see [“Specifying SLP Configuration Options”](#) in the *OES 11 SP1: Installation Guide*.

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

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

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

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

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

```
cluster offline <resource_name>
```

or

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

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

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

```
rcncdsd restart
```

```
rcslpd restart
```

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

```
cluster online <resource_name>
```

or

```
cluster migrate <resource_name> <node_name>
```

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

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

For more information about configuring and managing SLP, see “[Configuring OpenSLP for eDirectory](#)” in the *Novell eDirectory 8.8 SP7 Administration Guide*.

4.6.5 Novell iManager

Novell iManager 2.7.5 or later is required for configuring and managing clusters on OES 11 SP1. 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 2.7.6 documentation Web site \(http://www.netiq.com/documentation/imanager27/\)](http://www.netiq.com/documentation/imanager27/).

For SFCB (Small Footprint CIM Broker) and CIMOM requirements, see [Section 4.6.8, “SFCB and CIMOM,” on page 61](#).

For browser configuration requirements, see “[Web Browser](#)” on page 63.

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 61](#).

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 [Novell Business Continuity Clustering product page \(http://www.novell.com/products/businesscontinuity/\)](http://www.novell.com/products/businesscontinuity/).

BCC Release	OES Support	iManager and Clusters Plug-In
BCC 2.0	OES 11 SP1	Novell iManager 2.7.6 Requires the Clusters plug-in for OES 11 SP1 with the latest patches applied. See the BCC 2.0: Administration Guide for OES 11 SP1 (http://www.novell.com/documentation/bcc/bcc20_admin_lx/data/bookinfo.html) .
BCC 1.2.2	OES 2 SP3	Novell 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 the BCC 1.2.2: Administration Guide for OES 2 SP3 (http://www.novell.com/documentation/bcc/bcc122_admin_lx/data/bookinfo.html) .
BCC 1.2.1	OES 2 SP2	Novell iManager 2.7.3 or later Requires the Clusters plug-in in the OES 2 SP2 January 2010 Maintenance patch. See the BCC 1.2.1: Administration Guide for OES 2 SP2 (http://www.novell.com/documentation/bcc/bcc121_admin_lx/data/bookinfo.html) .
BCC 1.1 SP2	NetWare 6.5 SP8	Novell 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 the 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 and later, the following storage-related plug-ins for Novell 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 2.7.6 Administration Guide*.

For information about working with storage-related plug-ins for iManager, see “[Understanding Storage-Related Plug-Ins](#)” in the *OES 11 SP1: NSS File System Administration Guide for Linux*.

4.6.8 SFCB and CIMOM

The Small Footprint CIM Broker (SFCB) replaces OpenWBEM for CIMOM activities in OES 11 and later. SFCB provides the default CIMOM and CIM clients. When you install any OES components that depend on WBEM, SFCB and all of its corresponding packages are installed with the components. For information, see “[Small Footprint CIM Broker \(SFCB\)](#)” in the *OES 11 SP1: Planning and Implementation Guide*.

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

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

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


The Clusters plug-in (and all other storage-related plug-ins) for iManager require CIMOM connections for tasks that transmit sensitive information (such as a user name and password) between iManager and the `_admin` volume on the OES 11 SP1 server that you are managing. Typically, CIMOM is running, so this should be the normal condition when using the server. CIMOM connections use Secure HTTP (HTTPS) for transferring data, and this ensures that sensitive data is not exposed.

IMPORTANT: SFCB is automatically PAM-enabled for Linux User Management (LUM) as part of the OES installation. Users not enabled for LUM cannot use the CIM providers to manage OES. The user name that you use to log in to 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 49](#).

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 `rcsfcb` command to help resolve CIMOM and SFCB issues:

To perform this task	At a command prompt, enter as the root user
To start SFCB	<code>rcsfcb start</code>
To stop SFCB	<code>rcsfcb stop</code>
To check SFCB status	<code>rcsfcb status</code>
To restart SFCB	<code>rcsfcb 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 1.7

Novell Cluster Services requires CASA (Common Authentication Service Adapter) 1.7 or later 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 command 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 2.7.6 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 63](#)
- ♦ [Section 4.7.2, “Novell Storage Services File System for Linux,” on page 64](#)
- ♦ [Section 4.7.3, “LVM Volume Groups and Linux POSIX File Systems,” on page 65](#)
- ♦ [Section 4.7.4, “NCP Volumes on Linux POSIX File Systems,” on page 65](#)
- ♦ [Section 4.7.5, “Dynamic Storage Technology Shadow Volume Pairs,” on page 65](#)
- ♦ [Section 4.7.6, “NCP File Access,” on page 65](#)
- ♦ [Section 4.7.7, “Novell AFP,” on page 65](#)
- ♦ [Section 4.7.8, “Novell CIFS,” on page 66](#)
- ♦ [Section 4.7.9, “Novell Samba,” on page 66](#)
- ♦ [Section 4.7.10, “Novell Domain Services for Windows,” on page 66](#)

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 SP1: 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 SP1: 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 SP1: 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 67.
- ♦ Creating and cluster-enabling a shared pool. For information, see [Chapter 12, “Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes,”](#) on page 217.
- ♦ Creating and cluster-enabling a shared Linux Logical Volume Manager (LVM) volume group. For information, see [Chapter 13, “Configuring and Managing Cluster Resources for Shared LVM Volume Groups,”](#) on page 305.

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 63.
- ♦ Novell Apple Filing Protocol (AFP). For information, see “[Novell AFP](#)” on page 65.
- ♦ Novell CIFS. For information, see “[Novell CIFS](#)” on page 66.

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 installation.

After the cluster is configured, you can create LVM volume group cluster resources as described in [Chapter 13, “Configuring and Managing Cluster Resources for Shared LVM Volume Groups,”](#) on page 305.

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 63.

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 63.

For information about creating and cluster-enabling NCP volumes, see [“Configuring NCP Volumes with Novell Cluster Services”](#) in the *OES 11 SP1: 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 and later. DST is installed automatically when you install NCP Server for Linux. To use cluster-enabled DST volume pairs, select the *NCP Server and Dynamic Storage Technology* option during the install.

For information about creating and cluster-enabling Dynamic Storage Technology volumes on Linux, see [“Configuring DST Shadow Volume Pairs with Novell Cluster Services”](#) in the *OES 11 SP1: 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 63.

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 [OES 11 SP1: Novell AFP for Linux 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* field and the *CIFS* check box to be available. Select the check box to enable Novell CIFS as an advertising protocol for the NSS pool cluster resource. A default *CIFS Virtual Server Name* is suggested, but you can modify it. The CIFS daemon should also be running before you bring resources online that have CIFS enabled.

To install 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 SP1: 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 SMB/CIFS 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 SP1: 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 SP1: 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 67](#)
- ♦ [Section 4.8.2, “SBD Partitions,” on page 67](#)

- ♦ [Section 4.8.3, “Shared iSCSI Devices,” on page 69](#)
- ♦ [Section 4.8.4, “Shared RAID Devices,” on page 69](#)

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 SP1 documentation or the shared disk system documentation for troubleshooting information.

Prepare the device for use in a cluster resource:

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

All devices that contribute space to a clustered pool must be able to fail over with the pool cluster resource. You must use the device exclusively for the clustered pool; do not use space on it for other pools or for Linux volumes. A device must be marked as Shareable for Clustering before you can use it to create or expand a clustered pool.

- ♦ **Linux LVM volume group:** For new devices, use an unpartitioned device that has been initialized. Do not mark the device as shared because doing so creates a small partition on it. LVM uses the entire device for the volume group. For an existing volume group, 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 67](#).

4.8.2 SBD Partitions

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

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

For information about how SBD partitions work and how to create an SBD partition for an existing cluster, see [Section 9.16, “Creating or Deleting Cluster SBD Partitions,” on page 157](#).

- ♦ [“Preparing the SAN Devices for the SBD” on page 68](#)
- ♦ [“Initializing and Sharing a Device for the SBD” on page 68](#)
- ♦ [“Determining the SBD Partition Size” on page 68](#)

Preparing the SAN Devices for the SBD

Use the SAN storage array software to carve a LUN to use exclusively for the SBD partition. The device should have at least 20 MB of free available space. The minimum size is 8 MB. Connect the LUN device to all nodes in the cluster.

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

The device you use to create the SBD partition must not be a software RAID device. A hardware RAID configured in a SAN array is seen as a regular device by the server.

Initializing and Sharing a Device for the SBD

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

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

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

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

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

Determining the SBD Partition Size

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

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

4.8.3 Shared iSCSI Devices

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

4.8.4 Shared RAID Devices

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

Consider the following when using software RAIDs:

- ♦ NSS software RAID is supported for shared disks for NSS pools. Any RAID0/5 device that is used for a clustered pool must contribute space exclusively to that pool; it cannot be used for other pools. This allows the device to fail over between nodes with the pool cluster resource. Ensure that its component devices are marked as Shareable for Clustering before you use a RAID0/5 device to create or expand a clustered pool
- ♦ Linux software RAID can be used in shared disk configurations that do not require the RAID to be concurrently active on multiple nodes. Linux software RAID cannot be used underneath clustered file systems (such as OCFS2, GFS, and CXFS) because 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 70](#)
- [Section 4.10.2, “Modifying the Port Down Retry Setting in the modprobe.conf.local File,” on page 70](#)
- [Section 4.10.3, “Modifying the Polling Interval, No Path Retry, and Failback Settings in the multipath.conf File,” on page 71](#)
- [Section 4.10.4, “Modifying the Port Down Retry and Link Down Retry Settings for an HBA BIOS,” on page 72](#)

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 DM-MP works correctly with Novell Cluster Services. For information, see [Section 4.10.2, “Modifying the Port Down Retry Setting in the modprobe.conf.local File,” on page 70](#) and [Section 4.10.3, “Modifying the Polling Interval, No Path Retry, and Failback Settings in the multipath.conf File,” on page 71](#).

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 72](#).

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. In addition, you want the multipath I/O fail back to be set to manual in order to avoid a ping-pong of resources because of path failures.

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

- ♦ [“Polling Interval” on page 71](#)
- ♦ [“No Path Retry” on page 71](#)
- ♦ [“Failback” on page 71](#)
- ♦ [“Example of Multipath I/O Settings” on page 72](#)

Polling Interval

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

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

No Path Retry

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

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

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

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

Failback

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

```
failback "manual"
```

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

Example of Multipath I/O Settings

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

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

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

For information about configuring the `multipath.conf` file, see “[Managing Multipath I/O for Devices](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 SP1

This section describes how to install the Novell Cluster Services software on Novell Open Enterprise Server 11 SP1 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 49.

See the following resources for information about rolling cluster upgrades or conversions to the latest version of Novell Cluster Services for OES 11 SP1:

Upgrading from	For information, see:
OES 11	Chapter 6, “Upgrading OES 11 Clusters,” on page 101
OES 2 SP3	Chapter 7, “Upgrading Clusters from OES 2 SP3 to OES 11x,” on page 105
NetWare 6.5 SP8	OES 11 SP1: Novell Cluster Services 2.1 NetWare to Linux Conversion Guide

- [Section 5.1, “Novell Cluster Services Licensing,”](#) on page 73
- [Section 5.2, “Extending the eDirectory Schema to Add Cluster Objects,”](#) on page 74
- [Section 5.3, “Assigning Install Rights for Container Administrators or Non-Administrator Users,”](#) on page 76
- [Section 5.4, “Installing Novell Cluster Services,”](#) on page 77
- [Section 5.5, “Configuring Novell Cluster Services,”](#) on page 80
- [Section 5.6, “Adding a Node That Was Previously in the Cluster,”](#) on page 95
- [Section 5.7, “Configuring Additional Administrators,”](#) on page 96
- [Section 5.8, “Installing or Updating the Clusters Plug-in for iManager,”](#) on page 96
- [Section 5.9, “Patching Novell Cluster Services,”](#) on page 98
- [Section 5.10, “What’s Next,”](#) on page 99

5.1 Novell Cluster Services Licensing

Novell Cluster Services supports up to 32 nodes in a single cluster. OES 11 SP1 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 Cluster object container and the following types of objects in it:

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

A tree administrator user with the eDirectory credentials to do so can extend the eDirectory schema before a cluster is installed anywhere in a tree. This allows container administrators (or non-administrator users) to install a cluster in a container in that same tree without needing full administrator rights for the tree. After the schema has been extended, you must assign some eDirectory rights to the container administrators (or non-administrator users) who will install 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 76.

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 74
- ♦ [Section 5.2.2, “Extending the Schema,”](#) on page 75

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

Parameter	Description	Example
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 command prompt:

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

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

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

Replace *configuration_filename* with the actual name of the file that you created in [Step 2](#).

- 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 76.

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 85.

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

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

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

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

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

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

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

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

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

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

For information about eDirectory rights, see “[eDirectory Rights](#)” in the [Novell eDirectory 8.8 SP7 Administration Guide](#).

5.4 Installing Novell Cluster Services

Novell Cluster Services for Linux is included on the add-on media for OES 11 SP1. 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 73](#).

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 SP1, or you can install it later on an existing OES 11 SP1 server.

- ♦ [Section 5.4.1, “Before You Install Novell Cluster Services,” on page 77](#)
- ♦ [Section 5.4.2, “Installing Novell Cluster Services during an OES 11 SP1 Installation,” on page 78](#)
- ♦ [Section 5.4.3, “Installing Novell Cluster Services on an Existing OES 11 SP1 Server,” on page 79](#)

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 49](#). Before you install Novell Cluster Services, verify that your setup meets these requirements:

- ♦ [“Using a Local eDirectory Database” on page 77](#)
- ♦ [“Rights to Extend the Schema” on page 77](#)
- ♦ [“Shared Devices for the SBD” on page 77](#)
- ♦ [“SBD Required before You Install the Second Node in the Cluster” on page 78](#)

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 74](#), the user who installs Novell Cluster Services must be an administrator user with the rights to extend the schema, such as the tree administrator.

Shared Devices for the SBD

If you have a shared disk system and the server where you are installing 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 86](#). A one-node cluster that has shared disk storage can be configured without an SBD, but the SBD must be created before you add another node. For information about manually creating an SBD, see:

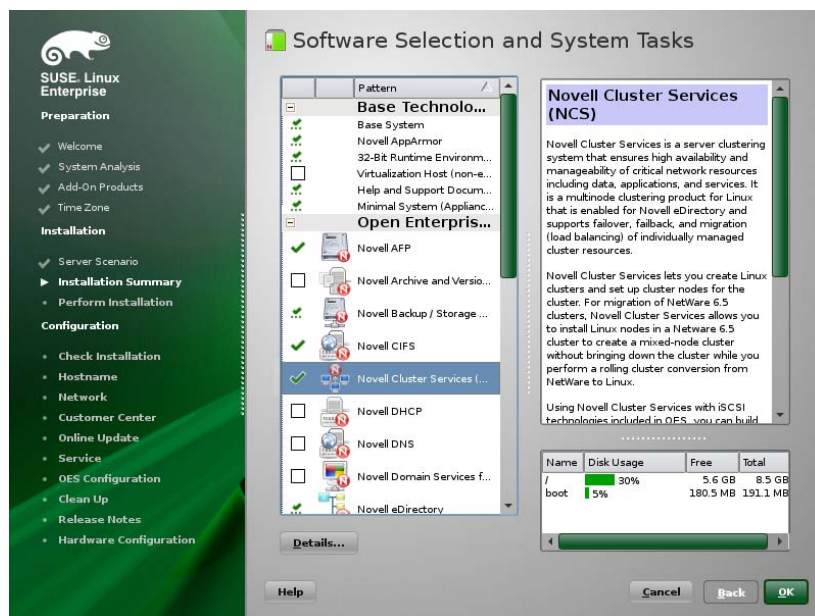
- ♦ [Section 9.16.3, “Creating a Non-Mirrored Cluster SBD Partition with SBDUTIL,” on page 160](#)
- ♦ [Section 9.16.5, “Creating a Mirrored Cluster SBD Partition with SBDUTIL,” on page 167](#)

5.4.2 Installing Novell Cluster Services during an OES 11 SP1 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 SP1, see the [OES 11 SP1: 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 63.](#)

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.

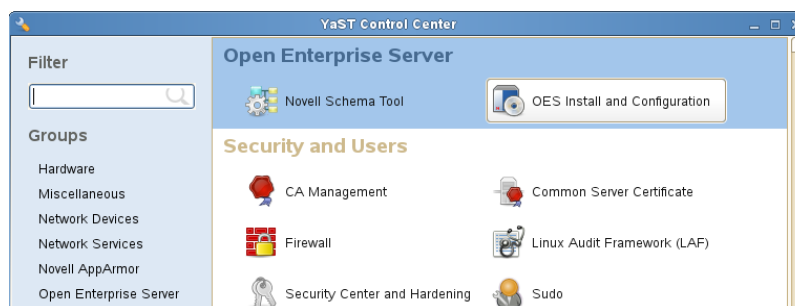
For information, see “[Updating \(Patching\) an OES 11 SP1 Server](#)” in the [OES 11 SP1: Installation Guide](#).

- 7 Ensure that your SAN storage is working properly for the server.
- 8 Continue with [Section 5.5, “Configuring Novell Cluster Services,” on page 80.](#)

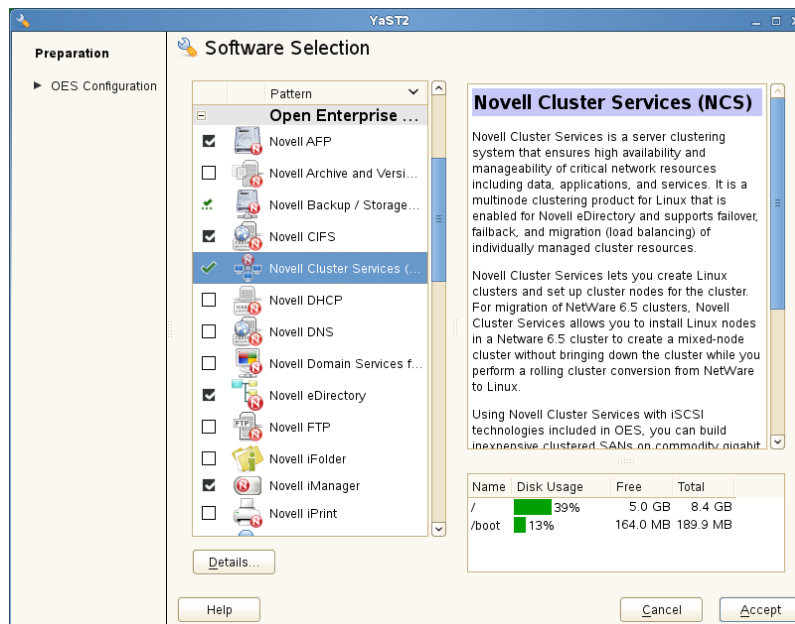
5.4.3 Installing Novell Cluster Services on an Existing OES 11 SP1 Server

If you did not install Novell Cluster Services during the OES 11 SP1 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 an OES 11 SP1 Installation,” on page 78.](#)

- 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.

For information, see [“Updating \(Patching\) an OES 11 SP1 Server”](#) in the *OES 11 SP1: Installation Guide*.

- 7 Ensure that your SAN storage is working properly for the server.
- 8 Continue with [Section 5.5, “Configuring Novell Cluster Services,” on page 80.](#)

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 the Node IP Addresses, LDAP Servers, or Administrator Credentials for a Cluster,” on page 134.](#)

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:

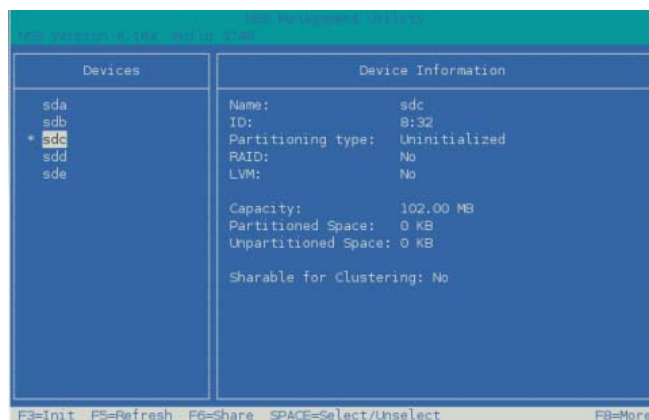
- ♦ [Section 5.5.1, “Initializing and Sharing a Device to Use for the SBD Partition,” on page 81](#)
- ♦ [Section 5.5.2, “Opening the Novell Open Enterprise Server Configuration Page,” on page 84](#)
- ♦ [Section 5.5.3, “Using Different LDAP Credentials for the Cluster Configuration,” on page 85](#)
- ♦ [Section 5.5.4, “Accessing the Novell Cluster Services Configuration Page in YaST,” on page 85](#)
- ♦ [Section 5.5.5, “Configuring a New Cluster,” on page 86](#)
- ♦ [Section 5.5.6, “Adding a Node to an Existing Cluster,” on page 91](#)

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 command prompt, enter


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



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

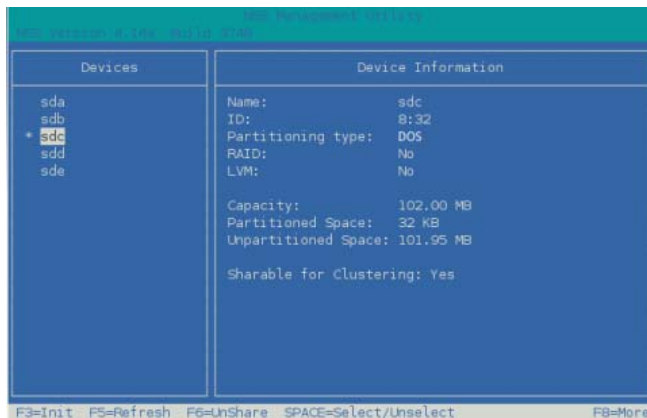
WARNING: Initializing a disk destroys all of the data on it.

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



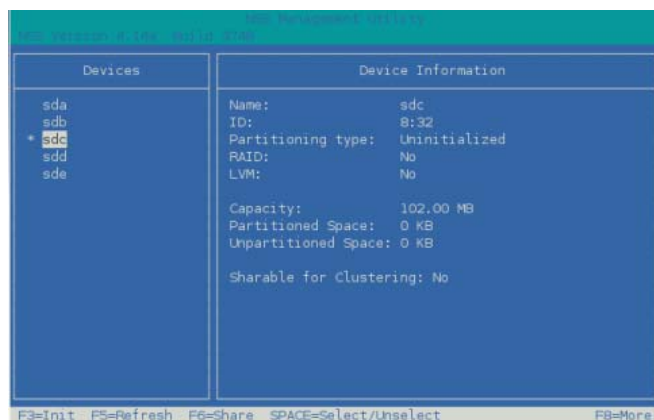
Wait for the page to refresh before continuing.

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



- 8 If you plan to mirror the SBD, repeat [Step 4](#) through [Step 7](#) for the second device.
- 9 Exit NSSMU.
- 10 Continue with [Section 5.5.2, “Opening the Novell Open Enterprise Server Configuration Page,” on page 84](#).
 - 1 Log in as the root user on the server that will be the first node of the cluster, then open a terminal console.
 - 2 At the command prompt, enter


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



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

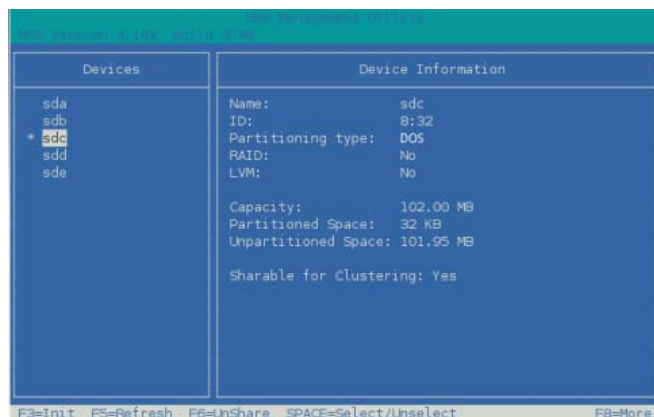
WARNING: Initializing a disk destroys all of the data on it.

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



Wait for the page to refresh before continuing.

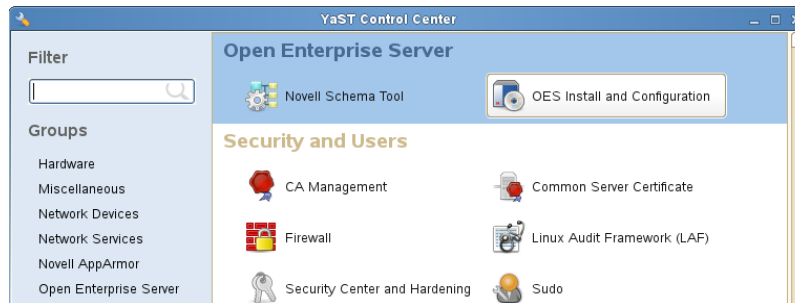
- 7 Press F6 to mark the device as shareable for clustering. The *Shareable for Clustering* value changes from *No* to *Yes*.



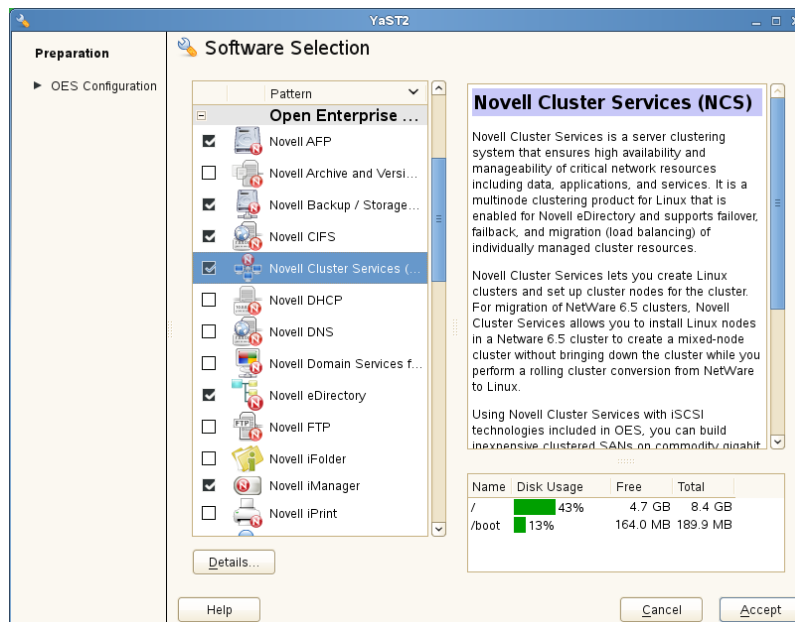
- 8 If you plan to mirror the SBD, repeat [Step 4](#) through [Step 7](#) for the second device.
- 9 Exit NSSMU.
- 10 Continue with [Section 5.5.2, “Opening the Novell Open Enterprise Server Configuration Page,”](#) on page 84.

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 85.](#)
 - ♦ **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 85.](#)

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 76.

- 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 85.

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*.



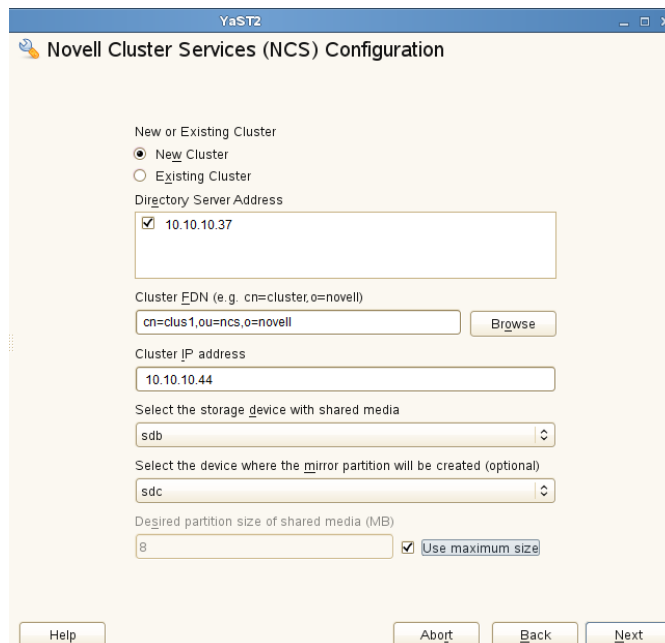
This is either the LDAP administrator identity that was used when the Open Enterprise Server were installed, or the identity configured after the install in [Section 5.5.3, “Using Different LDAP Credentials for the Cluster Configuration,”](#) on page 85.

- 5 On the Novell Cluster Services Configuration page, continue with one of the following:
 - ♦ [Section 5.5.5, “Configuring a New Cluster,”](#) on page 86
 - ♦ [Section 5.5.6, “Adding a Node to an Existing Cluster,”](#) on page 91

5.5.5 Configuring a New Cluster

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

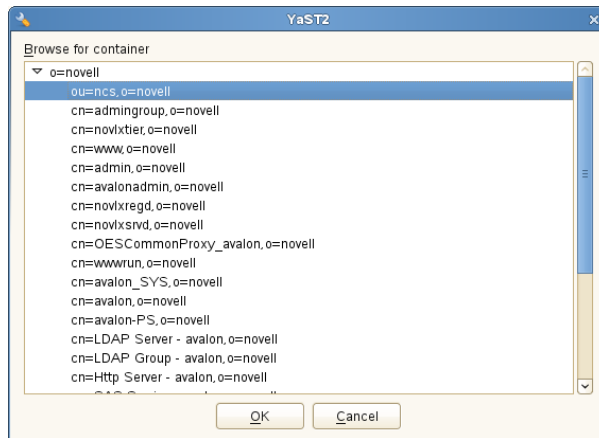
- 1 Go to the Novell Cluster Services Configuration page as described in [Section 5.5.4, “Accessing the Novell Cluster Services Configuration Page in YaST,”](#) on page 85.



- 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 134.</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>




You can also specify the typeful FDN (fully distinguished name) of the existing cluster. The name is case sensitive. Use the comma format illustrated in the example. Do not use dots.

For example:

```
cn=clus1,ou=ncs,o=mycompany
```

You must specify an existing context. Specifying a new context does not create a new context.

Cluster names must be unique. You cannot create two clusters with the same name in the same eDirectory tree.

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> <pre>10.10.10.44</pre> <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 initialized and 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 Section 4.8.2, "SBD Partitions," on page 67.</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 81.</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*.

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*.

YaST2

Novell Cluster Services (NCS) Configuration

Name of this node: avalon

IP address of this node


10.10.10.37

☒ Start Clustering Services now

Help Abort Back Finish

Parameter	Action
IP address of this node	<p>From the drop-down list, select the IP address that Novell Cluster Services will use for this node.</p> <p>Some servers have multiple IP addresses. This step lets you choose which IP address Novell Cluster Services uses.</p>
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>	<p>Start Novell Cluster Services manually by using one of these methods:</p> <ul style="list-style-type: none"> ◆ Reboot the cluster server. ◆ At a command 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 command prompt, enter the following as the <code>root</code> user: <pre>rcnovell-ncs start</pre>

- 8 Continue with [Section 5.7, “Configuring Additional Administrators,”](#) on page 96.

5.5.6 Adding a Node to an Existing Cluster

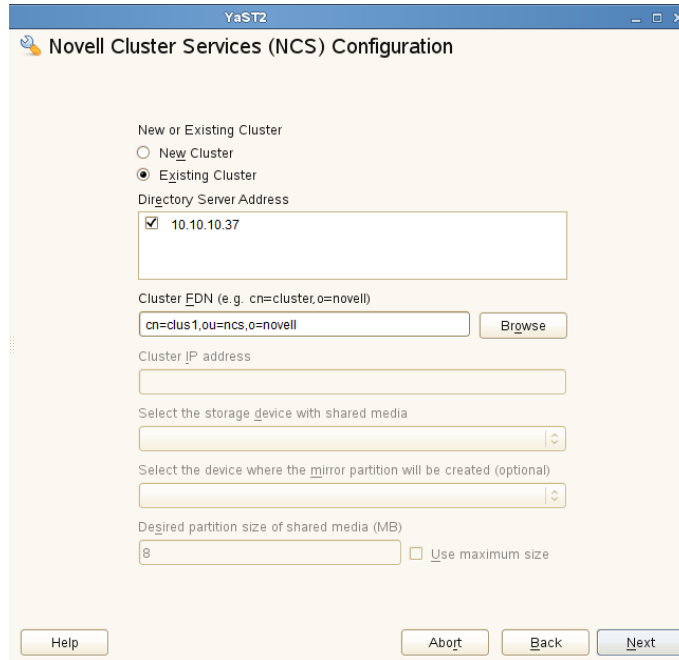
Perform the following configuration for each node that you add to an existing cluster:

- 1 If you have not previously configured the SBD partition for the cluster, create an SBD partition for the cluster by using one of the following procedures:
 - ◆ [Section 9.16.3, “Creating a Non-Mirrored Cluster SBD Partition with SBDUTIL,”](#) on page 160
 - ◆ [Section 9.16.5, “Creating a Mirrored Cluster SBD Partition with SBDUTIL,”](#) on page 167

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.

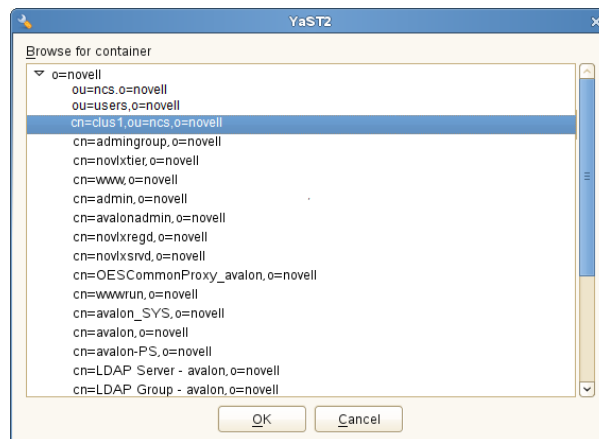
- 2 Go to the Novell Cluster Services Configuration page as described in [Section 5.5.4, “Accessing the Novell Cluster Services Configuration Page in YaST,”](#) on page 85.



- 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> .
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 134.</p>

Parameter	Action
Cluster FDN	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.



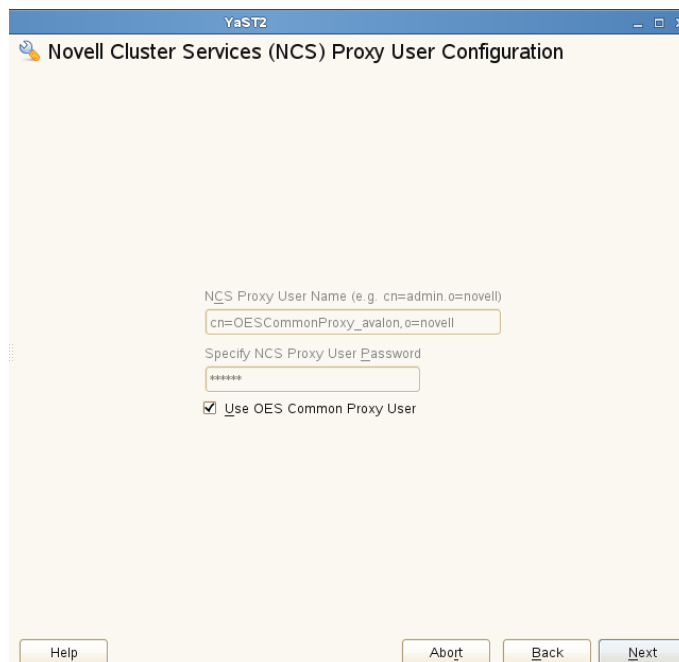
You can also specify the typeful FDN (fully distinguished name) of the existing cluster. The name is case sensitive. Use the comma format illustrated in the example. Do not use dots.

For example:

`cn=clus1,ou=ncs,o=mycompany`

Select the device where the mirror partition will be created	If you want to mirror the SBD, select a different shared device from the drop-down list.
--	--


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



Parameter	Action
OES Common Proxy User	If the OES Common Proxy User is enabled in eDirectory, the <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 by 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>	<p>Start Novell Cluster Services manually by using one of these methods:</p> <ul style="list-style-type: none"> ♦ Reboot the cluster server. ♦ At a command 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 command prompt, enter the following as the <code>root</code> user: <pre>rcnovell-ncs start</pre>

- 9 Continue with [Section 5.7, “Configuring Additional Administrators,”](#) on page 96.

5.6 Adding a Node That Was Previously in the Cluster

You can use the procedure in this section if you need to rebuild a node and want to use the same node name and IP address of the node it replaces.

- 1 If the Cluster object for the old server is still present, delete the object and force the cluster to recognize that the node has been permanently removed.
For information, see [Section 9.15, “Deleting a Cluster Node from a Cluster,”](#) on page 157.
- 2 In the SAN, assign the shared storage array and devices to the replacement server before you install Novell Cluster Services.
- 3 Install OES 11 SP1 and apply the latest patches. Ensure that you use the same node name and IP address of the node that was previously in the cluster.
- 4 Install Novell Cluster Services on the replacement server.
[Section 5.4.3, “Installing Novell Cluster Services on an Existing OES 11 SP1 Server,”](#) on page 79.
- 5 On the Novell Open Enterprise Server Configuration page under *Novell Cluster Services*, locate the Novell Cluster Services link, enable the configuration, then click the Novell Cluster Services link to open the Novell Cluster Services Configuration wizard.
For information see, [Section 5.5.4, “Accessing the Novell Cluster Services Configuration Page in YaST,”](#) on page 85.
- 6 Configure the replacement node for the existing cluster. The settings should be the same as those for the old server.
For information, see [Section 5.5.6, “Adding a Node to an Existing Cluster,”](#) on page 91.
After you restart Novell Cluster Services, the replacement server assumes its old identity in the cluster.

5.7 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 8.11.1, "Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster,"](#) on page 134.

5.8 Installing or Updating the Clusters Plug-in for iManager

In OES 11 SP1, the Clusters plug-in for Novell iManager 2.7.5 has been reorganized. On the Roles and Tasks page under *Clusters*, the *Cluster Manager*, *BCC Manager*, *Cluster Event Log*, and *Cluster Options* menu options have been replaced with two options: *My Clusters* and *My Resources*.

- ♦ **My Clusters:** The logged-in cluster administrator can set up a personalized list of clusters to manage, which allows their status to be viewed at-a-glance. The list persists between the administrator's logins to iManager in the same tree. For information, see [Section 8.2, "Setting Up a Personalized List of Clusters to Manage,"](#) on page 117.

My Resources: The logged-in cluster administrator can set up a personalized list of resources to manage, which allows their status to be viewed at-a-glance. The list persists between the administrator's logins to iManager in the same tree. For information, see [Section 10.2, "Setting Up a Personalized List of Resources to Manage,"](#) on page 179.

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

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

- ♦ [Section 5.8.1, "Prerequisites for Installing or Updating the Clusters Plug-In,"](#) on page 97
- ♦ [Section 5.8.2, "Installing the Clusters Plug-In,"](#) on page 97
- ♦ [Section 5.8.3, "Uninstalling and Reinstalling the Clusters Plug-In,"](#) on page 97
- ♦ [Section 5.8.4, "Updating Role-Based Services for the Clusters Plug-In for OES 11 SP1,"](#) on page 98

5.8.1 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 61.

5.8.2 Installing the Clusters Plug-In

If you are installing a new instance of Novell iManager 2.7.5 or later and the Clusters plug-in on a computer, follow the instructions in [“Downloading and Installing Plug-in Modules”](#) in the *Novell iManager 2.7.6 Administration Guide*. Ensure that you install the Clusters plug-in (`ncsmgmt.rpm`) as well as the common code for storage-related plug-ins (`storagemgmt.rpm`).

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

To configure role-based services for the Clusters plug-in, see [“RBS Configuration”](#) in the *Novell iManager 2.7.6 Administration Guide*.

5.8.3 Uninstalling and Reinstalling the Clusters Plug-In

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

- 1 In iManager, uninstall the currently installed storage-related plug-ins.
- 2 Copy the new version of the plug-in files to the iManager plug-ins location, manually overwriting the older version of the plug-in in the `packages` folder with the newer version of the plug-in.
- 3 In iManager, install all of the storage-related plug-ins, or install the plug-ins you need, plus the common code.
- 4 If you are updating the Clusters plug-in, delete the contents of the `/var/opt/novell/tomcat6/work/Catalina/localhost/nps` work folder.
You must manually delete the old compiled JSPs that are cached in the `nps` work folder. This ensures that only the latest files are cached.
- 5 Verify that the `/var/opt/novell/iManager/nps/WEB-INF/lib` directory contains the class loader file `commons-lang-2.6.jar` (or later version), then remove all earlier versions of the file from the directory.

The Clusters plug-in for Novell iManager 2.7.5 requires `commons-lang-2.6.jar` (or later version) of the class loader file. Novell iManager uses one class loader for all of its plug-ins. If multiple versions of the class loader file exist in the `/var/opt/novell/iManager/nps/WEB-`

INF/lib directory, iManager loads whichever one comes first alphabetically. You must remove the earlier versions (such as commons-lang.jar or commons-lang-2.4.jar) to ensure that iManager loads the 2.6 or later version of the class loader file.

- 6 Restart Tomcat by entering the following command, or reboot the server if a reboot is required by any of the patches applied.

```
rcnovell-tomcat6 restart
```

- 7 If you use Role-Based Services, continue with [Section 5.8.4, “Updating Role-Based Services for the Clusters Plug-In for OES 11 SP1,”](#) on page 98.

5.8.4 Updating Role-Based Services for the Clusters Plug-In for OES 11 SP1

In OES 11 SP1, the Clusters plug-in replaces the *Clusters*, the *Cluster Manager*, *BCC Manager*, *Cluster Event Log*, and *Cluster Options* menu options with the *My Clusters* and *My Resources* options. The plug-in is updated automatically to show new tasks, and hide old ones. However, the Role-Based Services role for Clusters is not updated automatically.

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

- 1 In iManager, click *Configuration*, then select *Role Based Services > RBS Configuration*.
- 2 Select the *iManager 2.x Collections* tab.
- 3 On the iManager 2.x Collections page, view the entry for *Role Based Services 2.novell*.
There are numbers in the columns for *Modules*, *Installed*, *Out-of-Date*, and *Not-Installed*.
- 4 Click the number link in the *Out-of-Date* column for *Role Based Services 2.novell*.
- 5 Select the check box next to *Clusters*, then click *Update*.
This updates the RBS configuration information for the Clusters plug-in.
- 6 Click OK to confirm the reinstallation.
- 7 After the reinstallation is complete, click OK to confirm the success message.
- 8 Log out of iManager.
- 9 Restart Tomcat by entering the following command, or reboot the server if a reboot is required by any of the patches applied.

```
rcnovell-tomcat6 restart
```

- 10 Log in to iManager to use the updated Clusters plug-in.

5.9 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 SP1 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.

5.10 What's Next

After installing Novell Cluster Services on OES 11 SP1 servers (physical servers or virtual guest servers), you can configure and manage the cluster and cluster resources. For information, see the following:

- ♦ [Chapter 8, “Configuring Cluster Policies, Protocols, and Properties,” on page 113](#)
- ♦ [Chapter 9, “Managing Clusters,” on page 139.](#)
- ♦ [Chapter 10, “Configuring and Managing Cluster Resources,” on page 177](#)
- ♦ [Chapter 12, “Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes,” on page 217](#)
- ♦ [Chapter 13, “Configuring and Managing Cluster Resources for Shared LVM Volume Groups,” on page 305](#)
- ♦ [Chapter 14, “Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers,” on page 371](#)

If you install Novell Cluster Services at the host level of an OES 11 SP1 virtualized server, you can create cluster resources for the virtual machines. For information, see [Section 15.2, “Virtual Machines as Cluster Resources,” on page 392.](#)

6 Upgrading OES 11 Clusters

This section describes how to upgrade a Novell Cluster Services cluster and the file system resources from Novell Open Enterprise Server (OES) 11 to OES 11 SP1. See the individual service guides for information about upgrading clustered services.

- ♦ [Section 6.1, “Requirements and Guidelines for Upgrading Clusters from OES 11 to OES 11 SP1,” on page 101](#)
- ♦ [Section 6.2, “Adding OES 11 SP1 Nodes to an OES 11 Cluster \(Rolling Cluster Upgrade\),” on page 102](#)
- ♦ [Section 6.3, “Upgrade Issues for OES 11 SP1,” on page 104](#)

6.1 Requirements and Guidelines for Upgrading Clusters from OES 11 to OES 11 SP1

In addition to the [Chapter 4, “Planning for Novell Cluster Services,” on page 49](#), consider the following rules and recommendations for mixed OES 11 and OES 11 SP1 clusters:

- ♦ Mixed-mode OES 11 and OES 11 SP1 clusters should be considered a temporary configuration that exists only during an upgrade.
- ♦ Adding a new OES 11 cluster node to a mixed-mode OES 11 and OES 11 SP1 cluster is not supported.
- ♦ NSS pool cluster resources can fail over to any node in the mixed-mode OES 11 and OES 11 SP1 cluster.
- ♦ LVM volume group cluster resources can fail over to any node in the mixed-mode OES 11 and OES 11 SP1 cluster.
- ♦ CSM cluster resources can fail over to any node in the mixed-mode OES 11 and OES 11 SP1 cluster.
- ♦ No storage management functions should be executed while a cluster is in mixed-mode unless you must do so as part of the documented upgrade process. Do not attempt to create, delete, expand, or modify the properties for partitions, pools, or volumes for any shared resources in the cluster.

6.2 Adding OES 11 SP1 Nodes to an OES 11 Cluster (Rolling Cluster Upgrade)

Performing a rolling cluster upgrade from OES 11 to OES 11 SP1 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 SP1 server is upgraded or 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 11 cluster nodes.
You will probably want to install the same OES components on any new OES 11 SP1 nodes that you add to the cluster.
- 2 If you are adding new OES 11 SP1 nodes, prepare the OES 11 SP1 nodes for the cluster, but do not allow them to join the cluster.
 - 2a Install SUSE Linux Enterprise Server (SLES) 11 SP2 and the OES Services from the OES 11 SP1 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 85.
 - 2d Configure the node for the existing OES 11 cluster as described in [Section 5.5.6, "Adding a Node to an Existing Cluster,"](#) on page 91. Complete the following fields:

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 85.</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 134.</p>
Cluster FDN	Click <i>Browse</i> , navigate the tree to select the existing Cluster object for the OES 11 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 For each of the new OES 11 SP1 nodes in turn, start Novell Cluster services manually by using one of the following methods, and allow the node to join the OES 11 cluster.
 - ♦ Reboot the cluster server.
 - ♦ At a command prompt, go to the `/etc/init.d` directory, then enter the following as the root user:

```
./novell-ncs start
```
 - ♦ At a command prompt, enter the following as the root user:

```
rcnovell-ncs start
```
- 4 For each new OES 11 SP1 node, ensure that you modify the *Preferred Nodes* list accordingly for each of the cluster resources.
- 5 For each of the existing OES 11 cluster nodes in turn:
 - 5a Log in to the node as the root user, then open a terminal console.
 - 5b Use the `cluster leave` command to remove the node from the cluster.

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

You can also manually cluster migrate the resources to another server in the cluster prior to bringing down the server.
 - 5c Stop Novell Cluster Services by entering

```
rcnovell-ncs stop
```
 - 5d If the autostart of Novell Cluster Services is enabled, you must disable autostart in order to prevent it from automatically loading Novell Cluster Services during the upgrade. Enter

```
chkconfig -s novell-ncs off
```
 - 5e Apply the next SP release as an upgrade, and fully patch the system.

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

```
chkconfig -s novell-ncs on
```

5g (Optional) Enable autostart of Novell Cluster Services for the node, enter

```
chkconfig -s novell-ncs on
```

5h Start Novell Cluster Services by entering

```
rcnovell-ncs start
```

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

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

5k Repeat these steps for each node running OES 11.

6 After the last cluster resource has been moved to the OES 11 SP1 nodes and all of the OES 11 nodes have been upgraded to OES 11 SP1, the conversion is complete.

6.3 Upgrade Issues for OES 11 SP1

This section contains known issues for upgrading clusters from OES 11 Linux to OES 11 SP1.

- ♦ [Section 6.3.1, “Updating the Clusters Plug-in and Role-Based Services in Novell iManager,” on page 104](#)
- ♦ [Section 6.3.2, “DHCP Cluster Resource: Path Change for dhcpd.pid,” on page 104](#)

6.3.1 Updating the Clusters Plug-in and Role-Based Services in Novell iManager

In OES 11 SP1, the Clusters plug-in replaces the *Cluster Manager*, *BCC Manager*, *Cluster Event Log*, and *Cluster Options* menu options with the *My Clusters* and *My Resources* options. The former menu options appear as tabs when you select a cluster from the *My Clusters* page. The plug-in is updated automatically to show new tasks, and hide old ones. However, the Role-Based Services role for Clusters is not updated automatically.

If you use Role-Based Services, the Clusters plug-in must be updated manually on the Role-Based Services Configuration page in order to pick up the modified menu options. For information, see [Section 5.8, “Installing or Updating the Clusters Plug-in for iManager,” on page 96](#).

6.3.2 DHCP Cluster Resource: Path Change for dhcpd.pid

The DHCP process ID file location was changed from `/var/run/dhcpd.pid` to `/var/lib/dhcp/var/run/dhcpd.pid`. During a cluster upgrade from OES 11 to OES 11 SP1, you must modify the location of the `dhcpd.pid` file in the DHCP cluster resource's unload script. After you modify the script, you should bring the resource online only on OES 11 SP1 nodes. This is necessary to ensure a graceful shutdown of the `dhcpd.pid` process when the DHCP resource fails over to a different node. For information, see [“Path Change for dhcpd.pid”](#) and [“Changing the Path for dhcpd.pid”](#) in the *OES 11 SP1: Novell DNS/DHCP Services for Linux Administration Guide*.

7 Upgrading Clusters from OES 2 SP3 to OES 11x

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 or to OES 11 SP1. See the individual service guides for information about converting cluster resources.

IMPORTANT: You can upgrade a cluster from OES 2 SP3 to OES 11 or from OES 2 SP3 to OES 11 SP1. The same release version of OES 11 must be installed and running on each new or upgraded node in the cluster.

The procedures in this section refer to OES 11x. A specific version of OES 11 is mentioned only if the procedure applies only to that version of the platform.

- ♦ [Section 7.1, “What’s New and Changed for Clustered Storage,” on page 105](#)
- ♦ [Section 7.2, “Supported Upgrade Paths from OES 2 SP3 to OES 11x,” on page 107](#)
- ♦ [Section 7.3, “Requirements and Guidelines for Upgrading Clusters from OES 2 SP3,” on page 108](#)
- ♦ [Section 7.4, “Adding OES 11x Nodes to an OES 2 SP3 Cluster \(Rolling Cluster Upgrade\),” on page 108](#)
- ♦ [Section 7.5, “Updating the Clusters Plug-in for Novell iManager,” on page 112](#)
- ♦ [Section 7.6, “Upgrading a Cluster with DST Resources from OES 2 SP3 to OES 11x,” on page 112](#)

7.1 What’s New and Changed for Clustered Storage

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

- ♦ [Section 7.1.1, “NSS Pools,” on page 105](#)
- ♦ [Section 7.1.2, “Linux POSIX File Systems on an EVMS CSM \(Compatibility Only\),” on page 106](#)
- ♦ [Section 7.1.3, “Linux POSIX File Systems on an LVM Volume Group,” on page 107](#)

7.1.1 NSS Pools

Novell Cluster Services supports NSS pools. In a mixed-mode cluster, existing NSS pool cluster resources can be cluster migrated to any node in the cluster.

NSS pools are managed by different storage managers on each platform, as shown in [Table 7-1](#). In a mixed-mode cluster with older operating systems, the OES 11 (or later) nodes are not able to recognize storage management actions performed on a pool on the alternate management system.

Table 7-1 Storage Managers for NSS Pools

Platform	Storage Manager
OES 11 and later	Novell Linux Volume Manager (NLVM)
OES 2 SP3	Enterprise Volume Management System (EVMS)
NetWare 6.5 SP8	Media Manager

In NetWare 6.5 SP8 and OES 2 SP3, NSS management tools recognize only DOS devices that are up to 2 TB in size. In OES 11 and later, 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 of any size, which allows you to create NSS pools up to 8 TB (terabytes) in size on a single device. In a mixed-mode cluster, the older platforms do not recognize devices that are partitioned with the GPT partitioning scheme.

WARNING: In a mixed-mode cluster, do not perform any storage management actions (create, modify, or delete) for NSS pools while OES 11 (or later) nodes are active in the cluster. Any such action can result in data corruption.

You must not perform storage management actions on pools while the cluster is in mixed mode. Wait to create new NSS pools until after all of the nodes in the cluster are upgraded. 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 11x nodes, create the NSS pool cluster resource on the older nodes, then restart the upgraded nodes. This allows the upgraded nodes to see the new configuration. This is critical so that pools are able to cluster migrate across platforms, and to prevent possible corruption.

For information about creating and managing NSS pool cluster resources on OES 11 or later, see [Chapter 12, “Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes,”](#) on page 217.

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

Novell Cluster Services for OES 11 and later 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 SP3, 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 SP3 to OES 11 and later:

- ◆ 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 OES 11x nodes join an OES2 SP3 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 11x nodes to join the OES 2 SP3 cluster, then modify each resources' preferred nodes list to use only the OES 11x 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 only on OES 11x nodes. As you add other OES 11x nodes, you can add those to the preferred nodes lists for failover.

After a resource has been moved from an OES 2 SP3 node to an OES 11x node, it is not supported to move it back to an OES 2 SP3 node. The resource should fail over only to other OES 11x 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 11x.

For information about managing CSM cluster resources on OES 11 and later servers, see [Chapter 14, “Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers,”](#) on page 371.

7.1.3 Linux POSIX File Systems on an LVM Volume Group

Beginning in OES 11, Novell Cluster Services supports Linux POSIX file systems on Linux Volume Manager (LVM) volume groups, including Btrfs, Ext2, Ext3, ReiserFS, and XFS. For information about the file systems, see the [“Overview of Linux POSIX File Systems”](#) in the *OES 11 SP1: Linux POSIX Volume Administration Guide*.

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

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

Beginning in OES 11 SP1, NSSMU and Novell Linux Volume Manager (NLVM) commands can be used to create cluster enabled LVM volume group cluster resources. For information, see [Chapter 13, “Configuring and Managing Cluster Resources for Shared LVM Volume Groups,”](#) on page 305.

Linux LVM management tools are also available in the YaST Expert Partitioner and in native Linux LVM command line commands. For information about using Linux commands to create a shared LVM volume group and cluster enabling it for Novell Cluster Services, see [Section 13.4.1, “Creating a Shared LVM Volume with LVM Commands,”](#) on page 330.

7.2 Supported Upgrade Paths from OES 2 SP3 to OES 11x

The following upgrade paths are supported from OES 2 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
OES 2 SP3 on SUSE Linux Enterprise Server 10 SP3	OES 11 SP1 on SUSE Linux Enterprise Server 11 SP2
OES 2 SP3 on SUSE Linux Enterprise Server 10 SP4	OES 11 SP1 on SUSE Linux Enterprise Server 11 SP2

7.3 Requirements and Guidelines for Upgrading Clusters from OES 2 SP3

In addition to the [Chapter 4, “Planning for Novell Cluster Services,” on page 49](#), consider the following rules and recommendations when upgrading from OES 2 SP3:

- ♦ Mixed-mode clusters should be considered a temporary configuration that exists only during an upgrade.
- ♦ Adding a new OES 2 SP3 cluster node to a mixed-mode cluster is not supported.
- ♦ Linux POSIX cluster resources on OES 2 SP3 nodes must be taken offline before you allow any of the OES 11x 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 11x nodes. You should also disable the *Resource Follows Master* settings until after the cluster conversion is finalized. For instructions, see [Chapter 14, “Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers,” on page 371](#).

- ♦ After OES 11x nodes are added to the OES 2 SP3 cluster, the Linux POSIX cluster resources must be mounted on and fail over to only OES 11x nodes.
- ♦ NSS pool cluster resources can fail over to any node in the mixed-mode OES 2 SP3 and OES 11x cluster.
- ♦ No storage management functions should be executed while a cluster is in mixed-mode unless you must do so as part of the documented upgrade process. Do not attempt to create, delete, expand, or modify the properties for partitions, pools, or volumes for any shared resources in the cluster.

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 11x nodes, create the NSS pool cluster resource on an OES 2 SP3 node, then restart the OES 11x nodes. This allows the OES 11x 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.

7.4 Adding OES 11x Nodes to an OES 2 SP3 Cluster (Rolling Cluster Upgrade)

Performing a rolling cluster upgrade from OES 2 SP3 to OES 11x 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 11x server is added at a time, either by adding a new node or upgrading an OES 2 SP3 node:

- ♦ Add an OES 11x node and migrate resources to it.
- ♦ Remove (cluster leave command) an OES 2 SP3 node, upgrade it to OES 11x, return it to the cluster, and migrate resources to it.

After all desired OES 11x nodes are present in the cluster, remove any remaining OES 2 SP3 nodes that will not be upgraded, migrating their resources to OES 11x nodes. 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 SP3 cluster nodes.
You will probably want to install the same OES components on the OES 11x nodes that you add to the cluster.
- 2 Prepare some OES 11x nodes for the cluster, but do not allow them to join the cluster.

IMPORTANT: If you are upgrading an OES 2 SP3 cluster node to OES 11x, leave the cluster before the upgrade and do not rejoin until [Step 4](#).

- 2a Install SUSE Linux Enterprise Server (SLES) 11 SP2 and the OES Services from the OES 11x 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 85.
 - 2d Configure the node for the existing OES 2 SP3 cluster as described in [Section 5.5.6, "Adding a Node to an Existing Cluster,"](#) on page 91.
-

Parameter	Action
LDAP administrator user password	Specify the password of the user name shown, then click <i>OK</i> . 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 85.
New or existing cluster	Select <i>Existing Cluster</i> .
Directory Server Address	Select the check box next to the IP address of the LDAP server you want to use as the default for this node. The local LDAP server is selected by default. You can add, remove, or change the order of available LDAP servers for the node after the setup is complete. See Section 8.11.1, "Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster," on page 134.
Cluster FDN	Click <i>Browse</i> , navigate the tree to select the existing Cluster object for the OES 2 SP3 cluster, click <i>OK</i> , then click <i>Next</i> .
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.

Parameter	Action
Start Cluster Services now	<p>Deselect the check box.</p> <p>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.</p>

- 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 11x 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 14.4, "Modifying the Scripts for CSM Resources without a Segment Manager,"](#) on page 374. This is the typical configuration in OES 2 SP3.
 - ♦ **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 14.5, "Modifying the Scripts for CSM Resources with a Segment Manager,"](#) on page 378. This is an optional configuration in OES 2 SP3 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 11x nodes.
 - 3f Repeat [Step 3a](#) to [Step 3e](#) for each Linux POSIX cluster resource.
- 4 For each of the new OES 11x nodes, start Novell Cluster services manually by using one of the following methods, and allow the node to join the OES 2 SP3 cluster.
 - ♦ Reboot the cluster server.
 - ♦ At a command prompt, go to the `/etc/init.d` directory, then enter the following as the root user:


```
./novell-ncs start
```
 - ♦ At a command 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 11x nodes as described in [Section 14.7, "Configuring the Preferred Nodes and Cluster Settings for a CSM Cluster Resource,"](#) on page 386.
 - 5b In iManager, bring the resource online on its first preferred node. Click *Clusters > My Clusters*, then select the cluster. On the *Cluster Manager* page, select the resource object, then click *Online*.

You can also use the `cluster online` command. Open a terminal console as the `root` user, then enter

```
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 11x nodes, ensure that you modify the *Preferred Nodes* list accordingly for each of the Linux POSIX cluster resources.
- 7 For each of the existing OES 2 SP3 cluster nodes that you want to upgrade to OES 11x:
 - 7a Log in to the node as the `root` user, then open a terminal console.
 - 7b Use the `cluster leave` command to remove the node from the cluster.

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

You can also manually cluster migrate the resources to another server in the cluster prior to bringing down the server.
 - 7c Stop Novell Cluster Services by entering

```
rcnovell-ncs stop
```
 - 7d Apply the OES 11x and Novell Cluster Services upgrade, and fully patch the system.
 - 7e Start Novell Cluster Services by entering

```
rcnovell-ncs start
```
 - 7f Use the `cluster join` command to rejoin the node to the cluster.
 - 7g Modify the *Preferred Nodes* list accordingly for each of the cluster resources.
 - 7h Repeat these steps for each OES 2 SP3 node that you want to upgrade.
- 8 Bring down each of the OES 2 SP3 cluster nodes 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.
- 9 After the last cluster resource has been moved to the OES 11x nodes and all OES 2 SP3 nodes have left the cluster, the conversion is complete.

7.5 Updating the Clusters Plug-in for Novell iManager

Beginning in OES 11 SP1, the Clusters plug-in replaces the *Cluster Manager*, *BCC Manager*, *Cluster Event Log*, and *Cluster Options* menu options with the *My Clusters* and *My Resources* options. The former menu options appear as tabs when you select a cluster from the *My Clusters* page. The plug-in is updated automatically to show new tasks, and hide old ones. However, the Role-Based Services role for Clusters is not updated automatically.

If you use Role-Based Services, the Clusters plug-in must be updated manually on the Role-Based Services Configuration page in order to pick up the modified menu options. For information, see [Section 5.8, “Installing or Updating the Clusters Plug-in for iManager,” on page 96](#).

7.6 Upgrading a Cluster with DST Resources from OES 2 SP3 to OES 11x

You can upgrade a cluster from OES 2 SP3 to OES 11 or from OES 2 SP3 to OES 11 SP1. The same release version of OES 11 must be installed and running on each new or upgraded node in the cluster.

The procedures in this section refer to OES 11x. A specific version of OES 11 is mentioned only if the procedure applies only to that version of the platform.

If you replace an OES 2 SP3 node with an OES 11x node instead of using an in-place upgrade to OES 11x, you must manually configure the global policies and move the *All Shadowed Volumes* policies data to the OES 11x node.

- 1 On the replacement OES 11x node, set the same global configuration policies that are set on the OES 2 SP3 nodes.
For information, see “[Configuring DST Global Policies](#)” in the *OES 11 SP1: Dynamic Storage Technology Administration Guide*.
- 2 If you have *All Shadowed Volumes* policies set on the OES 2 SP3 nodes, copy the information in the `/usr/novell/sys/. _NETWARE/shadow_policy.xml` file to the same path on the replacement OES 11x nodes.
- 3 In iManager, use the Clusters plug-in to modify the preferred nodes list for each DST pool cluster resource to specify only OES 11x nodes as fail-over candidates.
- 4 Launch a terminal console, then cluster migrate the DST pool cluster resource to an OES 11x node:

```
cluster migrate <dst_volume_cluster_resource> [oes11x_node_name]
```

If the node name is not specified, the resource automatically is brought online on an available node in its preferred nodes list. To bring it online on a specific node, you can specify a node in its preferred nodes list. The resource will be mounted there if it is possible to do so; that is, if the node is available and there are no resource conflicts.

8 Configuring Cluster Policies, Protocols, and Properties

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 8.1, “Understanding Cluster Settings,” on page 113](#)
- ♦ [Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 117](#)
- ♦ [Section 8.3, “Configuring Quorum Membership and Timeout Policies,” on page 121](#)
- ♦ [Section 8.4, “Configuring Cluster Protocols,” on page 122](#)
- ♦ [Section 8.5, “Configuring Cluster Event Email Notification,” on page 124](#)
- ♦ [Section 8.6, “Configuring Cascade Failover Prevention,” on page 126](#)
- ♦ [Section 8.7, “Configuring the Cluster Node Reboot Behavior,” on page 129](#)
- ♦ [Section 8.8, “Configuring STONITH,” on page 130](#)
- ♦ [Section 8.9, “Viewing or Modifying the NCS Proxy User Assignment in the NCS_Management_Group,” on page 132](#)
- ♦ [Section 8.10, “Viewing or Modifying the Cluster IP Address or Port,” on page 133](#)
- ♦ [Section 8.11, “Moving a Cluster, or Changing the Node IP Addresses, LDAP Servers, or Administrator Credentials for a Cluster,” on page 134](#)
- ♦ [Section 8.12, “What’s Next,” on page 138](#)

8.1 Understanding Cluster Settings

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

- ♦ [Section 8.1.1, “Cluster Policies,” on page 114](#)
- ♦ [Section 8.1.2, “Cluster Priorities,” on page 115](#)
- ♦ [Section 8.1.3, “Cluster Protocols,” on page 115](#)
- ♦ [Section 8.1.4, “RME Groups,” on page 116](#)
- ♦ [Section 8.1.5, “BCC,” on page 116](#)
- ♦ [Section 8.1.6, “Cascading Failover Prevention,” on page 116](#)
- ♦ [Section 8.1.7, “Cluster Node Reboot Behavior,” on page 116](#)
- ♦ [Section 8.1.8, “STONITH,” on page 116](#)
- ♦ [Section 8.1.9, “Cluster Information,” on page 116](#)

8.1.1 Cluster Policies

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

Table 8-1 Cluster Policies

Property	Description
Cluster IP address	<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. The IP address is bound to the master node and remains with the master node regardless of which server is the master node.</p> <p>Rarely, you might need to modify this value. For information, see Section 8.10, "Viewing or Modifying the Cluster IP Address or Port," on page 133.</p> <p>IMPORTANT: If you modify the master IP address, a cluster restart is required before iManager can continue to manage the cluster.</p>
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. You can change the port number to any other value that does not cause a conflict. For information, see Section 8.10, "Viewing or Modifying the Cluster IP Address or Port," on page 133.</p>
Quorum membership - Number of nodes	<p>Specifies number of nodes that must be up and running in the cluster before clusters begin to load.</p> <p>Specify a value between 1 and the number of nodes. Set this value to a number greater than 1 so that all resources are not automatically loaded on the first server that is brought up in the cluster. For example, if you set the membership value to 4, there must be four serves up in the cluster before any resource will load and start.</p> <p>For instructions, see Section 8.3, "Configuring Quorum Membership and Timeout Policies," on page 121.</p>
Quorum timeout	<p>Specifies the maximum amount of time to wait for the specified quorum to be met. If the timeout period elapses before a quorum is achieved, resources automatically begin loading on whatever number of nodes are actually up and running in the cluster.</p> <p>You can specify the time in seconds or minutes.</p> <p>For instructions, see Section 8.3, "Configuring Quorum Membership and Timeout Policies," on page 121.</p>

Property	Description
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>Specifies the type of cluster events for notifications. You can receive notifications for only critical events such as node failure or a resource going comatose, or you can receive notifications for all cluster state changes and resource state changes.</p> <p>Specifies whether to receive messages in XML format. XML format messages can be interpreted with a parser and formatted to customize the message information for your specific needs.</p> <p>For instructions, see Section 8.5, “Configuring Cluster Event Email Notification,” on page 124.</p>

8.1.2 Cluster Priorities

Cluster priorities determine the load priority of individual cluster resources. You can manage cluster resource priorities in iManager by going to the *Clusters > My Clusters*, selecting a cluster, then selecting *Action > Edit Properties > Priorities*. For information, see [Section 10.10, “Configuring Resource Priorities for Load Order,” on page 200.](#)

8.1.3 Cluster Protocols

[Table 8-2](#) describes the configurable cluster protocols properties that govern inter-node communication transmission and tolerances. You can manage cluster protocols policies in Novell iManager by going to the *Clusters > My Clusters*, selecting a cluster, then selecting *Action > Edit Properties > Protocols*. For instructions, see [Section 8.4, “Configuring Cluster Protocols,” on page 122.](#)

Table 8-2 Cluster Protocols

Property	Description
Heartbeat	Specifies the interval of time in seconds between signals sent by each of the non-master nodes in the cluster to the master node to indicate that it is alive. The default is 1 second.
Tolerance	Specifies the maximum amount of time in seconds that a master node waits to get an alive signal from a non-master node before considering that node to have failed and removing it from the cluster. The default is 8 seconds.
Master watchdog	<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. The default is 1 second.</p> <p>Modify this parameter setting only when supervised by Novell Technical Support (http://www.novell.com/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. The default is 8 seconds.</p> <p>Modify this parameter setting only when supervised by Novell Technical Support (http://www.novell.com/support).</p>
Maximum retransmits	This value is set by default and should not be changed. The default is 30.

8.1.4 RME Groups

The Resource Mutual Exclusion (RME) Groups page allows you to view or define sets of resources that must not run on the same node at the same time. You can manage RME Groups in iManager by going to the *Clusters > My Clusters*, selecting a cluster, then selecting *Action > Edit Properties > RME Groups*. For information, see [Section 10.11, “Configuring Resource Mutual Exclusion Groups,” on page 201](#).

8.1.5 BCC

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

8.1.6 Cascading Failover Prevention

Novell Cluster Services added the cascading failover prevention function that detects if a node has failed because of a bad cluster resource and prevents that bad resource from failing over to other servers in the cluster. It is enabled by default. This function was previously available on NetWare, but was not available for OES 1 Linux and OES 2 Linux. For information, see [Section 8.6, “Configuring Cascade Failover Prevention,” on page 126](#).

8.1.7 Cluster Node Reboot Behavior

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. On certain occasions, you might want to prevent a downed cluster node from rebooting so you can troubleshoot problems. To control the cluster node reboot behavior, you can use the directive `kernel.panic` in the Linux `/etc/sysctl.conf` file to prevent a reboot, or to allow an automatic reboot and to specify the number of seconds to delay the reboot. For information, see [Section 8.7, “Configuring the Cluster Node Reboot Behavior,” on page 129](#).

8.1.8 STONITH

The STONITH (shoot-the-other-node-in-the-head) capability allows Novell Cluster Services to kill a suspect node by using remote power control. Unlike a poison pill, it does not require a response from the suspect node. STONITH is used after a poison pill is issued; it does not replace the poison pill. For information, see [Section 8.8, “Configuring STONITH,” on page 130](#).

8.1.9 Cluster Information

After you install 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 an existing cluster. The YaST-based configuration is not used to modify the settings for an existing cluster. For information about modifying the settings for an existing cluster, see:

- ♦ [Section 8.9, “Viewing or Modifying the NCS Proxy User Assignment in the NCS_Management_Group,” on page 132](#)

- [Section 8.10, “Viewing or Modifying the Cluster IP Address or Port,” on page 133](#)
- [Section 8.11, “Moving a Cluster, or Changing the Node IP Addresses, LDAP Servers, or Administrator Credentials for a Cluster,” on page 134.](#)

8.2 Setting Up a Personalized List of Clusters to Manage

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

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

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

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

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


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

- [Section 8.2.1, “Adding Clusters to Your My Clusters List,” on page 117](#)
- [Section 8.2.2, “Viewing Information about Clusters in Your My Clusters List,” on page 118](#)
- [Section 8.2.3, “Personalizing the My Clusters Display,” on page 119](#)
- [Section 8.2.4, “Removing Clusters from a My Clusters List,” on page 120](#)

8.2.1 Adding Clusters to Your My Clusters List

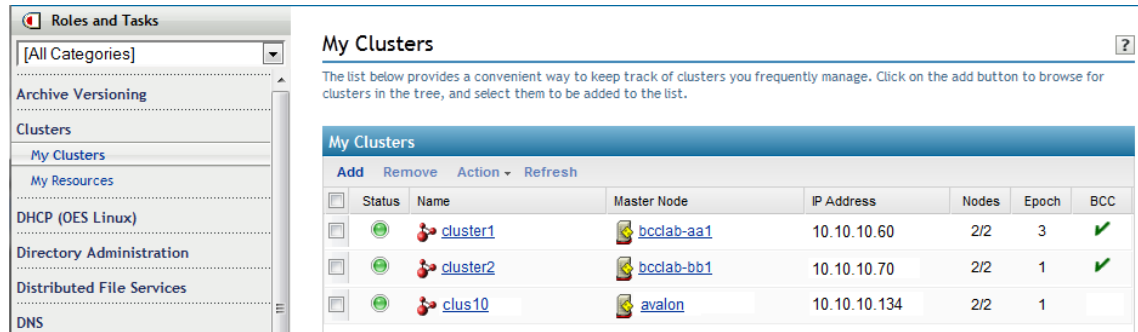
- 1 Log in to iManager as a cluster administrator.
- 2 In *Roles and Tasks*, select *Clusters > My Clusters*.

The list is initially empty.



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

Newly selected clusters are added to your personalized list.

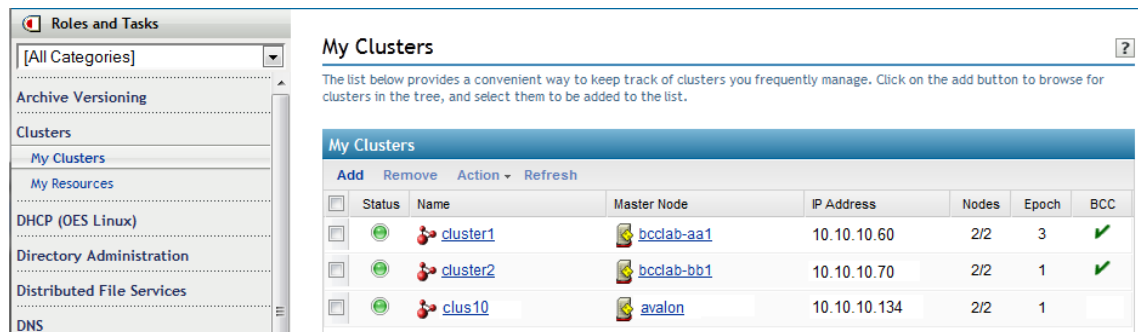


- 5 (Optional) Personalize the display as described in [Section 8.2.3, “Personalizing the My Clusters Display,”](#) on page 119.

8.2.2 Viewing Information about Clusters in Your My Clusters List

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

The My Clusters page displays your personalized list of clusters in the tree.



- 3 View the Status icon next to a Cluster object to understand the current overall state of the cluster.

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

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

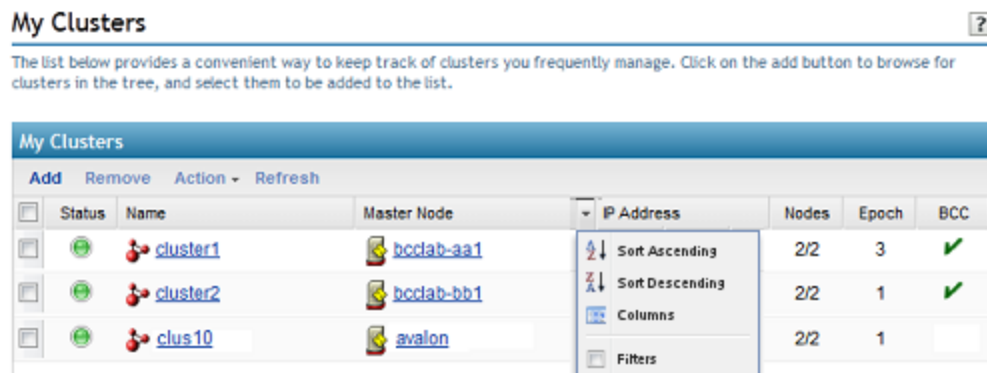
Parameter	Description
Name	The name assigned to the cluster.
Distinguished Name	(Not displayed by default) The dot-delimited fully distinguished name of the cluster, such as <code>clustername.context.domain</code> .
Master Node	The name of the currently assigned master node for the cluster.
IP Address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

- 5 (Optional) Personalize the display as described in [Section 8.2.3, “Personalizing the My Clusters Display,”](#) on page 119.

8.2.3 Personalizing the My Clusters Display

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

- 1 Log in to iManager as a cluster administrator.
- 2 In *Roles and Tasks*, select *Clusters > My Clusters*.
- 3 Select *Refresh*, then choose a refresh rate that is long enough to allow you to view all items in the list.
- 4 Click or mouse-over a column heading to activate its options, then click the arrow to access the menu for that column.
- 5 Perform one of the available actions when you access the column’s menu. Some actions are not available for all fields.
 - ♦ **Sort:** Select *Sort Ascending* (A to Z) or *Sort Descending* (Z to A) to specify the preferred sort order. Cluster objects are sorted based on values in the selected column. Numbers are treated as text values and are sorted alphabetically, not numerically.



- ♦ **Columns:** Select *Columns* to display the parameters, select or deselect a parameter's check box to add or remove the column, then press Enter to apply the changes. The parameters are described in [Step 4 in Section 9.14, "Viewing the Cluster Node Properties,"](#) on page 156.

My Clusters

The list below provides a convenient way to keep track of clusters you frequently manage. Click on the add button to browse for clusters in the tree, and select them to be added to the list.

My Clusters							
Add Remove Action Refresh							
<input type="checkbox"/>	Status	Name	Master Node	IP Address	Nodes	Epoch	BCC
<input type="checkbox"/>		cluster1	bcclab-aa1	Sort Ascending	2/2	3	
<input type="checkbox"/>		cluster2	bcclab-bb1	Sort Descending	2/2	1	
<input type="checkbox"/>		clus10		Columns	2/2	1	

- ♦ **Filter:** Specify a filter to display only the Cluster objects with values that match.

My Clusters

The list below provides a convenient way to keep track of clusters you frequently manage. Click on the add button to browse for clusters in the tree, and select them to be added to the list.

My Clusters							
Add Remove Action Refresh							
<input type="checkbox"/>	Status	Name	Master Node	IP Address	Nodes	Epoch	BCC
<input type="checkbox"/>		cluster1	bcclab-aa1	Sort Ascending	2/2	3	
<input type="checkbox"/>		cluster2	bcclab-bb1	Sort Descending	2/2	1	
<input type="checkbox"/>		clus10		Columns	2/2	1	

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

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

8.2.4 Removing Clusters from a My Clusters List

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

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

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

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

8.3 Configuring Quorum Membership and Timeout Policies

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

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the check box of the cluster that you want to manage.
If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, “Setting Up a Personalized List of Clusters to Manage,”](#) on page 117.
- 3 Select *Action > Edit Properties* to open the Properties dialog box.
- 4 Click the *Policies* tab.

The screenshot shows the 'Policies' tab of the 'Edit Properties' dialog for cluster 'clus1.ncs.novell'. The dialog has tabs for Policies, Priorities, Protocols, RME Groups, and Business Continuity. Below the tabs is a descriptive text: 'View or change cluster IP address, port numbers, quorum triggers, and cluster email notification. Changing the IP address requires a cluster restart before iManager can continue to manage the cluster.' The dialog is divided into three main sections: 'Cluster Information', 'Quorum Triggers', and 'Notification'. 'Cluster Information' includes fields for 'IP Address' (10.10.10.134) and 'Port' (7023). 'Quorum Triggers' includes 'Timeout' (60 Seconds) and 'Number of Nodes' (1). 'Notification' includes a checkbox for 'Enable Cluster Notification Events', an 'Email Addresses' field with a dropdown showing 'cluster@novell.com', and three radio buttons: 'Receive Only Critical Events' (selected), 'Verbose Messages', and 'XML Messages' (checked). At the bottom are 'OK', 'Cancel', and 'Apply' buttons.

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

For information, see [Section 8.3.1, “Quorum Membership \(Number of Nodes\),”](#) on page 122.

- 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 8.3.2, “Quorum Timeout,”](#) on page 122.

- 7 Click *Apply* or *OK* to save your changes.
- 8 Restart Cluster Services. Open a terminal console as the `root` user, then enter the following at the command prompt:

```
rcnovell-ncs restart
```


8.3.1 Quorum Membership (Number of Nodes)

A cluster quorum is the number of nodes that must be running in the cluster before resources begin to load. When you first bring up servers in your cluster, 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 do not automatically load on the first server that is brought up in the cluster. For example, if you set the *Number of Nodes* value to 4, there must be four servers up in the cluster before any resource loads and starts.

8.3.2 Quorum Timeout

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

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

8.4 Configuring Cluster Protocols

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

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

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

If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, “Setting Up a Personalized List of Clusters to Manage,”](#) on page 117.
- 3 Click *Action > Edit Properties* to open the Properties dialog box.

4 Click the *Protocols* tab.

The Protocols page allows you to view or change cluster protocol settings such as heartbeat, tolerance, master/slave watchdogs, and maximum retransmits. The Protocols page also lets you view the script used to configure the cluster protocol settings, but not to change it. Changes made to the protocols setting automatically update the scripts.

clus1.ncs.novell

Policies Priorities **Protocols** RME Groups Business Continuity

View or change cluster protocol settings.

Heartbeat (seconds): Master Watchdog (seconds):

Tolerance (seconds): Slave Watchdog (seconds):

Maximum Retransmits:

```
panning clusterid 1623659656
heartbeat rate_usecs 1000000
censustaker tolerance 8000000
sequencer master_watchdog 1000000
sequencer slave_watchdog 8000000
sequencer retrans_max 30

nodeid 137.65.67.37 0
```

OK Cancel Apply

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.

8.4.1 Heartbeat

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

8.4.2 Tolerance

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

8.4.3 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 \(http://www.novell.com/support\)](http://www.novell.com/support).

8.4.4 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 \(http://www.novell.com/support\)](http://www.novell.com/support).

8.4.5 Maximum Retransmits

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

8.5 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
```

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

- [Section 8.5.1, “Configuring an Email Port,” on page 124](#)
- [Section 8.5.2, “Configuring the Notification Policies,” on page 125](#)

8.5.1 Configuring an Email Port

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)

8.5.2 Configuring the Notification Policies

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

If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, “Setting Up a Personalized List of Clusters to Manage,”](#) on page 117.

- 3 Click *Action > Edit Properties* to open the Properties dialog box.
- 4 Click the *Policies* tab.

clus1.ncs.novell

Policies | Priorities | Protocols | RME Groups | Business Continuity

View or change cluster IP address, port numbers, quorum triggers, and cluster email notification. Changing the IP address requires a cluster restart before iManager can continue to manage the cluster.

Cluster Information

IP Address: 10.10.10.134

Port: 7023

Quorum Triggers

Timeout: 60 Seconds

Number of Nodes: 1

Notification

☐ Enable Cluster Notification Events

Email Addresses: cluster@novell.com

☒ Receive Only Critical Events

☐ Verbose Messages

☒ XML Messages

OK Cancel Apply

- 5 Under *Notification*, select or deselect the *Enable Cluster Notification Events* check box to enable or disable email notification.

Notification

☒ Enable Cluster Notification Events

Email Addresses: cluster@example.com

☒ Receive Only Critical Events

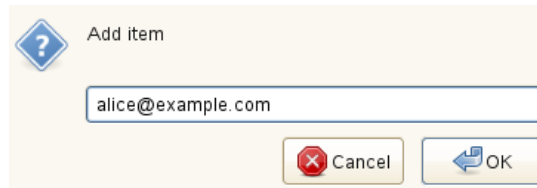
☐ Verbose Messages

☒ XML Messages

6 If you enable email notification, configure the notification settings:

6a Add up to eight email addresses in the field provided.

You can click the buttons next to the field to add, delete, or edit email addresses. Repeat this process for each email address you want on the notification list.



6b Specify the type of cluster events you want administrators to receive messages for.

- ♦ **Only Critical Events:** (Default) Select this option to receive notification only of critical events like a node failure or a resource going comatose.
- ♦ **Verbose Messages:** Select this option to receive notification of all cluster state changes including critical events, resource state changes, and nodes joining and leaving the cluster.

6c Specify whether you want to receive notification of all cluster state changes in XML format by selecting the *XML Messages* option.

XML is selected by default. XML format messages can be interpreted and formatted with a parser that lets you customize the message information for your specific needs.

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

8.6 Configuring Cascade Failover Prevention

A cascading failover occurs when a bad cluster resource causes a server to fail, then fails over to another server causing it to fail, and then continues failing over to and bringing down additional cluster servers until possibly all servers in the cluster have failed.

Novell Cluster Services provides a Cascade Failover Prevention function. It detects if a node has failed because of a bad cluster resource, and prevents that bad resource from failing over to other servers in the cluster.

- ♦ [Section 8.6.1, “Understanding the Cascade Failover Prevention Quarantine,” on page 126](#)
- ♦ [Section 8.6.2, “Releasing a Resource from Quarantine,” on page 128](#)
- ♦ [Section 8.6.3, “Enabling or Disabling Cascade Failover Prevention,” on page 128](#)

8.6.1 Understanding the Cascade Failover Prevention Quarantine

The cascade failover prevention quarantine puts a resource in a comatose state rather than letting it load on (and potentially cause to fail) other cluster nodes. A resource can be quarantined if the systematic analysis of the logged node failures determines the following conditions to be true:

- ♦ The resource is likely responsible for several consecutive server failures, unrelated to interference from failures of other resources.

The consecutive failures might have occurred on different nodes in the cluster. If the resource loads successfully on any node, the failure count for the resource starts over.

- ♦ Loading the resource will put the cluster in grave danger.

Novell Cluster Services does the following to determine if a resource should be put into quarantine:

1. Traces the history of node failures for the suspected bad resource. This includes:
 - ♦ What node the resource was running on or loading on
 - ♦ If the node failed
 - ♦ The state the resource was in when the node failed
 - ♦ If there were other resources trying to load when the node failed
2. Repeats the above process until one of the following happens:
 - ♦ The end of the cluster log file is reached
 - ♦ Enough consecutive node failures are found
 - ♦ Found that the node did not fail
 - ♦ Found that the whole cluster was down
 - ♦ The entries in the log file are more than 365 days old

If the resource attempts to load on a node where it was previously loaded and there are additional nodes still available in the cluster, it will not be quarantined and will be allowed to load. Also, a resource is not quarantined when it is initially brought online.

Factors that might contribute to a resource being quarantined include:

- ♦ A large number of consecutive node failures
- ♦ No other resources are causing node failures
- ♦ The resource never reaches a running state

Factors that might help prevent a resource from being quarantined include:

- ♦ A small number of consecutive node failures
- ♦ The resource has failed on this node previously
- ♦ Other resources are causing node failures
- ♦ The resource reaches a running state
- ♦ There is one node left up and running in the cluster

The resource quarantine is disabled if:

- ♦ Cascade Failover Prevention is turned off.

For information, see [“Releasing a Resource from Quarantine” on page 128](#).

- ♦ There is no shared storage (SAN) or SBD partition.
- ♦ There are enough nodes in the cluster to form a quorum.

8.6.2 Releasing a Resource from Quarantine

While a resource is in quarantine, you can still manually take the resource from the comatose state to an offline state, and then bring it online or cluster migrate it to other cluster nodes.

To get the resource out of quarantine so that is once again able to automatically fail over:

- 1 Log in to the node as the root user.
- 2 Disable Cascade Failover Prevention. See [“Disabling Cascade Failover Prevention” on page 128](#).
- 3 Re-enable Cascade Failover Prevention. See [“Manually Enabling Cascade Failover Prevention” on page 129](#).

8.6.3 Enabling or Disabling Cascade Failover Prevention

The Cascade Failover Prevention function is enabled by default when you install Novell Cluster Services. You can control Cascade Failover Prevention by creating a configuration file in the `/etc/modprobe.d` folder, and using it to disable and enable the `ncs_cascade_failover_detection_flag`. Two detection modes are supported:

- ♦ A value of 0 disables the function.
- ♦ A value of 2 enables the function. This value is assumed if the configuration file does not exist, or if the flag line is not present in the file. The file does not exist by default.

The setting in the file is specific to a node, and is set separately on each node. After you modify the setting on a node, you must manually unload and reload Novell Cluster Services software on the server to apply change. The setting persists through patches and upgrades.

- ♦ [“Disabling Cascade Failover Prevention” on page 128](#)
- ♦ [“Manually Enabling Cascade Failover Prevention” on page 129](#)

Disabling Cascade Failover Prevention

You might need to disable the Cascade Failover Prevention function for the following reasons:

- ♦ To get a resource out of quarantine and allow it to once again automatically fail over, you can temporarily disable the Cascade Failover Prevention function on the node where the resource was taken offline.
- ♦ To stop using the Cascade Failover Prevention function in the cluster, you can disable the function on every node in the cluster.

To disable Cascade Failover Prevention:

- 1 Log in to the node as the root user.
- 2 Navigate to the `/etc/modprobe.d` folder.
- 3 In a text editor, create a configuration file (such as `novell-ncs.conf`) under the `/etc/modprobe.d` folder.
If the file already exists, open it.
- 4 Add the following content to the file. If the line already exists, change the flag value to 0.

```
options crm ncs_cascade_failover_detection_flag=0
```

The flag value 0 disables Cascade Failover Prevention for the node.
- 5 Save the file, then close the text editor.

- 6 Open a terminal console, then restart Novell Cluster Services software to apply the modified Cascade Failover Prevention setting on this node:

```
rcnovell-ncs restart
```

- 7 Repeat this procedure on other nodes where you want to disable Cascade Failover Prevention.
- 8 (Optional) Re-enable Cascade Failover Prevention as described in [“Manually Enabling Cascade Failover Prevention” on page 129](#).

Manually Enabling Cascade Failover Prevention

You might need to manually re-enable Cascade Failover Prevention:

- ♦ If you disabled the function to release resources from the quarantine
- ♦ If you disabled the function for other reasons

To enable Cascade Failover Prevention:

- 1 Log in to the node as the root user.
- 2 Navigate to the `/etc/modprobe.d` folder.
- 3 Use one of the following methods to enable Cascade Failover Prevention:
 - ♦ **Delete the file:** Delete the `/etc/modprobe.d/novell-ncs.conf` file that you created in [“Disabling Cascade Failover Prevention” on page 128](#).
 - ♦ **Change the value to 2:** In a text editor, open the `/etc/modprobe.d/novell-ncs.conf` file that you created in [“Disabling Cascade Failover Prevention” on page 128](#), change the flag value from 0 to 2, then save the file.

```
options crm ncs_cascade_failover_detection_flag=2
```

The flag value 2 enables Cascade Failover Prevention for the node.

- ♦ **Delete the line:** In a text editor, open the `/etc/modprobe.d/novell-ncs.conf` file that you created in [“Disabling Cascade Failover Prevention” on page 128](#), remove the `ncs_cascade_failover_detection_flag` line, then save the file.
- 4 Open a terminal console, then restart Novell Cluster Services software to apply the modified Cascade Failover Prevention setting on this node:

```
rcnovell-ncs restart
```

- 5 Repeat this procedure on other nodes where you want to enable Cascade Failover Prevention.

8.7 Configuring the Cluster Node Reboot Behavior

If LAN connectivity is lost between a cluster node and the other nodes in the cluster, it is possible that the lost node will be automatically shut down by the other cluster nodes. This is normal cluster operating behavior, and it prevents the lost node from trying to load cluster resources because it cannot detect the other cluster nodes. The 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. On certain occasions, you might want to prevent a downed cluster node from rebooting so you can troubleshoot problems.

You can use the directive `kernel.panic` in the Linux `/etc/sysctl.conf` file to prevent a reboot, or to allow an automatic reboot and to specify the number of seconds to delay the reboot. Set the `kernel.panic` value to 0 for no reboot after a node shutdown. Set the `kernel.panic` value to a

positive integer value to indicate that the server should be rebooted after waiting the specified number of seconds. For information about using the Linux `sysctl`, see the Linux man pages on `sysctl(8)` and `sysctl.conf(5)`.

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

```
kernel.panic = 0
```

- 3 Set the `kernel.panic` value to 0 or to a positive integer value, depending on the desired behavior.

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

```
kernel.panic = 0
```

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

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

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

```
kernel.panic = 60
```

- 4 Save your changes.

8.8 Configuring STONITH

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 8.8.1, “Sample Script for HP iLO,” on page 131](#)
- ♦ [Section 8.8.2, “Sample Script for Web-Based Power Switches,” on page 131](#)

8.8.1 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
```

8.8.2 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
```

8.9 Viewing or Modifying 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 86](#):

- ♦ 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 command prompt, enter

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


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

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

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

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

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

8.10 Viewing or Modifying the Cluster IP Address or Port

The cluster IP address is assigned when you install Novell Cluster Services. It is used to access the master node when you are managing the cluster. The cluster IP address is separate from the server IP address and is required for some external programs to get cluster status alerts. The IP address is bound to the master node and remains with the master node regardless of which server is the master node. The Master IP Address Resource is created automatically when you install Cluster Services on the first node of the cluster.

The cluster IP address normally does not need to be changed, but it can be if needed. Changing the IP address requires a cluster restart before iManager can continue to manage the cluster.

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

If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, "Setting Up a Personalized List of Clusters to Manage,"](#) on page 117.

- 3 Click *Action > Edit Properties* to open the Properties dialog box.
- 4 Click the *Policies* tab.

clus1.ncs.novell

Policies **Priorities** **Protocols** **RME Groups** **Business Continuity**

View or change cluster IP address, port numbers, quorum triggers, and cluster email notification. Changing the IP address requires a cluster restart before iManager can continue to manage the cluster.

Cluster Information

IP Address: 10.10.10.134

Port: 7023

Quorum Triggers

Timeout: 60 Seconds

Number of Nodes: 1

Notification

☐ Enable Cluster Notification Events

Email Addresses: cluster@novell.com

☒ Receive Only Critical Events

☐ Verbose Messages

☒ XML Messages

OK Cancel Apply

- 5 Under *Cluster Information*, specify the new value for the IP address or port.
- 6 Click *Apply* or *OK* to save your changes.
- 7 Restart Cluster Services. Open a terminal console as the `root` user, then enter the following at a command prompt:

```
rcnovell-ncs restart
```

A cluster restart is required before iManager can continue to manage the cluster.

8.11 Moving a Cluster, or Changing the Node IP Addresses, LDAP Servers, or Administrator Credentials for a Cluster

Use the instructions in this section to change the IP address of the cluster, information about the LDAP servers that the cluster uses, or the credentials used to administer the cluster.

- ♦ [Section 8.11.1, “Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster,” on page 134](#)
- ♦ [Section 8.11.2, “Moving a Cluster or Changing IP Addresses of Cluster Nodes and Resources,” on page 136](#)

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. If Novell eDirectory is not installed on a node, it looks to the LDAP server list for information about which LDAP server to use.

You must modify this cluster information in the following cases:

- ♦ Changing the Administrator user name and password for the cluster
- ♦ Changing the password for the existing Administrator user name
- ♦ Changing the IP address information about the existing LDAP servers
- ♦ Assigning a different LDAP server for the cluster to use

- ♦ Changing the order of the servers in the LDAP server list
As a best practice, you should list the LDAP servers in the following order:
 - ♦ local to the cluster
 - ♦ closest physical read/write replica
- ♦ Adding LDAP servers to the list of ones that the cluster can use.

You can modify these settings at any time. 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 administrator user name that you plan to use meets the requirements specified in [Section 4.1, “Cluster Administration Requirements,” on page 49](#).
- 2 Ensure that the LDAP servers that you want to specify for the cluster are in the same Novell eDirectory tree, and that the IP addresses meet the requirements in [Section 4.2, “IP Address Requirements,” on page 52](#).
- 3 Log in as the root user to the master node of the cluster.
- 4 In a text editor, create a text file, specify the configuration information for the 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"
```

- 5 As the root user, enter the following command at a command prompt:

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

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

- 6 Delete the configuration file that you created.
- 7 For each of the other nodes in the cluster, log in as the root user for the node, then repeat [Step 4](#) to [Step 6](#).

Modifying the information on each node allows iManager to manage the cluster after a different node becomes the master. This step is necessary because credentials are stored on CASA, and CASA does not synchronize across cluster nodes.

- 8 Push this update to all nodes on the cluster by entering the following as the root user on one of the cluster nodes:

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

8.11.2 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 136](#)
- ♦ [“Changing the IP Addresses of Cluster Resources” on page 136](#)
- ♦ [“Changing the IP Addresses of Servers in a Cluster” on page 137](#)
- ♦ [“Modifying the Cluster Configuration Information” on page 137](#)

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 52](#). 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 52](#).

Changing the IP Addresses of Cluster Resources

Before you modify the server IP address for a server in a cluster, you must change the IP addresses of all of the cluster resources that run on it:

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the cluster that you want to manage.

If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 117](#).
- 3 Offline the cluster resources whose IP addresses are changing. On the Cluster Manager page, select the check boxes next to the cluster resources, then click *Offline*.
- 4 For each NSS pool cluster resource, modify the IP address on the resource’s Protocols page.
 - 4a On the Cluster Manager or Cluster Options page, click the pool cluster resource to open its Properties page.
 - 4b Click the *Protocols* tab.
 - 4c Modify the IP address for the NSS pool cluster resource.
 - 4d Click *OK* to save the changes.

This changes the IP address for the NCS :NCP Server object for the resource and updates the instances of the IP address in the resource's load, unload, and monitor script.

Do not online the resource at this time.

- 5 For the master IP address resource, modify the IP address on the resource's Protocols page.
- 6 For each non-NSS cluster resource, including the master IP address resource, modify the IP address information for the RESOURCE_IP variable as needed in the resource's load, unload, and monitor scripts.
 - 6a On the Cluster Manager page or Cluster Options page, click the non-NSS cluster resource to open its Properties page.
 - 6b Click the *Scripts* tab, then click the *Load Script* link.
 - 6c Edit the script by replacing the RESOURCE_IP variable's value with new IP address. You might also need to edit the values used in the command lines if you did not use the variable there.

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

- 6d Click *Apply* to save the changed script.
 - 6e Make similar changes to the non-NSS resource's *Unload Script* and *Monitor Script*.
 - 6f Repeat this process for each non-NSS cluster resource and the master IP address resource.
- Do not online the resources at this time.
- 7 Stop Novell Cluster Services for every node in the cluster by entering the following at the command prompt as the root user:

```
rcnovell-ncs stop
```

- 8 Continue with [“Changing the IP Addresses of Servers in a Cluster” on page 137](#).

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 136](#)), 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 SP1 Server's IP Address”](#) in the *OES 11 SP1: 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 137](#).

Modifying the Cluster Configuration Information

Before restarting Novell Cluster Services, you must update the cluster configuration information in Novell Cluster Services and Novell 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 134.

- 2 For each node in the cluster, including that of the master IP address resource, modify the *NCS: Network Address* attribute of its Cluster Node object.
 - 2a In iManager, select *Directory Administration > Modify Object*.
 - 2b Browse to locate and select the Cluster Node object of the cluster node you want to manage.
 - 2c In the *Valued Attributes* list, select the attribute *NCS: Network Address*, click *Edit*, modify the IP address, then click *OK*.
 - 2d Repeat this process for each node in the cluster and the master IP resource.
- 3 For the cluster container of the cluster you want to manage, modify the *NCS: Network Address* and *Network Address* attributes of its Cluster object to specify the new IP address information. Both *TCP* and *UDP* addresses need to be replaced.
 - 3a In iManager, select *Directory Administration > Modify Object*.
 - 3b Browse to locate and select the Cluster object of the cluster you want to manage.
 - 3c In the *Valued Attributes* list, select the attribute *NCS: Network Address* (the attribute for the TCP address), click *Edit*, modify the IP address, then click *OK*.
 - 3d In the *Valued Attributes* list, select the attribute *Network Address* (the attribute for the UDP address), click *Edit*, modify the IP address, then click *OK*.
- 4 Ensure that the LDAP server is running before restarting 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 onlining.
- 6 Start Novell Cluster Services by entering the following command at a command prompt as the root user:

```
rcnovell-ncs start
```
- 7 Online the cluster resources:
 - 7a In iManager, select *Clusters > My Clusters*.
 - 7b Select the cluster that you want to manage.
 - 7c On the Cluster Manager page, select the check boxes next to the cluster resources, then click *Online*.

8.12 What's Next

For information about managing the cluster, see [Chapter 9, “Managing Clusters,”](#) on page 139.

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

- ♦ [Chapter 10, “Configuring and Managing Cluster Resources,”](#) on page 177
- ♦ [Chapter 11, “Creating and Managing Service Cluster Resources,”](#) on page 215
- ♦ [Chapter 12, “Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes,”](#) on page 217
- ♦ [Chapter 13, “Configuring and Managing Cluster Resources for Shared LVM Volume Groups,”](#) on page 305
- ♦ [Chapter 14, “Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers,”](#) on page 371

9 Managing Clusters

After you have installed, set up, and configured Novell Cluster Services on Novell Open Enterprise (OES) 11 SP1 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 417.

- ♦ [Section 9.1, “Starting and Stopping Novell Cluster Services,”](#) on page 140
- ♦ [Section 9.2, “Taking the Cluster Down,”](#) on page 142
- ♦ [Section 9.3, “Monitoring Cluster and Resource States,”](#) on page 142
- ♦ [Section 9.4, “Responding to a Start, Failover, or Failback Alert,”](#) on page 145
- ♦ [Section 9.5, “Viewing the Cluster Event Log in iManager,”](#) on page 146
- ♦ [Section 9.6, “Viewing the Cluster Event Log at the Command Line,”](#) on page 149
- ♦ [Section 9.7, “Generating a Cluster Configuration Report,”](#) on page 150
- ♦ [Section 9.8, “Cluster Migrating Resources to Different Nodes,”](#) on page 151
- ♦ [Section 9.9, “Onlining and Offlining \(Loading and Unloading\) Cluster Resources from a Cluster Node,”](#) on page 152
- ♦ [Section 9.10, “Removing \(Leaving\) a Node from the Cluster,”](#) on page 154
- ♦ [Section 9.11, “Joining a Node to the Cluster \(Rejoining a Cluster\),”](#) on page 155
- ♦ [Section 9.12, “Shutting Down Novell Cluster Services When Servicing Shared Storage,”](#) on page 155
- ♦ [Section 9.13, “Enabling or Disabling Cluster Maintenance Mode,”](#) on page 156
- ♦ [Section 9.14, “Viewing the Cluster Node Properties,”](#) on page 156
- ♦ [Section 9.15, “Deleting a Cluster Node from a Cluster,”](#) on page 157
- ♦ [Section 9.16, “Creating or Deleting Cluster SBD Partitions,”](#) on page 157
- ♦ [Section 9.17, “Customizing Cluster Services Management,”](#) on page 174

9.1 Starting and Stopping Novell Cluster Services

Novell Cluster Services automatically starts after it is installed if you choose the *Start Now* option in the configuration. You can also configure it to start automatically when you reboot the server.

If you need to restart `adminfs`, you must stop Novell Cluster Services before you stop `adminfs`. Start `adminfs`, and then start Novell Cluster Services. Alternatively, you can reboot the server to stop and restart both in the proper order.

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
```

- ♦ [Section 9.1.1, “Starting Novell Cluster Services,” on page 140](#)
- ♦ [Section 9.1.2, “Stopping Novell Cluster Services,” on page 140](#)
- ♦ [Section 9.1.3, “Enabling and Disabling the Automatic Start of Novell Cluster Services,” on page 141](#)

9.1.1 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 command prompt, go to the `/etc/init.d` directory and enter

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

- ♦ At the command prompt, enter

```
rcnovell-ncs start
```

- 3 Verify the cluster status:

```
cluster status
```

9.1.2 Stopping Novell Cluster Services

Novell Cluster Services is a kernel-space application. As with any kernel-space application, unloading it is a best effort and is not guaranteed to stop the modules. If Novell Cluster Services is busy providing services from the kernel when you attempt to stop it, the process might not stop, and a reboot might be required.

- 1 Open a terminal console, then log in as the root user.
- 2 Check whether the node is active in the cluster. Enter

```
cluster view
```

If the node is in the cluster (and not still in a process of joining the cluster), it is relatively safe to attempt to stop it.

- 3 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 command prompt, enter

```
rcnovell-ncs stop
```

- 4 Verify the cluster status:

```
cluster status
```

- 5 If you get in a state where the node has not joined the cluster and Novell Cluster Services cannot be stopped:

- 5a Disable autostart on reboot. Enter

```
chkconfig novell-ncs off
```

- 5b Reboot the server.

- 5c Start Novell Cluster Services manually. Enter

```
rcnovell-ncs start
```

- 5d Diagnose why the node had trouble joining the cluster.

- 5e (Optional) Re-enable autostart for the node. Enter

```
chkconfig novell-ncs on
```

9.1.3 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.

9.2 Taking the Cluster Down

The `cluster down` command concurrently removes all cluster nodes from the cluster. This command has the same effect as executing the `cluster leave` command on every server in the cluster. You are prompted to confirm the cluster down. For example, you might need to take the cluster down if an SBD partition fails, and you need to delete and recreate an SBD partition.

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

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

9.3 Monitoring Cluster and Resource States

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

- 1 In iManager, select *Clusters > My Clusters*.
View the overall status of the each of the clusters in your list.
- 2 In your personalized *My Clusters* list, select the name of the cluster you want to manage.
If the cluster does not appear in your list, add the cluster to your list. Click *Add*, browse to select the Cluster object, then click *OK*.
- 3 View the status of the cluster nodes and resources.

My Clusters > cluster1.Servers.novell

cluster1.Servers.novell ?

View the status of both servers and resources in the cluster. You can also migrate cluster resources to different servers in the cluster, or change the state of resources to offline or online.

Cluster Manager | BCC Manager | Cluster Event Log | Cluster Options

[Run Report](#)

Epoch: 5

[SYSTST-TS-94](#) [SYSTST-TS-97](#) [SYSTST-TS-149](#)

Cluster State						
Online Offline Migrate Respond to Alert Refresh						
6 Item(s)						
<input type="checkbox"/>	Type	Name	State	Location	Lives	Up Since
<input type="checkbox"/>		Master_IP_Address_Resource	Running	SYSTST-TS-149	3	Apr 19, 2012 2:41:34 PM
<input type="checkbox"/>		POOL_RAID10_SERVER	Running	SYSTST-TS-149	2	Apr 19, 2012 2:41:58 PM
<input checked="" type="checkbox"/>		POOL_SH_SERVER	Failover Alert	SYSTST-TS-94	1	

If scrolling is needed to view the complete list, you can use the *Refresh* option to set a longer refresh rate for the page to allow enough time to view the status for all items.







Master Server

The master server in the cluster is identified by a yellow diamond in the middle of the server icon (). The master server is initially the first server in the cluster, but another server can become the master if the first server fails.

Operating State

An icon next to the cluster resource or cluster indicates its operating state. [Table 9-1](#) describes each operating state and icon. When a resource is state is blank or has no colored icon, it is unassigned, offline, changing state, or in the process of loading or unloading.

Table 9-1 *Operating States for a Cluster and Its Resources*

State	Icon	Description
Normal		A green ball indicates that the node is running, or the resource is online.
Stopped		A red ball with a horizontal white line indicates that the node or resource is stopped and waiting for administrator intervention.
Offline		A white ball with a horizontal red line indicates that the node or resource is offline.
Critical		A white ball with a red X indicates that the node has failed, the resource is comatose, or a resource in a My Resources list cannot be not found (renamed or deleted).
Warning		A white ball with a yellow diamond indicates that an alert condition has occurred, and the resource needs administrator attention.
Unknown		A yellow ball with a question mark indicates that the state of the node or resource cannot be determined. Either the server is not currently a member of the cluster or its state is unknown.

Epoch

The Epoch number indicates the number of times the cluster state has changed. The cluster state changes every time a server joins or leaves the cluster.

Cluster Resource States


[Table 9-2](#) identifies the different resource states and gives descriptions and possible actions for each state.

Table 9-2 Cluster Resource States

Resource State	Description	Possible Actions
<p>Alert</p> <p>Alert types are:</p> <p>Failback Alert</p> <p>Failover Alert</p> <p>Start Alert</p>	<p>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.</p>	<p>On the Cluster Manager page in Novell iManager, select the check box next to the resource, click the <i>Respond to Alert</i> option, then resolve the alert. Depending on the resource state, you are prompted to start, fail over, or fail back the resource. For information, see Section 9.4, "Responding to a Start, Failover, or Failback Alert," on page 145.</p> <p>After you clear the <i>Start Alert</i>, you can take the resource offline, then bring the resource online.</p>
Comatose	The resource is not running properly and requires administrator intervention.	<p>Click the <i>Comatose</i> status indicator and take the resource offline. After resource problems have been resolved, you can bring the resource online (return it to the running state).</p> <p>A comatose resource cannot be taken directly to an online state.</p>
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.	If desired, you can click the <i>Offline</i> status indicator and 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.	You can 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 is not an assigned node available that the resource can be loaded on.	If desired, you can click the <i>Unassigned</i> status indicator and offline the resource. Offlining the resource prevents it from running on any of its preferred nodes if any of them join the cluster.

Resource State	Description	Possible Actions
Unloading	The resource is in the process of unloading from the server it was running on.	None.

9.4 Responding to a Start, Failover, or Failback Alert

You might receive an alert () for a resource if you have configured the resource's Start mode, Failover mode, or Failback mode settings to Manual. The alert types correspond to the mode setting:

- ♦ Start Alert
- ♦ Failover Alert
- ♦ Failback Alert

The resource is waiting for the administrator to intervene to start, fail over, or fail back the resource on the specified server.

If you attempt to offline a resource that is in the *Start Alert* state, nothing happens and the following message is displayed:


This operation cannot be completed. It is only available when the resource is in an offline state.

You must clear the *Start Alert* before you can offline the resource. After the resource is offline, you can online the resource.

To respond to an alert:





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

If the cluster you want to manage does not appear in your list, add the cluster to your list as described in [Section 8.2, "Setting Up a Personalized List of Clusters to Manage,"](#) on page 117.

- 2 In your personalized list of clusters, a critical () status for a cluster indicates that a resource might be in a comatose state and require administrator intervention. Click the resource name to manage the resource.

My Clusters

The list below provides a convenient way to keep track of clusters you frequently manage. Click on the add button to browse for clusters in the tree, and select them to be added to the list.

My Clusters							
		Action ▾		Refresh			
<input type="checkbox"/>	Status	Name	Master Node	IP Address	Nodes	Epoch	BCC
<input type="checkbox"/>		 cluster1	 SYSTST-TS-149	10.10.183.220	2/3	5	

- On the Cluster Manager page, select the check box next to the resource that has the alert, then click *Respond to Alert*.

My Clusters > cluster1.Servers.novell




cluster1.Servers.novell ?

View the status of both servers and resources in the cluster. You can also migrate cluster resources to different servers in the cluster, or change the state of resources to offline or online.

Cluster Manager BCC Manager Cluster Event Log Cluster Options

Run Report

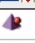



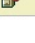
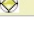
Epoch: 5

[SYSTST-TS-94](#) [SYSTST-TS-97](#) [SYSTST-TS-149](#)

Cluster State

Online | Offline | Migrate | Respond to Alert | Refresh 6 Item(s)

<input type="checkbox"/>	Type	Name	State	Location	Lives	Up Since
<input type="checkbox"/>		Master_IP_Address_Resource	 Running	SYSTST-TS-149	3	Apr 19, 2012 2:41:34 PM
<input type="checkbox"/>		POOL_RAID10_SERVER	 Running	SYSTST-TS-149	2	Apr 19, 2012 2:41:58 PM
<input checked="" type="checkbox"/>		POOL_SH_SERVER	 Failover Alert	SYSTST-TS-94	1	


- On the Respond to Alert page, click *Yes*, *No*, or *Cancel* to perform the desired action for the selected resource's alert state.

My Clusters > cluster1.Servers.novell > Respond to Alert

Respond to Alert ?

Click Yes, No, or Cancel to perform the desired action for the selected resources.

Selected Resources

Name	State	Yes Action	No Action	Cancel Action
POOL_SH_SERVER		Failover to next available node	Offline resource	Stay in Alert state

Do you want to continue?

- After clearing a *Start Alert*, go to the Cluster Manager page, select the check box next to the resource, then click *Offline*. After the resource is offline, select the check box again and click *Online* to bring the resource online.

9.5 Viewing the Cluster Event Log in iManager

The Cluster Event Log in Clusters-plugin for Novell iManager lets you view recent cluster events logged in the `/admin/Novell/Cluster/EventLog.xml` file. Events can be node specific (a node joins the cluster or leaves the cluster) or they can be resource specific (a resource comes online or goes offline).

By default, the Cluster Event Log displays events for the current date in chronological order for all severity levels, all nodes, and all resources. You can use the Cluster Event Log Filter page to filter out events that are not of interest so that only the desired log entries are displayed in the report. You can filter out events according to these categories:

- Severity (Error, Normal, Warning)
- Node (by node name)
- Resource (by resource name)

Selecting a check box causes that entry type to appear in the log. Deselecting a check box filters out the associated entry type.

In the *Event Log Range* section of the filter, you can choose to report all events in the log, or you can get a report of logged events that occurred since a specified date and time.

In the *Details* section of the filter, you can control how many event log messages are displayed per page. The default is 10 events per page. If the number of entries exceeds the one page, you can click the right-arrow button and left-arrow button to page through the report. Click the double-left-arrow button to jump to the first page of the report. Click the double-right-arrow button to jump to the last page of the report.

In the *Details* section of the filter, you can also choose whether the log message portion of a log entry is displayed in the log.

To view the logged cluster events:

- 1 In iManager, select *Clusters > My Clusters*, then select the cluster.
- 2 Select the *Cluster Event Log* to view the latest cluster events.

My Clusters > cl134.ncs.novell

cl134.ncs.novell

Use the Cluster Event Log to view recent cluster events. Events can be node specific (a node joined or left the cluster) or they can be resource specific (a resource changed state from loading to running).

Cluster Manager BCC Manager Cluster Event Log Cluster Options

Event Log

Filter | Clear Filter | Clear Log 10 Item(s)

Timestamp	Node	Resource	Severity	Cluster Event
May 30, 2013 11:52:02 AMMDT		LXVOLG_resource	Normal	Running on avalon
May 30, 2013 11:51:57 AMMDT		Master_IP_Address_Resource	Normal	Running on avalon
May 30, 2013 11:51:56 AMMDT		LXVOLF_resource	Normal	Running on avalon
May 30, 2013 11:51:55 AMMDT	lvm245		Normal	Running on avalon
May 30, 2013 11:51:53 AMMDT		LXVOLG_resource	Normal	Loading on avalon
May 30, 2013 11:51:53 AMMDT		POOLD_SERVER	Normal	Running on avalon
May 30, 2013 11:51:39 AMMDT		POOLD_SERVER	Normal	Loading on avalon
May 30, 2013 11:51:39 AMMDT	lvm245		Normal	Loading on avalon
May 30, 2013 11:51:39 AMMDT		LXVOLF_resource	Normal	Loading on avalon
May 30, 2013 11:51:39 AMMDT		Master_IP_Address_Resource	Normal	Loading on avalon

Print Close

3 (Optional) Generate a customized event report.

3a At the top of the Cluster Event Log page, click *Filter* to open the Event Log Filter page.

My Clusters > c1134.ncs.novell > Event Log Filter

Event Log Filter ?

Modify log attributes. Change the severity of logged events. Filter out nodes and resources.

Severity		Node	
<input checked="" type="checkbox"/> Types		<input checked="" type="checkbox"/> Name	
<input checked="" type="checkbox"/> Error		<input checked="" type="checkbox"/> avalon	
<input checked="" type="checkbox"/> Normal		<input checked="" type="checkbox"/> pemberly	
<input checked="" type="checkbox"/> Warning		<input checked="" type="checkbox"/> stratford	

Resource	
<input checked="" type="checkbox"/> Name	
<input checked="" type="checkbox"/> POOLD_SERVER	
<input checked="" type="checkbox"/> LXVOLF_resource	
<input checked="" type="checkbox"/> lvm245	
<input checked="" type="checkbox"/> LXVOLG_resource	

Event Log Range

☐ Entire Log

☒ Start Date and Time

Thu 30 May 2013 11:52:02 AM MDT

Details

Number of Messages to Display:

☒ Display Log Messages

OK Close

3b In the *Severity* section, deselect the check box next to message types that are not of interest.

- ♦ *Error*
- ♦ *Normal*
- ♦ *Warning*

For example, deselect *Normal* to omit normal events. Only error and warning messages are displayed.

3c In the *Node* section, deselect the check box next to nodes that are not of interest. Only events for selected nodes will appear in the report.

3d In the *Resource* section, deselect the check box next to cluster resources that are not of interest. Only events for selected resources will appear in the report.

3e In the *Event Log Range* section, select one of the following:

- ♦ Entire log
- ♦ Start Date and Time

3f In the *Details* section, specify how many messages to display per page and whether the log message portion of a log entry is displayed in the log.

3g Click *OK* to apply the filter.

3h View the customized report.

For example, the following report shows only error or warning messages for the POOLD_SERVER cluster resource that appear in the entire log.

My Clusters > cl134.ncs.novell

cl134.ncs.novell

Use the Cluster Event Log to view recent cluster events. Events can be node specific (a node joined or left the cluster) or they can be resource specific (a resource changed state from loading to running).

Cluster Manager BCC Manager Cluster Event Log Cluster Options

Event Log

Filter | Clear Filter | Clear Log 4 item(s)

Timestamp	Node	Resource	Severity	Cluster Event
April 11, 2013 1:41:53 PMMDT		POOLD_SERVER	Warning	Unassigned on unknown
April 11, 2013 1:41:52 PMMDT		POOLD_SERVER	Warning	NDS Sync on unknown
April 11, 2013 1:03:11 PMMDT		POOLD_SERVER	Warning	Unassigned on unknown
April 11, 2013 1:03:07 PMMDT		POOLD_SERVER	Warning	NDS Sync on unknown

Print Close

4 Click *Clear Filter* to clear the current filter.

5 (Optional) Click *Clear Log* to remove all entries from the log.

All prior log entries are deleted and are no longer available.

9.6 Viewing the Cluster Event Log at the Command Line

You can view cluster events logged in the /admin/Novell/Cluster/EventLog.xml file by copying the /admin/Novell/Cluster/EventLog.xml file to a working location, and using the cat command to display the file.

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

2 Enter

```
cp /admin/Novell/Cluster/EventLog.xml /tmp/
```

```
cat /tmp/EventLog.xml
```

This returns the log content in an XML format. For example:

```
<logEntry>
  <timeStamp> 5/12/2013 23:28:38</timeStamp>
  <node>CG-03</node>
  <clusterEvent>Joined at epoch 2</clusterEvent>
</logEntry>

<logEntry>
  <timeStamp> 5/12/2013 23:27:16</timeStamp>
  <resource>Master_IP_Address_Resource</resource>
  <clusterEvent>Running on CG-01</clusterEvent>
</logEntry>

<logEntry>
  <timeStamp> 5/12/2013 23:27:15</timeStamp>
  <node>CG-02</node>
  <clusterEvent>Joined at epoch 1</clusterEvent>
</logEntry>

<logEntry>
  <timeStamp> 5/12/2013 23:14:32</timeStamp>
  <resource>Master_IP_Address_Resource</resource>
```



```

    <clusterEvent>Running on CG-01</clusterEvent>
</logEntry>

<logEntry>
  <timeStamp> 5/12/2013 23:14:19</timeStamp>
  <resource>Master_IP_Address_Resource</resource>
  <clusterEvent>Loading on CG-01</clusterEvent>
</logEntry>

<logEntry>
  <timeStamp> 5/12/2013 23:14:13</timeStamp>
  <node>CG-01</node>
  <clusterEvent>Joined at epoch 0</clusterEvent>
</logEntry>

```

9.7 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 Monitor script Protocols Virtual server name CIFS server name IP address Advertising protocols

To generate a cluster configuration report:

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

If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, "Setting Up a Personalized List of Clusters to Manage,"](#) on page 117.

- 3 Click *Action > Run Report*.

When you select *Run Report*, the generated report opens in a new window, instead of in the child frame. This allows you to print and save the report from any Web browser.

- 4 Use the scroll bar to view the report in the browser.
- 5 (Optional) Use the browser options to print the report.
- 6 (Optional) Save the report to a text file.

Because some of the output is scripted, you must copy and paste the report content instead of using the browser's *File > Save as* option.

- 6a Click the cursor in the browser window where the report results are displayed, press Ctrl+A to select all, press Ctrl+C to copy the report to the clipboard.
- 6b In a text editor, open a new document, then press Ctrl+V to paste the report in the document.
- 6c Save the report document.

9.8 Cluster Migrating Resources to Different Nodes

You can migrate cluster resources to different servers in your cluster without waiting for a failure to occur. You might want to migrate resources to lessen the load on a specific server, to free up a server so it can be brought down for scheduled maintenance, or to increase the performance of the resource or application by putting it on a faster machine. The node you choose must be in the resource's preferred nodes list and available to run the resource.

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

Using iManager

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the cluster that you want to manage.

If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, "Setting Up a Personalized List of Clusters to Manage,"](#) on page 117.

- 3 On the Cluster Manager page, select the check box next to the resource you want to cluster migrate, then click *Migrate*.

A page appears, displaying a list of possible servers that you can migrate this resource to.

- 4 Select a server from the list to migrate the resource to, then click *OK* to migrate the resource to the selected server.

Using Console Commands

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

```
cluster migrate <resource_name> {<node_name> | -list | -most | -next}
```

Migrates the specified resource from the node where it is currently running to the node you specify in the command. The node you migrate the resource to must be running in the cluster and also be in the resource's *Assigned Nodes* list.

Options

<node_name>

Migrates the resource to the specified node if possible. No action is taken if the specified node is not in the resource's preferred nodes list, there is a Resource Mutual Exclusion conflict on the specified node, or the specified node is not currently active in the cluster.

-l, -list

Shows the preferred nodes list for the specified resource.

-m, -most

Migrates the specified resource to the most preferred node currently in the cluster. If the resource is running on such a node, no action is taken.

-n, -next

Migrates the resource to the node in the preferred node list that is next in order to the node where the resource is currently running. If there is no next node or all such nodes are not in the cluster currently, it searches from the beginning of the resource's preferred nodes list. No action is taken if a new destination could not be determined after the search is exhausted.

Examples

Move `res1` from the current node to `node2` in the resource's preferred nodes list. The resource is migrated only if there are no Resource Mutual Exclusion conflicts on the target node.

```
cluster migrate res1 node2
```

List the preferred nodes for resource `AUTO_POOL_14_SERVER`.

```
cluster migrate AUTO_POOL_14_SERVER -l
```

Preferred nodes for resource 'AUTO_POOL_14_SERVER'

```
CG-06  
CG-05  
CG-04  
CG-03  
CG-02  
CG-01  
CG-08  
CG-07
```

Status for Resource: `AUTO_POOL_14_SERVER`

Running on CG-05 Lives: 65

Revision: 5

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

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

- ♦ [Section 9.9.1, "Guidelines for Onlining and Offlining Resources," on page 153](#)
- ♦ [Section 9.9.2, "Using iManager to Online and Offline Resources," on page 153](#)
- ♦ [Section 9.9.3, "Using the Cluster Online Command," on page 153](#)
- ♦ [Section 9.9.4, "Using the Cluster Offline Command," on page 153](#)
- ♦ [Section 9.9.5, "Using the Cluster Migrate Command," on page 154](#)

9.9.1 Guidelines for Onlining and Offlining Resources

Onlining a resource runs the load script, which loads the resource on its primary preferred node, or on an alternate preferred node if possible, according to the order in its *Preferred Nodes* list. No action is taken if the preferred nodes are not active in the cluster, or if there are Resource Mutual Exclusion conflicts on its active preferred nodes.

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

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

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

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

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

9.9.2 Using iManager to Online and Offline Resources

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the cluster that you want to manage.
If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, “Setting Up a Personalized List of Clusters to Manage,”](#) on page 117.
- 3 On the Cluster Manager page, select the check box next to the resource you want to manage, then click *Online* or click *Offline*.

9.9.3 Using the Cluster Online Command

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

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

```
cluster online resource_name
```

To bring a resource online on a specific active node:

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

```
cluster online resource_name node_name
```

9.9.4 Using the Cluster Offline Command

To take a resource offline:

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

```
cluster offline resource_name
```


9.9.5 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 another node. If a node name is not provided, the resource is brought online according to its *Preferred Nodes* list if it is possible. You can alternatively specify a destination node from its *Preferred Nodes* list that is running in the cluster. It migrates the resource to the specified node if it is possible. No action is taken if the specified node is not in the resource's preferred nodes list, there is a Resource Mutual Exclusion conflict on the specified node, or the specified node is not currently active in the cluster.

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

```
cluster migrate resource_name [node_name]
```

9.10 Removing (Leaving) a Node from the Cluster

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

- ♦ [Section 9.10.1, “Removing One Node at a Time,” on page 154](#)
- ♦ [Section 9.10.2, “Removing All Nodes at the Same Time \(Taking the Cluster Down\),” on page 154](#)

9.10.1 Removing One Node at a Time

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

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

```
cluster leave
```

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

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

9.10.2 Removing All Nodes at the Same Time (Taking the Cluster Down)

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

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

```
cluster down
```

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

9.11 Joining a Node to the Cluster (Rejoining a Cluster)

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

While the node is not active in the cluster, any changes in storage objects are not updated automatically for the node. For example, the following storage object changes are automatically updated only for nodes that are currently active in the cluster:

- ♦ Creating, deleting, or expanding a clustered NSS pool
- ♦ Creating, deleting, or expanding a clustered LVM volume group
- ♦ Adding or removing segments in an NSS software RAID (that is used by a clustered NSS pool)
- ♦ Creating, mirroring, or re-creating an SBD partition
- ♦ Creating or deleting shared partitions

If Novell Cluster Services is configured to autostart when a server is restarted, a node that has left the cluster automatically rejoins the cluster on reboot. A server restart automatically rescans devices, rescans storage objects, and remaps multipath devices. Thus, if you reboot a node that has left the cluster, a node automatically recognizes any storage object changes made while the node was not active in the cluster.

To rejoin a cluster without rebooting the server:

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

```
cluster join
```

- 2 (Storage object changes) If you modified shared storage objects while the node was not active in the cluster, you must rescan the storage objects in order for the node to recognize these changes:

- 2a Log in to the node as the `root` user and open a terminal console.

- 2b Scan for storage object changes. At the command prompt, enter

```
nlvm rescan
```

9.12 Shutting Down Novell Cluster Services When Servicing Shared Storage

If you need to power down or recycle your shared storage system (such as when you apply firmware changes to a SAN), you should shut down Novell Cluster Services prior to doing so. One way to do this is to use the `cluster down` command on any of the nodes in the cluster. You are prompted to confirm the cluster down.

9.13 Enabling or Disabling Cluster Maintenance Mode

Cluster maintenance mode lets you temporarily suspend the cluster heartbeat while hardware maintenance is being performed. This is useful if you want to reset or power down the LAN switch without bringing down cluster servers.

Enabling the cluster maintenance mode from one cluster node puts the entire cluster in maintenance mode.

- 1 Log in as the root user to a node in the cluster, then enter the following at a command prompt:

```
cluster maintenance on
```

If the master server in the cluster is up, disabling the cluster maintenance mode from one cluster node brings the entire cluster out of maintenance mode. If the master server in the cluster goes down while the cluster is in cluster maintenance mode, you must disable cluster maintenance mode on all remaining cluster nodes in order to bring the cluster out of maintenance mode.

- 1 Log in as the root user to a node in the cluster, then enter the following at a command prompt:

```
cluster maintenance off
```

9.14 Viewing the Cluster Node Properties

You can view the cluster node number and IP address of the selected node as well as the distinguished name of the Linux Server object.

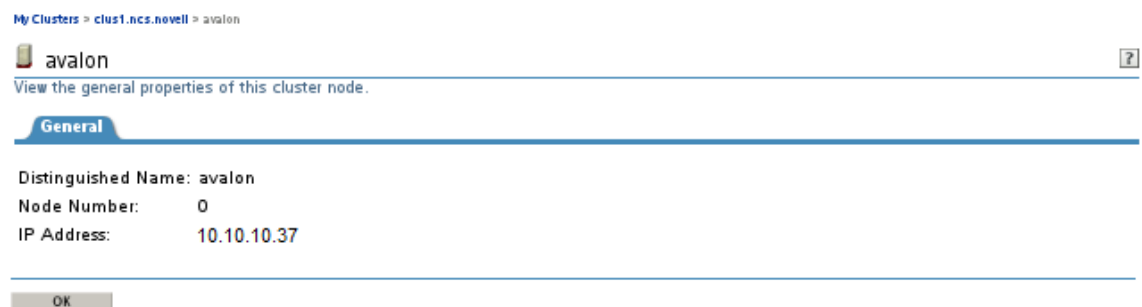
- 1 In iManager, select *Clusters > My Clusters*.

- 2 Click the cluster that you want to manage.

If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, “Setting Up a Personalized List of Clusters to Manage,”](#) on page 117.

- 3 Select *Cluster Options*.

- 4 Select the check box next to the cluster node whose properties you want to view, then click the *Details* link.



- 5 View the node information:

- ♦ **Distinguished Name:** The distinguished name of the Server object.
- ♦ **Node Number:** The cluster node number.
- ♦ **IP Address:** The IP address of the server.

- 6 Click OK to return to the Cluster Options page.

9.15 Deleting a Cluster Node from a Cluster

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

To 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 command prompt:

```
cluster leave
```

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

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

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

- 6 (Optional) Build a replacement server with the same node name and IP address, and add it to the cluster. For information, see [Section 5.6, "Adding a Node That Was Previously in the Cluster,"](#) on page 95.

9.16 Creating or Deleting Cluster SBD Partitions

If a single node (or group of nodes) somehow becomes isolated from other nodes, a condition called *split brain* results. Each side believes the other has failed, and forms its own cluster view that excludes the nodes it cannot see. Neither side is aware of the existence of the other. If the split brain is allowed to persist, each cluster will fail over the resources of the other. Since both clusters retain access to shared disks, corruption will occur when both clusters mount the same volumes.

Novell Cluster Services provides a split-brain detector (SBD) function to detect a split-brain condition and resolve it, thus preventing resources from being loaded concurrently on multiple nodes. The SBD partition contains information about the cluster, nodes, and resources that helps to resolve the split brain condition.

Novell Cluster Services requires an SBD partition for a cluster if its nodes use physically shared storage. Typically, you create the SBD when you configure the cluster on the first node. You can alternatively configure an SBD for the cluster after you configure the first node, but before you configure Novell Cluster Services on the second node of the cluster. You might also need to delete and re-create an SBD partition if the SBD becomes corrupted or its device fails.

An SBD must exist and the cluster must be enabled for shared disk access before you attempt to create shared storage objects in a cluster, such as pools and volumes. Novell Linux Volume Manager (NLVM) and other NSS management tools need the SBD to detect if a node is a member of the cluster and to get exclusive locks on physically shared storage.

This section describes how to use the Novell Cluster Services SBD Utility (`sbdutil`) to create and delete SBD partitions.

- ♦ [Section 9.16.1, “Requirements and Guidelines for Creating an SBD Partition,” on page 158](#)
- ♦ [Section 9.16.2, “Before You Create a Cluster SBD Partition,” on page 160](#)
- ♦ [Section 9.16.3, “Creating a Non-Mirrored Cluster SBD Partition with SBDUTIL,” on page 160](#)
- ♦ [Section 9.16.4, “Mirroring an Existing SBD Partition with NSSMU,” on page 164](#)
- ♦ [Section 9.16.5, “Creating a Mirrored Cluster SBD Partition with SBDUTIL,” on page 167](#)
- ♦ [Section 9.16.6, “Removing a Segment from a Mirrored Cluster SBD Partition,” on page 170](#)
- ♦ [Section 9.16.7, “Deleting a Non-Mirrored Cluster SBD Partition,” on page 172](#)
- ♦ [Section 9.16.8, “Deleting a Mirrored Cluster SBD,” on page 173](#)
- ♦ [Section 9.16.9, “Additional Information about SBD Partitions,” on page 174](#)

9.16.1 Requirements and Guidelines for Creating an SBD Partition

Consider the requirements and guidelines in this section when you create a Novell Cluster Services SBD (split-brain detector) partition for an existing cluster.

IMPORTANT: Check to see if a cluster SBD partition already exists before you create a new one. See [Section 9.16.2, “Before You Create a Cluster SBD Partition,” on page 160.](#)

- ♦ [“Preparing Novell Cluster Services” on page 159](#)
- ♦ [“Using a Shared Disk System” on page 159](#)
- ♦ [“Preparing a SAN Device for the SBD” on page 159](#)
- ♦ [“Initializing and Sharing a Device for the SBD” on page 159](#)
- ♦ [“Working with NLVM Commands in a Cluster” on page 159](#)
- ♦ [“Determining the SBD Partition Size” on page 160](#)
- ♦ [“Replacing an Existing SBD Partition” on page 160](#)

Preparing Novell Cluster Services

Before you create an SBD partition for an existing cluster, you must take the cluster down and stop Novell Cluster Services software on all nodes. Do not restart Novell Cluster Services and rejoin nodes to the cluster until after you create the new SBD and configure the Shared Disks flag attribute for the Cluster object.

You can mirror an existing SBD while the cluster is up and running.

Using a Shared Disk System

You must have a shared disk system (such as a Fibre Channel SAN or an iSCSI SAN) connected to your cluster nodes before you create a split-brain-detector (SBD) partition. For information, see [Section 4.8, “Shared Disk Configuration Requirements,” on page 66](#) for more information.

Preparing a SAN Device for the SBD

Use the SAN storage array software to carve a LUN to use exclusively for the SBD partition. The device should have at least 20 MB of free available space. Connect the LUN device to all nodes in the cluster.

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

The device you use to create the SBD partition must not be a software RAID device. A hardware RAID configured in a SAN array is seen as a regular device by the server.

If you attach new devices to the server while it is running, you should scan for new devices on each cluster node to ensure that the devices are recognized by all nodes. Log in as the `root` user, launch a terminal console, then enter

```
nlvm -s rescan
```

Initializing and Sharing a Device for the SBD

Before you use `sbdutil` to create an SBD, you must initialize each of the SAN devices that you created for the SBD and mark each device as Shareable for Clustering. When you mark the device as Shareable for Clustering, share information is added to the disk in a free-space partition that is about 4 MB in size. This space becomes part of the SBD partition.

When the cluster is down and Novell Cluster Services is stopped, you can use NSSMU, the Storage plug-in for iManager, the `nlvm init` command, or an NSS utility called `ncsinit` to initialize a device and set it to a shared state. To minimize the risk of possible corruption, you are responsible for ensuring that you have exclusive access to the shared storage at this time.

Working with NLVM Commands in a Cluster

If you have physically shared storage and the SBD does not exist, NSS management tools cannot detect if the node is a member of the cluster and cannot get exclusive locks to the physically shared storage. In this state, you can use the `-s` NLVM option with NLVM commands to override the shared locking requirement and force NLVM to execute the commands. To minimize the risk of possible corruption, you are responsible for ensuring that you have exclusive access to the shared storage at this time.

Determining the SBD Partition Size

You use the `size` option with the SBD Utility to specify the desired size of the SBD partition. You can specify how much free space to use for the SBD, or you can specify the `-1` option to use the entire device (the maximum size). If you specify a device to use as a mirror, the same amount of space is used. If you specify to use the maximum size and the mirror device is bigger than the SBD device, you will not be able to use the excess free space on the mirror for other purposes.

Because an SBD partition ends on a cylinder boundary, the partition size might be slightly smaller than the size you specify. When you use an entire device for the SBD partition, you can use the `-1` option as the size, and let the software determine the size of the partition.

Replacing an Existing SBD Partition

To replace an existing SBD partition, you must first delete the old SBD partition, and then create the new one. To reuse the SBD partition's device, you must remove the SBD partition, then re-initialize and share the device.

You must take the cluster down and stop Novell Cluster Services on all nodes before you delete the existing SBD partition. Do not restart Novell Cluster Services and rejoin nodes to the cluster until after you create the new SBD.

9.16.2 Before You Create a Cluster SBD Partition

Before you create a Novell Cluster Services SBD partition, you should ensure that an SBD does not already exist on your cluster. The new SBD partition is not recognized until the old partition is deleted.

- 1 Log in as the `root` user to any node in the cluster, and launch a terminal console.
- 2 At the command prompt, enter

```
sbdutil -f
```

The `-f` option tells you whether an SBD partition exists for the cluster, and identifies the SAN device where the SBD partition is located. It returns `NotFound` if the SBD partition does not exist.

- 3 If an SBD partition already exists, use one of the following methods to delete the existing SBD partition before attempting to create another one:
 - ♦ [Section 9.16.7, "Deleting a Non-Mirrored Cluster SBD Partition," on page 172](#)
 - ♦ [Section 9.16.8, "Deleting a Mirrored Cluster SBD," on page 173](#)

9.16.3 Creating a Non-Mirrored Cluster SBD Partition with SBDUTIL

If you did not create a cluster SBD partition during the Novell Cluster Services installation on the first node of the cluster, you can create it later by using the SBDUTIL utility (`/opt/novell/ncs/bin/sbdutil`). You might also need to delete and re-create an SBD partition if the SBD becomes corrupted or its device fails. See the man page for `sbdutil` for more information on how to use it.

If a cluster partition does not exist, create one by using the SBDUTIL:

- 1 Ensure that nobody else is changing any storage on any nodes at this time.
Until the SBD exists and the cluster is set up for shared disk access, you are responsible for ensuring that you have exclusive access to the shared storage.

2 Take the cluster down:

2a Log in to any node in the cluster as the `root` user, then open a terminal console.

2b At the command prompt, enter

```
cluster down
```

3 On each cluster node, stop Novell Cluster Services:

3a Log in to the cluster node as the `root` user, then open a terminal console.

3b At the command prompt, enter

```
rcnovell-ncs stop
```

3c After you have stopped Novell Cluster Services on all nodes, continue with the next step.

4 Prepare a SAN device to use for the SBD partition:

4a Use the SAN storage array software to carve a device to use exclusively for the SBD partition.

4b Attach the device to all nodes in the cluster.

4c On each node, log in as the `root` user and rescan for devices:

```
nlvm -s rescan
```

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

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

6 View a list of the devices and identify the leaf node name (such as `sdc`) of the SAN device that you want to use for the SBD partition. At the command prompt, enter

```
nlvm -s list devices
```

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

The device information shows the leaf node name, the size, the amount of free available space, the partitioning format (such as MSDOS or GPT), the shared state (whether it is marked as Shareable for Clustering), and the RAID state (whether the device is an NSS software RAID device). Do not use an NSS software RAID for the device.

For example, an uninitialized device reports a format of `None` and a shared state of `No`:

```
sdc size=102.00MB free=0KB format=None shared=No RAID=No
```

7 Initialize and share the device. At the command prompt, enter

```
nlvm -s init <device_name> format=msdos shared
```

WARNING: Initializing a device destroys all data on the device.

Replace *device_name* with the leaf node name (such as `sdc`) of the SAN device you want to use as the SBD partition.

Specify a partitioning format of `msdos`.

Specify the `shared` option to mark the device as Shareable for Clustering.

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

You can list the devices to visually verify that the device is formatted and shared:

```
nlvm -s list devices
```


For example, the formatted device `sdc` reports a format of MSDOS and a shared state of Yes:

```
sdc size=102.00MB free=101.98MB format=MSDOS shared=Yes RAID=No
```

- 8** As the `root` user, enter the following at the command prompt:

```
sbdutil -c -n <cluster_name> -d <device_name> -s <size>
```

For the `-n` option, replace *cluster_name* with the name of the cluster, such as `cluster1`. This name must match the name of an existing cluster that has a Cluster object in Novell eDirectory. The name is case sensitive.

For the `-d` option, replace *device_name* with the leaf node name of the device where you want to create the cluster partition, such as `sdc`.

For the `-s` option, use one of the following methods for specifying the size of the SBD partition. For information about size requirements, see [“Preparing a SAN Device for the SBD” on page 159](#).

- ♦ **Specify a value:** Replace *size* with the size (in MB) to use for the SBD partition.

For example, the following command creates the `/dev/nss/mycluster1.sbd` partition with a size of 200 MB:

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

- ♦ **Specify -1:** You can specify the size as `-1` to use all free space on the device. This option allows 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/nss/cl1.sbd` partition on the CX4-LUN000 device, and uses the entire device:

```
sbdutil -c -n cl1 -d CX4-LUN000 -s -1
```

- ♦ **Use default size:** If the `-s` option is not used, the default size is 8 MB.

For example, the following command creates the `/dev/nss/cl2.sbd` partition on the `sdd` device with the default size of 8 MB:

```
sbdutil -c -n cl2 -d sdd
```

- 9** (Optional) You can use the `sbdutil -f -s` command to view the path of the SBD partition.

- 10** Modify the Cluster object in eDirectory to enable its NCS: Shared Disk Flag attribute.

This step is required only if the cluster has never had an SBD partition. However, it does no harm to verify that the NCS: *Shared Disk Flag* attribute is enabled.


- 10a** 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.

- 10b** Select *Directory Administration*, then select *Modify Object*.

- 10c** Browse to locate and select the Cluster object of the cluster you want to manage, then click **OK**.

- 10d** Under *Valued Attributes*, select the NCS: Shared Disk Flag, then click *Edit*.



- 10e** Select (enable) the *NCS: Shared Disk Flag* check box, then click *OK*.



- 10f** Click *Apply* to save changes.
- 11** On each cluster node, start Novell Cluster Services:
- 11a** Log in to the cluster node as the root user, then open a terminal console.
- 11b** At the command prompt, enter
- ```
rcnovell-ncs start
```
- 11c** After you have restarted Novell Cluster Services on all nodes, continue with the next step.
- 12** On each cluster node, join the cluster. At the command prompt, enter
- ```
cluster join
```
- 13** (Optional) Continue with [Section 9.16.4, "Mirroring an Existing SBD Partition with NSSMU,"](#) on [page 164](#).

9.16.4 Mirroring an Existing SBD Partition with NSSMU

You can mirror an existing Novell Cluster Services SBD partition to provide device fault tolerance. It is not necessary to take the cluster down or stop the cluster software.

This section describes how to use NSSMU to mirror the SBD partition. For information about using Novell Linux Volume Manager (NLVM) commands to mirror an SBD, see [“Mirroring an Existing SBD Partition with NLVM”](#) in the *OES 11 SP1: NLVM Reference*.

1 Prepare a SAN device to use as the mirror segment for the SBD partition:

- 1a** Use the SAN storage array software to carve a device that is at least the size of the existing SBD partition's device.
- 1b** Attach the device to all nodes in the cluster.
- 1c** On each node, log in as the root user and rescan for devices:

```
nlvm rescan
```

2 Initialize and mark the device as *Shareable for Clustering*:

- 2a** Log in to any node of the cluster as the root user, then open a terminal console.
- 2b** Launch NSSMU:

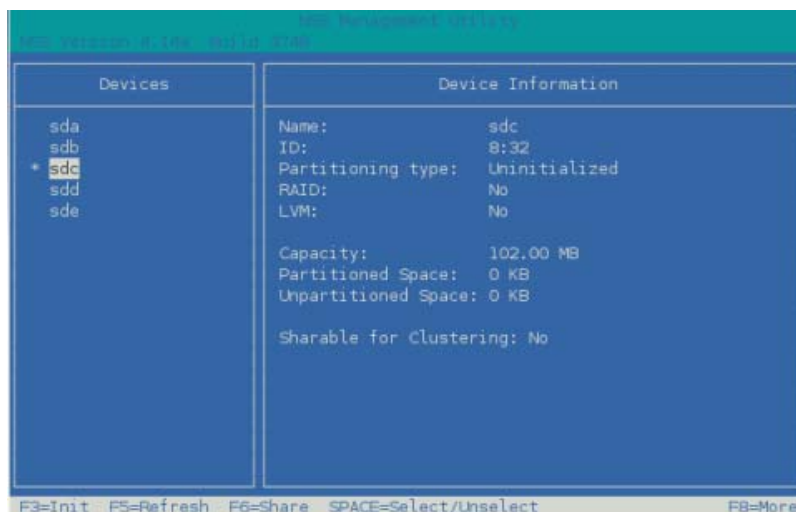
```
nssmu
```

- 2c** In the NSSMU Main Menu, select *Devices* and press Enter.




- 2d** In the *Devices* list, select the device that you want to use for the SBD mirror, press F3 to initialize the device, then press y (Yes) to confirm and continue.

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

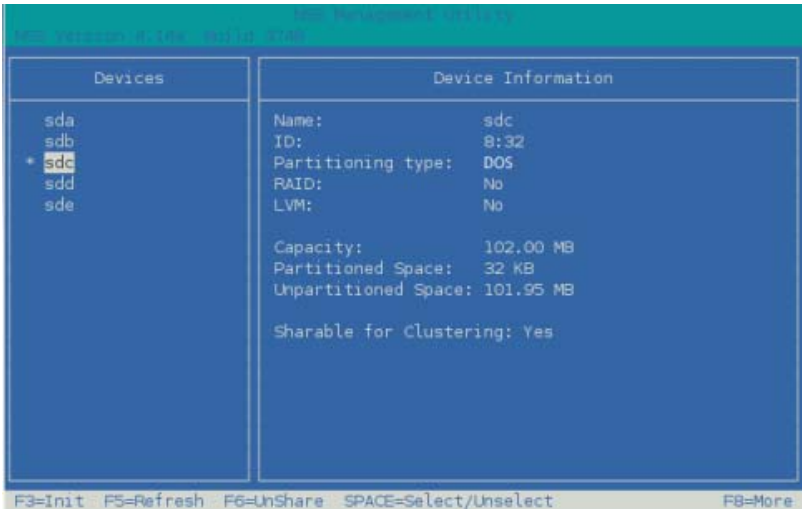


- 2e Select the DOS partitioning scheme for the device, then press Enter.
DOS supports devices up to 2 TB in size. GPT supports devices of any size.

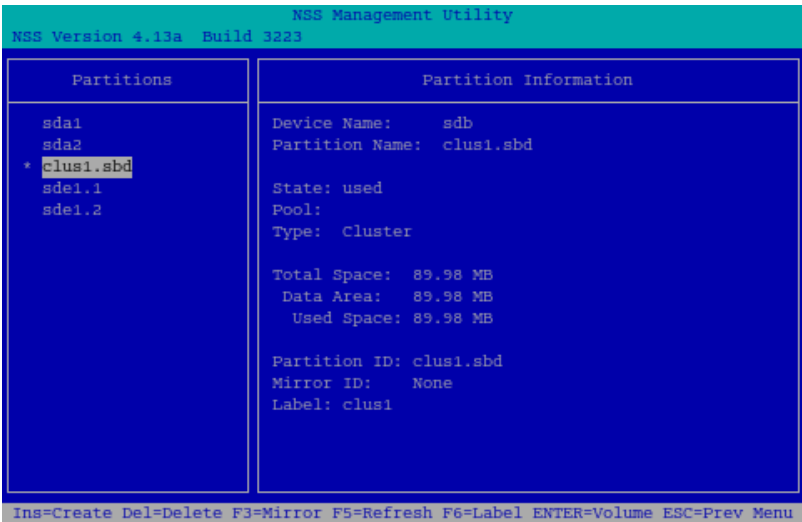


Wait for the page to refresh before continuing.

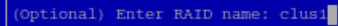
- 2f Press F6 to mark the device as shareable for clustering.
The *Shareable for Clustering* value changes from *No* to *Yes*.



- 2g Press Esc to return to the NSSMU Main Menu.
- 3 In NSSMU, mirror the SBD partition:
- 3a Select *Partitions* from the NSSMU Main Menu.
- 3b From the list of partitions, select the SBD partition that you want to mirror. The partition is named the same as the cluster with an *.sbd* extension, such as *clus1.sbd*.



- 3c** Press F3 to open the dialog to create the RAID1 mirror.
- 3d** Type the cluster name (such as `clus1`) for the RAID, then press Enter.
This is the same name as the SBD partition but without the `.sbd` extension.



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

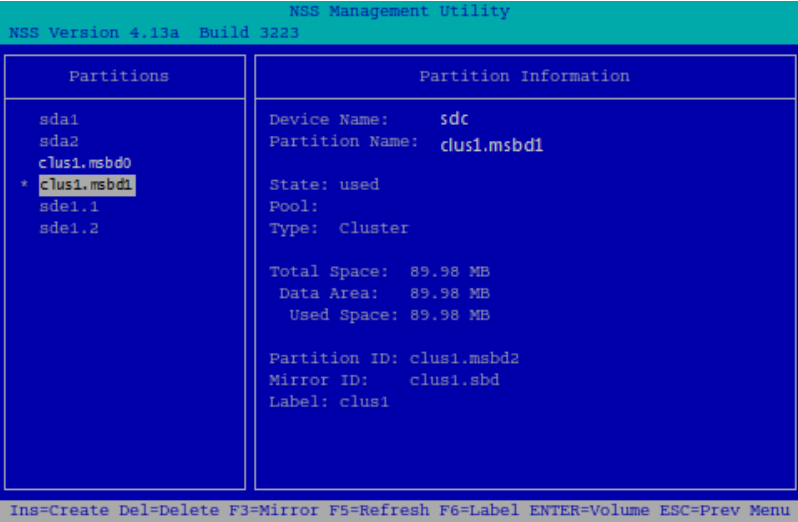
PartitionID	Size	Shared
* <code>sdc1.nwfreospace1</code>	101.98 MB	Yes
<code>sdd1.nwfreospace1</code>	511.98 MB	Yes

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

- 3f** Press F3 again to accept your selection and create the mirror.
- 3g** In the confirmation message, press `y` (Yes) to approve the RAID1 creation.

The SBD partitions are renamed with the partition name, and now have the extension of `.msbd0` (mirrored SBD partition 0) and `.msbd2` (mirrored SBD partition 1).

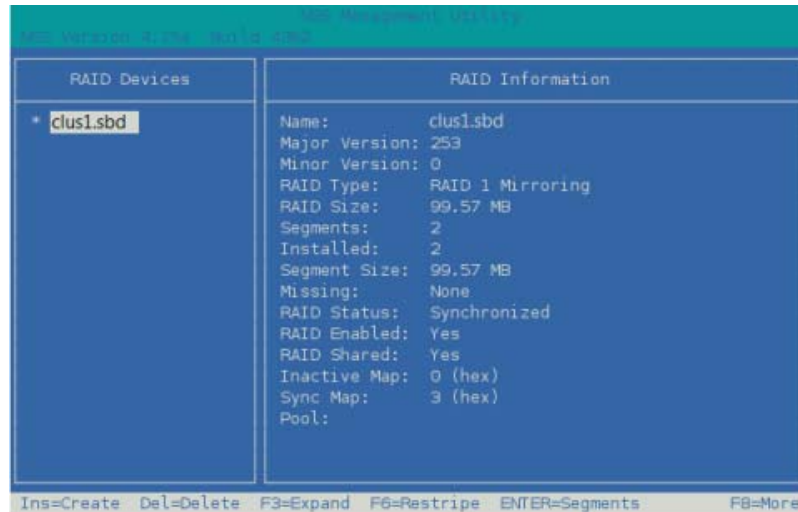
For example, for a cluster named `clus1`, the mirrored SBD partitions are `clus1.msbd0` and `clus1.msbd1`.



NSS Management Utility	
NSS Version 4.13a Build 3223	
Partitions	Partition Information
sda1	Device Name: sdc
sda2	Partition Name: clus1.msbd1
clus1.msbd0	State: used
* clus1.msbd1	Pool:
sde1.1	Type: Cluster
sde1.2	Total Space: 89.98 MB
	Data Area: 89.98 MB
	Used Space: 89.98 MB
	Partition ID: clus1.msbd2
	Mirror ID: clus1.sbd
	Label: clus1
Ins=Create Del=Delete F3=Mirror F5=Refresh F6=Label ENTER=Volume ESC=Prev Menu	

- 3h** Press Esc to return to the NSSMU Main Menu.

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



- 4c View the RAID Status to ensure that synchronization has begun.
Synchronization is complete when the status is *Synchronized*.
- 5 Press Esc twice to exit NSSMU.

9.16.5 Creating a Mirrored Cluster SBD Partition with SBDUTIL

You can create a mirrored Novell Cluster Services SBD partition to provide device fault tolerance for the SBD. You must take the cluster down and stop the cluster software.

Use the procedure in this section to create a new mirrored SBD partition by using the `sbdutil` utility. For information about using Novell Linux Volume Manager (NLVM) commands to create a new mirrored SBD partition, see [“Creating a Mirrored SBD Partition with NLVM”](#) in the *OES 11 SP1: NLVM Reference*.

- 1 Ensure that nobody else is changing any storage on any nodes at this time.
Until the SBD exists and the cluster is set up for shared disk access, you are responsible for ensuring that you have exclusive access to the shared storage.
- 2 Take the cluster down:
 - 2a Log in to any node in the cluster as the `root` user, then open a terminal console.
 - 2b At the command prompt, enter


```
cluster down
```
- 3 On each cluster node, stop Novell Cluster Services:
 - 3a Log in to the cluster node as the `root` user, then open a terminal console.
 - 3b At the command prompt, enter


```
rcnovell-ncs stop
```
 - 3c After you have stopped Novell Cluster Services on all nodes, continue with the next step.

- 4 Prepare two SAN devices to use for the SBD partition:
 - 4a Use the SAN storage array software to carve two devices of equal size to use exclusively for the mirrored SBD partition.
 - 4b Attach the devices to all nodes in the cluster.
 - 4c On each node, log in as the `root` user and rescan for devices:

```
nlvm -s rescan
```

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

- 5 Log in to any node in the cluster as the `root` user, then open a terminal console.
- 6 View a list of the devices and identify the leaf node name (such as `sdc`) of the two SAN devices that you want to use for the mirrored SBD partition. At the command prompt, enter

```
nlvm -s list devices
```

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

The device information shows the leaf node name, the size, the amount of free available space, the partitioning format (such as MSDOS or GPT), the shared state (whether it is marked as Shareable for Clustering), and the RAID state (whether the device is an NSS software RAID device). Do not use an NSS software RAID for the device.

For example, uninitialized devices report a format of `None` and a shared state of `No`:

```
sdc size=102.00MB free=0KB format=None shared=No RAID=No
sdd size=102.00MB free=0KB format=None shared=No RAID=No
```

- 7 Initialize and share the two devices. At the command prompt, enter

```
nlvm -s init <device_name1>,<device_name2> format=msdos shared
```

WARNING: Initializing a device destroys all data on the device.

Replace *device_name1* and *device_name2* with the leaf node names (such as `sdc` and `sdd`) of the two SAN devices you want to use for the mirrored SBD partition.

Specify a partitioning format of `msdos`.

Specify the `shared` option to mark the devices as Shareable for Clustering.

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

For example, to initialize devices `sdc` and `sdd`, enter

```
nlvm -s init sdc,sdd format=msdos shared
```

You can list the devices to visually verify that the device is formatted and shared:

```
nlvm -s list devices
```

For example, the formatted devices `sdc` and `sdd` report a format of `MSDOS` and a shared state of `Yes`:

```
sdc size=102.00MB free=101.98MB format=MSDOS shared=Yes RAID=No
sdd size=102.00MB free=101.98MB format=MSDOS shared=Yes RAID=No
```

- 8 Create a mirrored SBD partition. At the command prompt, enter

```
sbdutil -c -n <cluster_name> -d <device_name1> -d <device_name2> -s <size>
```


Replace *cluster_name* with the name of the cluster, such as `cluster1`. This name must match the name of an existing cluster that has a Cluster object in Novell eDirectory. The name is case sensitive.

Replace *device_name1* and *device_name2* with the leaf node names (such as `sdc` and `sdd`) of the two SAN devices you want to use for the mirrored SBD partition. The `cluster1.msbd0` mirrored SBD partition is created on the first device option instance in the command. The `cluster1.msbd1` mirrored SBD partition is created on the second device option instance in the command.

For the `-s` option, replace *size* with the size (in MB) to use for each of the SBD RAID1 segments. Specify the size only once. Both devices should be the same size, but if they are not, the size of the RAID segments is determined by the size of the smaller device. Use one of the following methods for specifying the size of the SBD partition. For information about size requirements, see [“Preparing a SAN Device for the SBD” on page 159](#).

- ♦ **Specify a value:** Replace *size* with the size (in MB) to use for the SBD partition.

For example, the following command creates the `cl1.sbd` mirrored RAID device and the `cl1.msbd0` and `cl1.msbd1` partitions. Each segment is up to 1020 MB in size, depending on where the last cylinder boundary lies.

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

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

For example, the following command uses devices `sdc` and `sdd` to create the `cluster1.sbd` mirrored RAID device and the SBD mirrored partitions `cluster1.msbd0` and `cluster1.msbd1`. It uses all available free space on the smaller device as the segment size.

```
sbdutil -c -n cluster1 -d sdc -d sdd -s -1
```

- ♦ **Use default size:** If the `-s` option is not used, the default size is 8 MB.

- 9 (Optional) You can use the `sbdutil -f -s` command to verify the path and name of the SBD RAID device.

You can also use NSSMU, the Storage plug-in to Novell iManager, or the `nlvm -s list partitions` command to view the partitions used by the RAID device.

- 10 Modify the Cluster object in eDirectory to enable its NCS: Shared Disk Flag attribute.

This step is required only if the cluster has never had an SBD partition. However, it does no harm to verify that the *NCS: Shared Disk Flag* attribute is enabled.


- 10a 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.

- 10b Select *Directory Administration*, then select *Modify Object*.

- 10c Browse to locate and select the Cluster object of the cluster you want to manage, then click OK.

- 10d** Under *Valued Attributes*, select the NCS: Shared Disk Flag, then click *Edit*.



- 10e** Select (enable) the *NCS: Shared Disk Flag* check box, then click *OK*.



- 10f** Click *Apply* to save changes.
- 11** On each cluster node, start Novell Cluster Services:
- 11a** Log in to the cluster node as the *root* user, then open a terminal console.
- 11b** At the command prompt, enter
- ```
rcnovell-ncs start
```
- 11c** After you have restarted Novell Cluster Services on all nodes, continue with the next step.
- 12** On each cluster node, join the cluster. At the command prompt, enter
- ```
cluster join
```

9.16.6 Removing a Segment from a Mirrored Cluster SBD Partition

You can remove a segment from a mirrored cluster SBD partition and keep the remaining SBD partition. The software RAID definition remains, so if you delete the remaining partition later, you must [delete the software RAID](#) instead of [simply deleting the partition](#) as with a stand-alone SBD partition.

IMPORTANT: You can use NLVM commands to completely unmirror the SBD RAID1 device, and keep the SBD partition. For information, see “[Unmirroring a Mirrored SBD Partition with NLVM](#)” in the [OES 11 SP1: NLVM Reference](#).

You can specify which segment to keep when you use NSSMU to remove a segment from a software RAID 1 (mirror) device.

- 1 Log in to any node as the root user, then launch a terminal console.
- 2 At the command prompt, enter

`cluster maintenance on`

This causes all cluster servers to enter maintenance mode.

- 3 Remove a segment from the SBD RAID1 device.

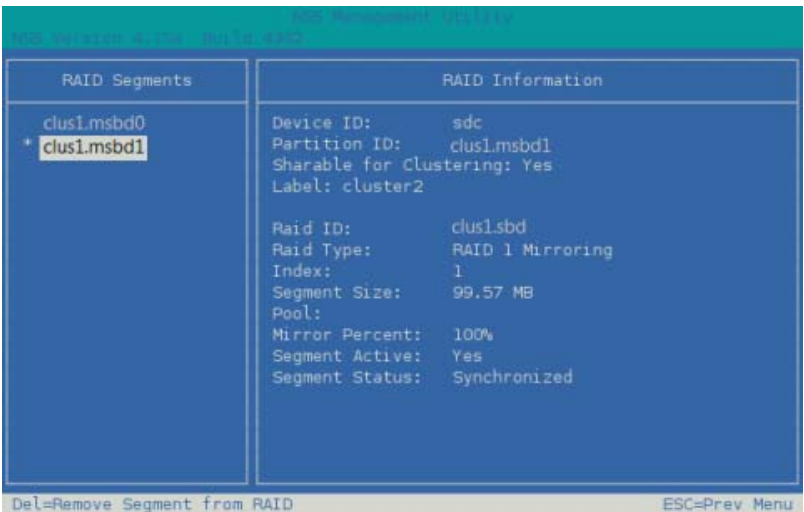
3a Launch NSSMU:

`nssmu`

3b In the NSSMU Main Menu, select *RAID Devices*.

3c Select the software RAID1 device for the cluster SBD that you want to manage.

3d Press Enter to show its member segments.



- 3e Select the member segment you want to delete, press Delete to remove the RAID segment, then press y (Yes) to confirm the removal.
- 3f Press Esc to return to the Software RAIDs page.

The RAID definition remains for the remaining segment of the mirrored SBD partition. The SBD RAID1 reports that it has one segment.

RAID Management Utility Novell Version 4.17a Build 4000	
RAID Devices	RAID Information
* clus1sbd	Name: clus1sbd Major Version: 253 Minor Version: 0 RAID Type: RAID 1 Mirroring RAID Size: 99.57 MB Segments: 1 Installed: 1 Segment Size: 99.57 MB Missing: None RAID Status: Synchronized RAID Enabled: Yes RAID Shared: Yes Inactive Map: 0 (hex) Sync Map: 1 (hex) Pool:
Ins=Create Del=Delete F3=Expand F6=Restripe ENTER=Segments F8=More	

3g Press Esc twice to exit NSSMU.

4 At the terminal console of one cluster server, enter

```
cluster maintenance off
```

This causes all cluster servers to return to normal mode.

9.16.7 Deleting a Non-Mirrored Cluster SBD Partition

You might need to delete and re-create a Novell Cluster Services SBD partition if the SBD becomes corrupted or its device fails. Use the procedure in this section to delete the SBD partition, then create a new SBD partition as described in [Section 9.16.3, “Creating a Non-Mirrored Cluster SBD Partition with SBDUTIL,” on page 160](#) or [Section 9.16.5, “Creating a Mirrored Cluster SBD Partition with SBDUTIL,” on page 167](#).

IMPORTANT: You must take the cluster down and stop Novell Cluster Services on all nodes before you delete the existing SBD partition. Do not restart Novell Cluster Services and rejoin nodes to the cluster until after you create a new SBD.

1 Ensure that nobody else is changing any storage on any nodes at this time.

Until the SBD exists and the cluster is set up for shared disk access, you are responsible for ensuring that you have exclusive access to the shared storage.

2 Take the cluster down:

2a Log in to any node in the cluster as the `root` user, then open a terminal console.

2b At the command prompt, enter

```
cluster down
```

3 On each cluster node, stop Novell Cluster Services:

3a Log in to the cluster node as the `root` user, then open a terminal console.

3b At the command prompt, enter

```
rcnovell-ncs stop
```

3c After you have stopped Novell Cluster Services on all nodes, continue with the next step.

- 4 Delete the SBD partition.
 - 4a Log in to any node in the cluster as the `root` user, then launch a terminal console.
 - 4b Launch NSSMU:


```
nssmu
```
 - 4c In the NSSMU Main Menu, select *Partitions*.
 - 4d Select the SBD partition you want to delete, such as `clus1.sbd`.
 - 4e Press Delete to delete the partition, then press `y` (*Yes*) to confirm the deletion.
 - 4f Press Esc twice to exit NSSMU.
- 5 Use one of the following methods to create a new SBD partition:
 - ♦ [Section 9.16.3, “Creating a Non-Mirrored Cluster SBD Partition with SBDUTIL,” on page 160](#)
 - ♦ [Section 9.16.5, “Creating a Mirrored Cluster SBD Partition with SBDUTIL,” on page 167](#)

Do not restart Novell Cluster services and rejoin nodes to the cluster until after you create the new SBD.

9.16.8 Deleting a Mirrored Cluster SBD

You might need to delete and re-create a Novell Cluster Services SBD partition if the SBD becomes corrupted or its device fails. Use the procedure in this section to delete a mirrored SBD partition, then create a new SBD partition as described in [Section 9.16.3, “Creating a Non-Mirrored Cluster SBD Partition with SBDUTIL,” on page 160](#) or [Section 9.16.5, “Creating a Mirrored Cluster SBD Partition with SBDUTIL,” on page 167](#).

IMPORTANT: You must take the cluster down and stop Novell Cluster Services on all nodes before you delete the existing mirrored SBD RAID1 device and its mirrored SBD partitions. Do not restart Novell Cluster Services and rejoin nodes to the cluster until after you create a new SBD.

- 1 Ensure that nobody else is changing any storage on any nodes at this time.

Until the SBD exists and the cluster is set up for shared disk access, you are responsible for ensuring that you have exclusive access to the shared storage.
- 2 Take the cluster down:
 - 2a Log in to any node in the cluster as the `root` user, then open a terminal console.
 - 2b At the command prompt, enter


```
cluster down
```
- 3 On each cluster node, stop Novell Cluster Services:
 - 3a Log in to the cluster node as the `root` user, then open a terminal console.
 - 3b At the command prompt, enter


```
rcnovell-ncs stop
```
 - 3c After you have stopped Novell Cluster Services on all nodes, continue with the next step.
- 4 Delete the mirrored software RAID that you used for the SBD partition:
 - 4a Log in to any node in the cluster as the `root` user, then launch a terminal console.
 - 4b Launch NSSMU:


```
nssmu
```


- 4c In the NSSMU Main Menu, select *RAID Devices*.
 - 4d Select the software RAID1 for the mirrored SBD you want to delete.
 - 4e Press Delete to delete the software RAID1 and its two mirrored SBD partitions, then press y (Yes) to confirm the deletion.
 - 4f Press Esc twice to exit NSSMU.
- 5 Use one of the following methods to create a new SBD partition:
- ♦ [Section 9.16.3, “Creating a Non-Mirrored Cluster SBD Partition with SBDUTIL,” on page 160](#)
 - ♦ [Section 9.16.5, “Creating a Mirrored Cluster SBD Partition with SBDUTIL,” on page 167](#)
- Do not restart Novell Cluster services and rejoin nodes to the cluster until after you create the new SBD.

9.16.9 Additional Information about SBD Partitions

See the following resources for additional information about SBD partitions:

- ♦ To understand general requirements for SBD partitions, see [Section 4.8.2, “SBD Partitions,” on page 67](#).
- ♦ To configure the SBD when you configure the cluster on the first node, see [Section 5.5.5, “Configuring a New Cluster,” on page 86](#).
- ♦ To use Novell Linux Volume Manager (NLVM) commands to create an SBD partition in a cluster, see [“Creating or Mirroring an SBD Partition” in the OES 11 SP1: NLVM Reference](#).
- ♦ To use NLVM commands to delete an SBD partition in a cluster, see [“Deleting an SBD Partition with NLVM” in the OES 11 SP1: NLVM Reference](#).
- ♦ For information about how the split brain detector works, see [NetWare Cluster Services: The Gory Details of Heartbeats, Split Brains, and Poison Pills \(TID 10053882\)](#) (<http://support.novell.com/docs/Tids/Solutions/10053882.html>).

9.17 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 9-3](#) lists the cluster-related virtual XML files and gives a brief description of each.

Table 9-3 *Cluster-Related Virtual XML Files*

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

[Table 9-4](#) lists the cluster-related CMD files and gives a brief description of each.

Table 9-4 *Cluster-Related CMD Files*

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

10 Configuring and Managing Cluster Resources

After you create and configure a 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 9, “Managing Clusters,”](#) on page 139.

- ♦ [Section 10.1, “Planning for Cluster Resources,”](#) on page 177
- ♦ [Section 10.2, “Setting Up a Personalized List of Resources to Manage,”](#) on page 179
- ♦ [Section 10.3, “Using Cluster Resource Templates,”](#) on page 184
- ♦ [Section 10.4, “Creating Cluster Resources,”](#) on page 189
- ♦ [Section 10.5, “Configuring a Load Script for a Cluster Resource,”](#) on page 190
- ♦ [Section 10.6, “Configuring an Unload Script for a Cluster Resource,”](#) on page 191
- ♦ [Section 10.7, “Enabling Monitoring and Configuring the Monitor Script,”](#) on page 192
- ♦ [Section 10.8, “Configuring the Start, Failover, and Failback Modes for Cluster Resources,”](#) on page 197
- ♦ [Section 10.9, “Configuring Preferred Nodes and Failover Order for a Resource,”](#) on page 198
- ♦ [Section 10.10, “Configuring Resource Priorities for Load Order,”](#) on page 200
- ♦ [Section 10.11, “Configuring Resource Mutual Exclusion Groups,”](#) on page 201
- ♦ [Section 10.12, “Controlling Resource Monitoring,”](#) on page 205
- ♦ [Section 10.13, “Changing the IP Address of a Cluster Resource,”](#) on page 206
- ♦ [Section 10.14, “Renaming a Cluster Resource,”](#) on page 208
- ♦ [Section 10.15, “Deleting Cluster Resources, or Disabling Clustering for a Pool, LVM Volume Group, or Service,”](#) on page 209
- ♦ [Section 10.16, “Additional Information for Creating Cluster Resources,”](#) on page 212

10.1 Planning for Cluster Resources

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

- ♦ [Section 10.1.1, “Naming Conventions for Cluster Resources,”](#) on page 178
- ♦ [Section 10.1.2, “Using Parameter Values with Spaces in a Cluster Script,”](#) on page 178
- ♦ [Section 10.1.3, “Using Double Quotation Marks in a Cluster Script,”](#) on page 178
- ♦ [Section 10.1.4, “Script Length Limits,”](#) on page 179
- ♦ [Section 10.1.5, “Planning Cluster Maintenance,”](#) on page 179

- ♦ [Section 10.1.6, “Number of Resources,” on page 179](#)
- ♦ [Section 10.1.7, “Linux POSIX File System Types,” on page 179](#)

10.1.1 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.

10.1.2 Using Parameter Values with Spaces in a Cluster Script

In cluster scripts, if a parameter value contains spaces, you should enclose the value with both double-quotes (") and single-quotes ('), such as `"'quoted text' "`. The preferred solution is to use both sets of quotation marks with the single-quotes on the inside. For OES 11 and later, the other order also works for spaces, but it reads a variable as text, not as a variable.

For example, either of the following techniques works for values with spaces:

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

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

You can alternatively use two sets of double-quotes and escape the inside set of double-quotes with a backslash (\). For example:

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

Another alternative is to escape each of the spaces in the value. For example:

```
exit_on_error echo Errors\ will\ be\ reported\ here.
```

10.1.3 Using Double Quotation Marks in a Cluster Script

In cluster scripts, if a command requires double-quotation marks, you should enclose the quoted part of the command with both double-quotes (") and single-quotes ('), such as `"'quoted text' "`. The key is to use both sets of quotation marks with the single-quotes on the inside. This works for spaces and allows you to pass variables to the command. For example:

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

You can alternatively use two sets of double-quotes and escape the inside set of double-quotes with a backslash (\). For example:

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


10.1.4 Script Length Limits

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.

10.1.5 Planning Cluster Maintenance

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

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

10.1.6 Number of Resources

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

10.1.7 Linux POSIX File System Types

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

10.2 Setting Up a Personalized List of Resources to Manage

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

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

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

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

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

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


- ♦ [Section 10.2.1, “Adding Resources to a My Resources List,” on page 180](#)
- ♦ [Section 10.2.2, “Viewing Information about Resources in a My Resources List,” on page 180](#)
- ♦ [Section 10.2.3, “Personalizing the My Resources Display,” on page 182](#)

- ♦ [Section 10.2.4, “Removing Resources from a My Resources List,”](#) on page 183
- ♦ [Section 10.2.5, “Updating the My Resources List after Renaming or Deleting a Resource,”](#) on page 184

10.2.1 Adding Resources to a My Resources List

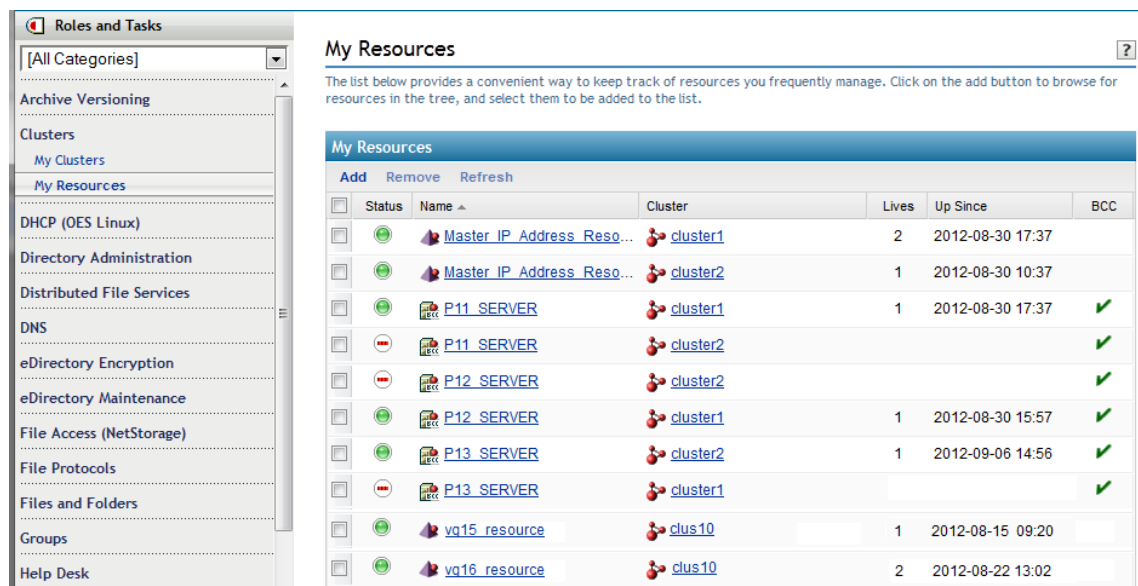
- 1 Log in to iManager as a cluster administrator.
- 2 In *Roles and Tasks*, select *Clusters > My Resources*.

The list is initially empty.



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

Newly selected resources are added to your personalized list.



- 5 (Optional) Personalize the display as described in [Section 10.2.3, “Personalizing the My Resources Display,”](#) on page 182.

10.2.2 Viewing Information about Resources in a My Resources List

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

The My Resources page displays your personalized list of cluster resources in the tree.

Roles and Tasks

[All Categories]

Archive Versioning

Clusters

My Clusters

My Resources

DHCP (OES Linux)

Directory Administration

Distributed File Services

DNS

eDirectory Encryption

eDirectory Maintenance

File Access (NetStorage)

File Protocols

Files and Folders

Groups

Help Desk

My Resources

The list below provides a convenient way to keep track of resources you frequently manage. Click on the add button to browse for resources in the tree, and select them to be added to the list.

My Resources

Add

Remove

Refresh

	Status	Name	Cluster	Lives	Up Since	BCC
<input type="checkbox"/>		Master IP Address Reso...	cluster1	2	2012-08-30 17:37	
<input type="checkbox"/>		Master IP Address Reso...	cluster2	1	2012-08-30 10:37	
<input type="checkbox"/>		P11_SERVER	cluster1	1	2012-08-30 17:37	
<input type="checkbox"/>		P11_SERVER	cluster2			
<input type="checkbox"/>		P12_SERVER	cluster2			
<input type="checkbox"/>		P12_SERVER	cluster1	1	2012-08-30 15:57	
<input type="checkbox"/>		P13_SERVER	cluster2	1	2012-09-06 14:56	
<input type="checkbox"/>		P13_SERVER	cluster1			
<input type="checkbox"/>		vq15_resource	clus10	1	2012-08-15 09:20	
<input type="checkbox"/>		vq16_resource	clus10	2	2012-08-22 13:02	

- View the Status icon next to a Cluster Resource object to understand the current state of the resource.

If scrolling is needed to view the complete list, you can use the *Refresh* option to set a longer refresh rate for the page to allow enough time to view the status for all items.

State	Icon	Description
Normal		The resource is online.
Offline		The resource is offline. Administrator intervention is needed.
Critical		<p>The resource is comatose, or it cannot be found (renamed or deleted). Administrator intervention is needed.</p> <p>If you rename or delete a resource, its status changes to "Resource not found" when you next access or refresh the My Resources page. For a deleted resource, remove the resource from the list. For a renamed resource, remove the resource from the list, then add the resource using its new name.</p>
Warning		An alert condition has occurred, and the resource needs administrator attention.
Unknown		The state of the cluster resource cannot be determined.

- 4 The My Resources list displays the following information about each resource:

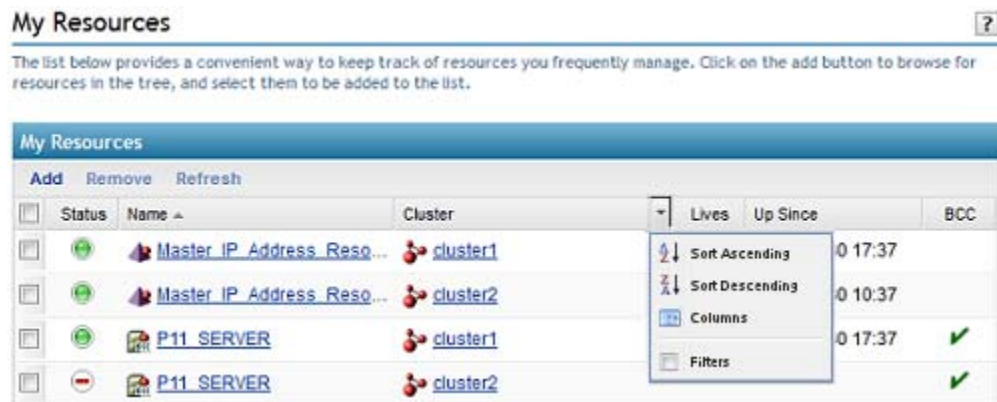
Parameter	Description
Name	The name assigned to the resource.
Distinguished Name	(Not displayed by default) The dot-delimited fully distinguished name of the resource, such as <code>resourcename.clustername.context.domain</code> .
Cluster	The name of cluster for the resource.
Lives	The number of times the resource has been brought online.
Up Since	The date and time the resource was brought online.
BCC	A check mark in this column indicates that the resource is enabled for a Business Continuity Cluster (BCC).

- 5 (Optional) Personalize the display as described in [Section 10.2.3, “Personalizing the My Resources Display,”](#) on page 182.

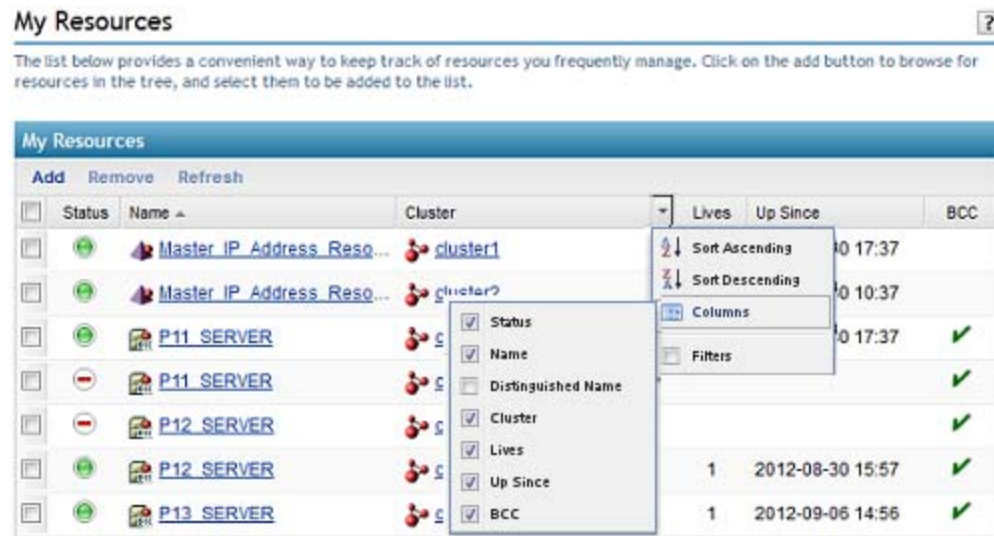
10.2.3 Personalizing the My Resources Display

The My Resources page provides menus in the column headings that allow you to personalize the display. You can sort the entries, modify the columns, or filter the entries. Your preferences are saved between logins to Novell iManager.

- 1 Log in to iManager as a cluster administrator.
- 2 In *Roles and Tasks*, select *Clusters > My Resources*.
- 3 Select *Refresh*, then choose a refresh rate that is long enough to allow you to view all items in the list.
- 4 Click or mouse-over a column heading to activate its options, then click the arrow to access the menu for that column.
- 5 Perform one of the available actions when you access the column’s menu. Some actions are not available for all fields.
 - ♦ **Sort:** Select *Sort Ascending* (A to Z) or *Sort Descending* (Z to A) to specify the preferred sort order. Resource objects are sorted based on values in the selected column. Numbers are treated as text values and are sorted alphabetically, not numerically.



- ♦ **Columns:** Select *Columns* to display the parameters, select or deselect a parameter's check box to add or remove the column, then press Enter to apply the changes. The parameters are described in [Step 4 in Section 10.2.2, "Viewing Information about Resources in a My Resources List,"](#) on page 180.



- ♦ **Filter:** Specify a filter to display only the Cluster objects with values that match.



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

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

10.2.4 Removing Resources from a My Resources List

- 1 Log in to iManager as a cluster administrator.
- 2 In *Roles and Tasks*, select *Clusters > My Resources*.
- 3 Select the check box next to the resource that you want to remove from your personalized list, then click *Remove*.

The resource is removed from your personalized list. This action does not delete the Cluster Resource object or its data.

10.2.5 Updating the My Resources List after Renaming or Deleting a Resource

In a [My Resources list](#), if you rename or delete a resource, its status changes to `Critical: Resource not found`. To update the list:

- ♦ **Deleted Resource:** Remove the resource entry from the list. Select the check box next to the resource entry, then click *Remove*.
- ♦ **Renamed Resource:** Remove the old resource entry from the list, then add the resource using its new name.
 1. Select the check box next to the resource entry, then click *Remove*.
 2. Click *Add*, select the newly named resource, then click *OK*.

10.3 Using Cluster Resource Templates

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

- ♦ [Section 10.3.1, “Default Resource Templates,” on page 184](#)
- ♦ [Section 10.3.2, “Viewing or Modifying a Resource Template in iManager,” on page 185](#)
- ♦ [Section 10.3.3, “Creating a Resource Template,” on page 187](#)
- ♦ [Section 10.3.4, “Synchronizing Locally Modified Resource Templates with Templates in Novell eDirectory,” on page 189](#)

10.3.1 Default Resource Templates

Novell Cluster Services provides several cluster resource templates that you can use on physical servers, virtualization host servers, and virtual machine (VM) guest servers (DomU). You can also create your own templates or personalize the default templates by using the Clusters plug-in for Novell iManager. For information, see [Section 10.3.3, “Creating a Resource Template,” on page 187](#). Third-party templates might also be available for third-party applications; see the vendor documentation.

[Table 10-1](#) identifies the cluster resource templates that can be used for OES services and storage on physical servers and virtual machine guest servers:

Table 10-1 Cluster Resource Templates for Physical Servers and VM Guest Servers

Cluster Resource Template	OES 11 and Later Product
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

Cluster Resource Template	OES 11 and Later Product
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 use in the virtualization host server environment. They are not supported for use by a VM guest server. The templates are generic and can be modified to start or stop different virtualization host server software.

Table 10-2 Cluster Resource Templates for Virtualization Host Environments

Cluster Resource Template	Use
Xen	Use to configure the cluster resource for a virtual machine. This template is available only on a virtualization host server. For information, see Section 15.2.1, “Creating a Xen Virtual Machine Cluster Resource,” on page 393.
XenLive	Use to configure the cluster resource for a virtual machine. It provides an additional function to allow a virtual machine resource migration (manual) without the need to boot or bring up the virtual machine on the cluster node where the virtual machine has been migrated. This template is available only on a virtualization host server. For information, see Section 15.2.3, “Setting Up Live Migration,” on page 399.


10.3.2 Viewing or Modifying a Resource Template in iManager

You can view or manage the Novell Cluster Services resource templates from the Cluster Options page in the Clusters plug-in for Novell iManager. Ensure that you use the Clusters plug-in to modify the properties of the resource template. The modified settings apply for all subsequent cluster resources that you create from the template.

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the cluster that you want to manage.


If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, “Setting Up a Personalized List of Clusters to Manage,”](#) on page 117.

- 3 Select the *Cluster Options* tab to access the templates.



- 4 Click the template name to view or modify the properties for the template.

For example, the following image shows the Properties page for the Generic File System template:



- 5 Click the tabs to view or modify the default property settings for the resource template.

The process for modifying a resource template is similar to modifying the settings for a cluster resource. If you modify a setting on the page, click *Apply* before you click the next tab in order to save the changes.

- ♦ **Policies**

Use the *Policies* tab to view or modify the default policies to use for the resource template. For information, see [Section 10.8, “Configuring the Start, Failover, and Failback Modes for Cluster Resources,”](#) on page 197.

- ♦ **Monitoring**

Use the *Monitoring* tab to enable or disable monitoring. If monitoring is enabled, you can set the default behaviors for monitoring, including the interval to poll the resource’s health, and the action to take if the resource fails to load on the maximum number of local restarts. For information, see [Section 10.7, “Enabling Monitoring and Configuring the Monitor Script,”](#) on page 192.

- ♦ **Preferred Nodes**

Use the *Preferred Nodes* tab to view or modify the preferred nodes to use for cluster resources that are created with the selected template. For information, see [Section 10.9, “Configuring Preferred Nodes and Failover Order for a Resource,”](#) on page 198.

- ♦ **Scripts**

Use the *Scripts* tab to view or modify the default load, unload, and monitor scripts for the selected template. For information, see the following:

- ♦ [Section 10.5, “Configuring a Load Script for a Cluster Resource,”](#) on page 190
- ♦ [Section 10.6, “Configuring an Unload Script for a Cluster Resource,”](#) on page 191
- ♦ [Section 10.7.3, “Example Monitor Scripts,”](#) on page 196

- ♦ **Business Continuity**

The *Business Continuity* tab contains information only if the selected cluster is enabled for Novell Business Continuity Clustering. If the page is enabled, you can view or set the default settings for resources created from the template. For information, see the [Novell Business Continuity Clustering documentation Web site \(http://www.novell.com/documentation/bcc\)](http://www.novell.com/documentation/bcc).

6 Click OK

10.3.3 Creating a Resource Template

Templates help ensure that all of the necessary definition information, dependent services to be loaded, unloaded, or monitored, and the shared service or storage are entered correctly. You can use the default templates as a guide for what types of information to include in your personalized template, as well as the sequence of commands for loading, unloading, or monitoring a resource.

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

2 Select the cluster that you want to manage.

If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, “Setting Up a Personalized List of Clusters to Manage,”](#) on page 117.

3 Select the *Cluster Options* tab to access the templates.

4 Click the *New* link.

5 Specify *Template* as the resource type you want to create by clicking the *Template* radio button, then click *Next*.

New Resource ?

Resource Type Select the type of cluster resource to create.

Type

- ☐ Pool
- ☐ Resource
- ☒ Template

- 6 In *Cluster Resource Name*, specify the name of the template you want to create.
- 7 If desired, in *Inherit from Template*, browse to the Cluster object and select the existing resource template in the Cluster container that you want to personalize for the new template.
- 8 Ensure that the *Define Additional Properties* check box is selected, then click *Next* to continue to the Load Script page.
- 9 On the Load Script page, configure the load script for the cluster resource template.
 - 9a Edit or add variables with example values for your template configuration, such as the mount point, IP address, volume group name, file system type, and mount device.
 - 9b Edit or add any lines to the load script that are required to load dependent services such as Web servers or file access protocols.
 - 9c Edit or add the necessary commands to the script to load the resource on the server.
For example, this might include bind command for the NCP service and the mount commands for the shared disks and file systems.
 - 9d Specify the default *Load Script Timeout* value, then click *Next* to continue to the Unload Script page.
The timeout value determines how much time the script is given to complete. If the script does not complete within the specified time, the resource becomes comatose. Cluster Services marks the process as failed right after the defined timeout expires, but it must wait for the process to conclude before it can start other resource operations.
- 10 On the Unload Script page, configure the unload script for the cluster resource template.
 - 10a Edit or add variables with example values for your template configuration, such as the mount point, IP address, volume group name, file system type, and mount device.
 - 10b Edit or add the necessary commands to the script to unload the resource from the server.
For example, this might include unbind command for the NCP service and the dismount commands for the shared disks and file systems.
 - 10c Edit or add any lines to the unload script that are required to unload the dependent services that are loaded by this cluster resource.
 - 10d Specify the default *Unload Script Timeout* value, then click *Next* to continue to the Monitor Script page.
The timeout value determines how much time the script is given to complete. If the script does not complete within the specified time, the resource becomes comatose when migrating to another node. Cluster Services marks the process as failed right after the defined timeout expires, but it must wait for the process to conclude before it can start other resource operations.
- 11 On the Monitor Script page, configure the monitor script for the cluster resource template.
 - 11a Edit or add the variables with example values for your template configuration, such as the mount point, IP address, volume group name, file system type, and mount device.
 - 11b Edit or add the necessary commands to the script to monitor the resource on the server.
You can use the same commands that are used at the Linux terminal console.

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

- 11c** Specify the default *Monitor Script Timeout* value, then click *Next*.

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

- 12** On the Resource Policies page, specify the default *Start*, *Failover*, and *Failback* modes, then click *Next*.
- 13** On the Resource Preferred Nodes page, specify the node assignments for the resource template, then click *Finish*.

The template you created is saved to the Cluster container of the cluster you selected. If you personalized an existing template, both the old template and the new template are in the container.

10.3.4 Synchronizing Locally Modified Resource Templates with Templates in Novell eDirectory

You can also view the resource templates outside of iManager. They are cached as `/var/opt/novell/ncs/*_Template.*` files on the master node of the cluster. Typically, you modify the templates by using the Clusters plug-in to Novell iManager as described in [Section 10.3.2, “Viewing or Modifying a Resource Template in iManager,” on page 185](#). However, if you add new template files or modify the files locally on the master cluster node, you must synchronize those changes with the resource templates and scripts that are held in the Cluster container in Novell eDirectory.

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

- 1** Log in to the master node of the cluster as the `root` user, then open a terminal console.
- 2** At the command prompt, enter the following commands:

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

10.4 Creating Cluster Resources

Cluster resources must be created for every shared file system or server that you run on servers in your cluster. Cluster resources can include 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, select *Clusters > My Clusters*.
- 2** Select the cluster that you want to manage.
If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 117](#).
- 3** Select the *Cluster Options* tab.
- 4** On the Cluster Options page, click the *New* link.
- 5** Specify *Resource* as the resource type you want to create by clicking the *Resource* radio button, then click *Next*.

- 6 Specify the name of the resource you want to create.

NOTE: Do not use periods in cluster resource names. Novell clients interpret periods as delimiters. If you use a space in a cluster resource name, that space is converted to an underscore.

- 7 In the *Inherit From Template* field, specify one of the available templates, such as the *Generic_FS_Template*.
For information about cluster resource templates, see [Section 10.3, “Using Cluster Resource Templates,” on page 184](#).
- 8 Select the *Define Additional Properties* check box, then click *Next*.
- 9 If you are creating a new cluster resource, continue with [“Configuring a Load Script for a Cluster Resource” on page 190](#).

10.5 Configuring a Load Script for a Cluster Resource

A load script is required for each resource, service, disk, or pool in your cluster. The load script specifies the commands to start the resource or service on a server.

Example load scripts are available in the following sections:

- ♦ [Section 12.6, “Configuring a Load Script for the Shared NSS Pool,” on page 240](#)
- ♦ [Section 14.4.3, “Sample Load Script for a CSM Resource without a Segment Manager,” on page 377](#)
- ♦ [“Sample Generic LVM Resource Load Script” on page 341](#)

If you are creating a new cluster resource, the load script page should already be displayed. You can start with [Step 6](#).

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the cluster that you want to manage.
If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 117](#).
- 3 Select the *Cluster Options* tab.
- 4 Select the cluster resource to view its Properties page.
You can also select the check box next to the cluster resource, then click *Details*.
- 5 On the Properties page, click the *Scripts* tab, then click the *Load Script* link.
- 6 Edit or add the necessary commands to the script to load the resource on a server.
You can then add any lines to the load script that are required to load needed services like Web servers, and so on.

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

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

- 7 Specify the *Load Script Timeout* value, then click *Apply* to save the script or, if you are creating a new cluster resource, click *Next*.

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

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

8 Do one of the following:

- ♦ If you are configuring a new resource, click *Next*, then continue with [Section 10.6, “Configuring an Unload Script for a Cluster Resource,” on page 191](#).
- ♦ Click *Apply* to save your changes.

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

10.6 Configuring an Unload Script for a Cluster Resource

Depending on your cluster application or resource, you can add an unload script to specify how the application or resource should terminate. An unload script is not required by all resources, but is required for cluster-enabled Linux partitions. Consult your application vendor or documentation to determine if you should add commands to unload the resource.

Programs should be unloaded in the reverse order of how they were loaded. This ensures that supporting programs are not unloaded before programs that rely on them in order to function properly.

Example unload scripts are available in the following sections:

- ♦ [Section 12.7, “Configuring an Unload Script for the Shared NSS Pool,” on page 242](#)
- ♦ [Section 14.4.4, “Sample Unload Script for CSM Resource without a Segment Manager,” on page 377](#)
- ♦ [“Sample Generic LVM Resource Unload Script” on page 342](#)

If you are creating a new cluster resource, the unload script page should already be displayed. You can start with [Step 6](#).

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

2 Select the cluster that you want to manage.

If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 117](#).

3 Select the *Cluster Options* tab.

4 Select the cluster resource to view its Properties page.

You can also select the check box next to the resource, then click *Details*.

5 On the Properties page, click the *Scripts* tab, then click the *Unload Script* link.

6 Edit or add the necessary commands to the script to unload the resource on the server.

You can add any lines to the unload script that are required to unload services that are loaded by this cluster resource.

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

- 7 Specify the *Unload Script Timeout* value, then click *Apply* to save the script or, if you are creating a new cluster resource, click *Next*.

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

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

- 8 Do one of the following:

- ♦ If you are configuring a new resource, click *Next*, then continue with [Section 10.7.2, “Configuring Resource Monitoring,” on page 194](#).
- ♦ Click *Apply* to save your changes.

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

10.7 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 10.7.1, “Understanding Resource Monitoring,” on page 192](#)
- ♦ [Section 10.7.2, “Configuring Resource Monitoring,” on page 194](#)
- ♦ [Section 10.7.3, “Example Monitor Scripts,” on page 196](#)
- ♦ [Section 10.7.4, “Monitoring Services that Are Critical to Clustering,” on page 196](#)

10.7.1 Understanding Resource Monitoring

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

- ♦ [“Polling Interval” on page 192](#)
- ♦ [“Failure Rate” on page 193](#)
- ♦ [“Failure Action” on page 193](#)
- ♦ [“Timeout Value” on page 193](#)
- ♦ [“How Resource Monitoring Works” on page 194](#)

Polling Interval

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

Failure Rate

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

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

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

Failure Action

The *Failover Action* indicates whether you want the resource to be set to a comatose state, to migrate to another server, or to reboot the hosting node (without synchronizing or unmounting the disks) if a failure action initiates. With resource monitoring, the *Start*, *Failover*, and *Failback* Modes have no effect on where the resource migrates. This means that a resource that has been migrated by the resource monitoring failure action does not migrate back (fail back) to the node it migrated from unless you manually migrate it back.

Set Resources as Comatose: (Default) If the failure action initiates, the resource is placed in a comatose state. Administrator action is required to take the resource offline, resolve the issue, and bring it online again on the same or different node.

Migrate the Resource Based on the Preferred Nodes List: If the failure action initiates and the resource is on its most preferred node, the resource migrates to the next available node in its *Preferred Nodes* list, which you previously ordered according to your failover order preferences. The resource is not automatically failed back to the original node. Administrator action is required to cluster migrate the resource to the node, as desired. Each time a failure action triggers a failover, the resource migrates to a different node, according to the order in its *Preferred Nodes* list and availability of the nodes.

Reboot the Hosting Node without Syncing or Unmounting Disks: If the failure action initiates, each of the resources on the hosting node will fail over to the next available node in its *Preferred Nodes* list because of the reboot. All resources on the node are failed over. This is a hard reboot, not a graceful one. The reboot option is normally used only for a mission-critical cluster resource that must remain available. The resources are not automatically failed back to the original node. Administrator action is required to cluster migrate them back to the node, as desired.

Timeout Value

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

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

How Resource Monitoring Works

- 1 The monitor script runs at the frequency you specify as the polling interval.
- 2 There are two conditions that trigger a response by 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)

10.7.2 Configuring Resource Monitoring

The resource monitoring function allows you to monitor the health of a specified resource by using a script that you create or customize. If you want 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 6](#).

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the cluster that you want to manage.

If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, "Setting Up a Personalized List of Clusters to Manage,"](#) on page 117.
- 3 Select the *Cluster Options* tab.
- 4 Click the cluster resource to open its Properties page.

You can also select the check box next to the resource, then click *Details*.

- 5 On the Properties page, click the *Monitoring* tab.

My Clusters > clus1.ncs.novell > P_D_SERVER

P_D_SERVER

Policies **Monitoring** Preferred Nodes Scripts Protocols 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. Changes other than business continuity changes made to a resource will not take affect until the resource is reloaded.

☒ Enable Resource Monitoring

Polling Interval:

Failure Rate

Maximum Local Failures:

Time Interval:

(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

- 6 Select the *Enable Resource Monitoring* check box to enable resource monitoring for the selected resource.

Resource monitoring is disabled by default.

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

You can specify the value in minutes or seconds.

- 8 Specify the number of failures (*Maximum Local Failures*) for the specified amount of time (*Time Interval*).

For information, see [“Failure Rate” on page 193](#).

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

For information, see [“Failure Action” on page 193](#).

- 10 Click the *Scripts* tab, then click the *Monitor Script* link.

- 11 Edit or add the necessary commands to the script to monitor the resource on the server.

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

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

You can use the same commands that would be used at the Linux terminal console. For example, see [Section 10.7.4, “Monitoring Services that Are Critical to Clustering,” on page 196](#).

- 12** Specify the *Monitor Script Timeout* value, then click *Apply* to save the script.

The timeout value determines how much time the script is given to complete. If the script does not complete within the specified time, the failure action you chose in [Step 9](#) initiates.

- 13** Do one of the following:

- ♦ If you are configuring a new resource, click *Next*, then continue with [Section 10.8.2, “Setting the Start, Failover, and Failback Modes for a Resource,” on page 197](#).
- ♦ Click *Apply* to save your changes.

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

10.7.3 Example Monitor Scripts

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

Example monitor scripts are available in the following sections:

- ♦ [“Configuring a Monitor Script for the Shared NSS Pool” on page 244](#)
- ♦ [“Sample Generic LVM Resource Monitor Script” on page 342](#)
- ♦ [“Sample Monitor Script for a CSM Resource without a Segment Manager” on page 378](#)

10.7.4 Monitoring Services that Are Critical to Clustering

Monitor scripts can also be used for monitoring critical services needed by the resources, such as Linux User Management (*namcd*) and Novell eDirectory (*ndsd*). However, the monitoring is in effect only where the cluster resource is running.

IMPORTANT: The monitor script runs only on the cluster server where the cluster resource is currently online. The script does not monitor the critical services on its assigned cluster server when the resource is offline. The monitor script does not monitor critical services for any other cluster node.

For example, to monitor whether the *namcd* and *ndsd* services are running, add the following commands to the Monitor script:

```
# (optional) status of the eDirectory service
exit_on_error rcndsd status
```

```
# (optional) status of the Linux User Management service
exit_on_error 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.

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

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

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

- ♦ [Section 10.8.1, “Understanding Cluster Resource Modes,” on page 197](#)
- ♦ [Section 10.8.2, “Setting the Start, Failover, and Failback Modes for a Resource,” on page 197](#)

10.8.1 Understanding Cluster Resource Modes

With the resource *Start* mode set to AUTO, the resource automatically starts on a server when the cluster is first brought up. If the resource *Start* mode is set to MANUAL, you can manually start the resource on a server when you want, instead of having it automatically start when servers in the cluster are brought up.

With the resource *Failover* mode set to AUTO, the resource automatically starts on the next server in the Assigned Nodes list 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.

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

If you are creating a new cluster resource, the Resource Policies page should already be displayed. You can start with [Step 6](#).

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the cluster that you want to manage.

If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 117](#).

- 3 Select the *Cluster Options* tab.
- 4 Select the cluster resource to open its Properties page.

You can also select the check box next to the cluster resource, then click *Details*.
- 5 On the Properties page, click the *Policies* tab.
- 6 (Optional) Select the *Resource Follows Master* check box if you want to ensure that the resource runs only on the master node in the cluster.

If the master node in the cluster fails, the resource fails over to whichever node becomes the master.
- 7 (Optional) Select the *Ignore Quorum* check box if you don't want the cluster-wide timeout period and node number limit enforced.

The quorum default values were set when you installed 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.
- 8 Specify the *Start*, *Failover*, and *Failback* modes for this resource.

The default for both *Start* and *Failover* modes is AUTO, and the default for *Failback* mode is DISABLE.
- 9 Do one of the following:
 - ♦ If you are configuring a new resource, click *Next*, then continue with [Section 10.9, "Configuring Preferred Nodes and Failover Order for a Resource,"](#) on page 198.
 - ♦ Click *Apply* to save your changes.

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

10.9 Configuring Preferred Nodes and Failover Order for a Resource

The Preferred Nodes page allows you to assign nodes to use for the resource. You also sequence the list of nodes to specify the preferred order that the nodes will be tried when a resource is brought online after its current node fails or if you do not specify a node when you bring the resource online. Changes are not allowed for the Preferred Nodes list for the *Master_IP_Address_Resource*.

By default, a new resource is configured to run on all nodes, including future nodes. If its preferred node list is ever changed (directly or indirectly) by administrators, future nodes are not automatically added to the resource's preferred nodes list.

IMPORTANT: Ensure that you prepare the node for the services in the resource before you migrate or fail over the resource to it.

If you are creating a new cluster resource, the Preferred Nodes page should already be displayed. If you are assigning nodes for an existing resource, the Preferred Nodes page is displayed as part of the Resource Policies page. You can start with [Step 6](#).

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the cluster that you want to manage.

If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, "Setting Up a Personalized List of Clusters to Manage,"](#) on page 117.

3 Select the *Cluster Options* tab.

4 Select the cluster resource to open its Properties page.

You can also select the box next to the cluster resource, then click the *Details* link.

5 On the Properties page, click the *Preferred Nodes* tab.

6 From the *Unassigned Nodes* area, select a node that the resource can use, then click the right-arrow button to move the selected node to the *Assigned Nodes* area.

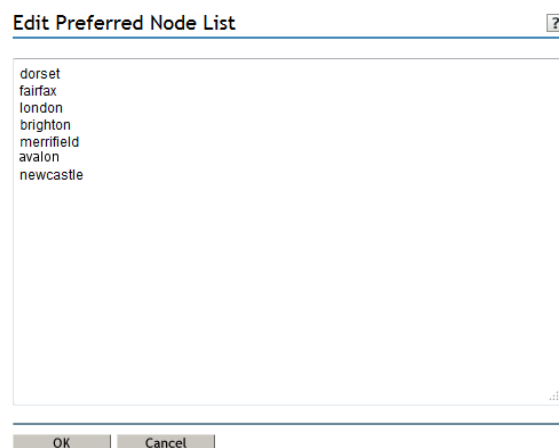
Repeat this step for all of the cluster nodes you want to assign to the resource.

7 From the *Assigned Nodes* area, select a node that you want to unassign from the resource, then click the left-arrow button to move the selected node to the *Unassigned Nodes* area.

8 Use one of the following methods to change the preferred failover order of the nodes assigned to a resource:

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

Be careful to not alter the names of the nodes while editing the list. If you remove nodes from the list in the Edit view, the removed nodes are automatically added to the *Unassigned Nodes* area when you save.



Edit Preferred Node List ?

dorset
fairfax
london
brighton
merrifield
avalon
newcastle

OK Cancel

9 At the bottom of the Preferred Nodes page, click *Apply* or *OK* to save the changes you made.

The new list and failover order apply immediately for OES 11 SP1 clusters, with the latest patches applied.

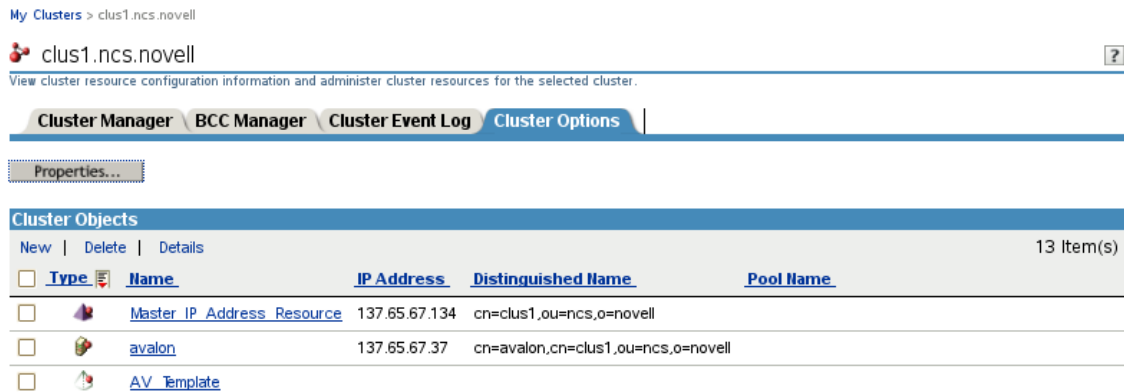
10 (OES 11) Offline the resource, then online it again to apply the revised preferred nodes list and failover order.

10.10 Configuring Resource Priorities for Load Order

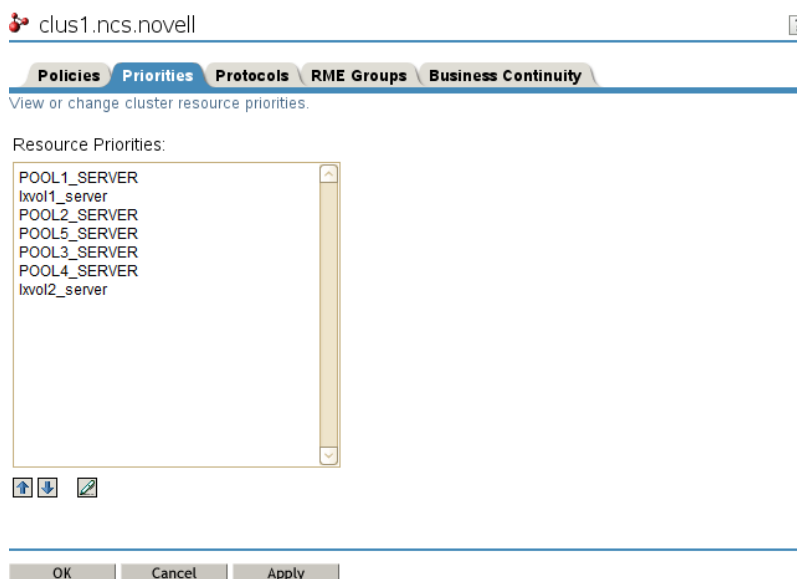
Cluster resource priorities control the load order of a resource relative to other cluster resources on the same cluster node when bringing up a cluster, or during a failover or failback. This is useful for ensuring that the most critical resources load first and are available to users before less critical resources.

The Resource Priority setting controls the order in which multiple resources start on a given node when the cluster is brought up or during a failover or failback. For example, if a node fails and two resources fail over to another node, the resource priority determines which resource loads first.

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the cluster that you want to manage.
If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, “Setting Up a Personalized List of Clusters to Manage,”](#) on page 117.
- 3 Select the *Cluster Options* tab.
- 4 On the Cluster Options page, click the *Properties* button above the *Cluster Objects* table to view the Cluster Properties dialog box.

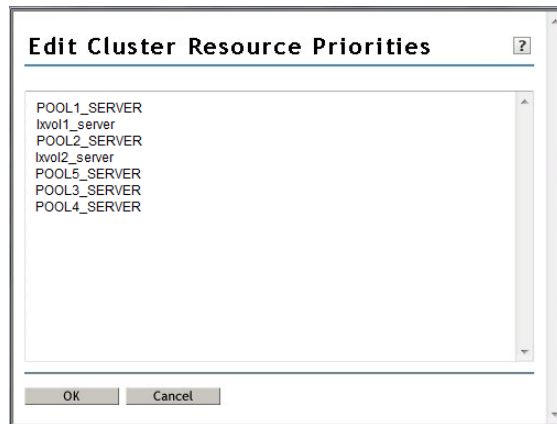


- 5 In the Cluster Properties dialog box, click the *Priorities* tab.



- 6 Use either of the following methods to change the load order (from highest priority to lowest priority) of a resource relative to other cluster resources on the same node:
- ♦ **Arrows:** Select a resource in the list by clicking on it, then click the up-arrow or down-arrow buttons to move the resource up or down in the list. The page refreshes between each click on an arrow.
 - ♦ **Edit:** Open the Resource Priority Edit function by clicking the *Edit* pen icon, list the resources in the preferred load order with one resource per line, then click *OK* to accept the revised order.

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



- 7 Click *Apply* or *OK* to save changes and close the Cluster Properties dialog box.

10.11 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 10.11.1, “Understanding Resource Mutual Exclusion,” on page 202](#)
- ♦ [Section 10.11.2, “Setting Up RME Groups,” on page 204](#)
- ♦ [Section 10.11.3, “Viewing RME Group Settings,” on page 204](#)

10.11.1 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 202](#)
- ♦ [“Rules for Node Assignment” on page 202](#)
- ♦ [“Examples” on page 202](#)
- ♦ [“Planning for RME Groups” on page 203](#)

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 onlineing, 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.

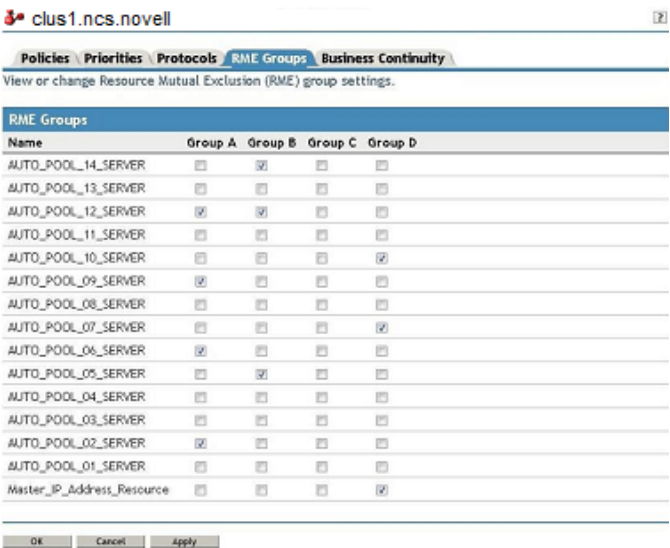
Planning for RME Groups

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.

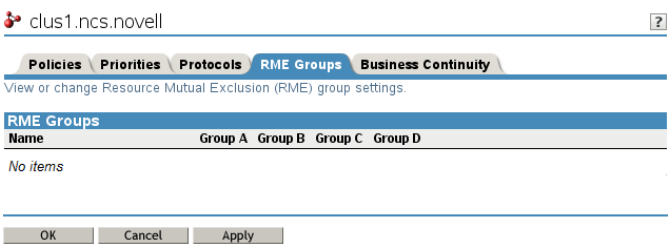
Figure 10-1 Sample RME Groups



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"/>

RME Groups are supported when all cluster nodes are running 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 for iManager, the tables on the RME Groups page are shown as *No items*. You cannot set up RME Groups for clusters running on older platforms.

Figure 10-2 Sample RME Groups Page for an Unsupported Platform



Name	Group A	Group B	Group C	Group D
No items				

If you are upgrading from an OES 2 SP2 or NetWare cluster, the RMS Groups are not available. You should wait until the rolling cluster upgrade has been completed on all nodes before you define RME Groups. If you define RME Groups while the cluster is in a mixed-mode condition, the RME Groups are honored only if the resource fails over to a cluster node where the new OS is running.

10.11.2 Setting Up RME Groups

To define sets of resources that must not be assigned to the same node at the same time:

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the cluster that you want to manage.
If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, “Setting Up a Personalized List of Clusters to Manage,”](#) on page 117.
- 3 Select the *Cluster Options* tab.
- 4 On the Cluster Options page, click *Properties*, then click the *RME Groups* tab.

The RME Groups page lists the resources in the cluster and columns for four possible RME Groups for the cluster. In each Group’s column, the check boxes correspond to the resources listed at the left.

If you are managing an older version of Novell Cluster Services (without the RME code) from a new Clusters plug-in for iManager, the tables are shown as *No items*. You cannot set up RME Groups for clusters running on older platforms.

- 5 To set up the members in a group, select the check boxes in the same column for two or more resources that must not be assigned on the same node at the same time.

The selected resources in a column are not allowed to run concurrently on a node. A resource can be a member of more than one group. You can define up to four groups (A, B, C, and D).

RME Groups				
Name	Group A	Group B	Group C	Group D
lvol1_server	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
lvol2_server	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Master_IP_Address_Resource	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
POOL1_SERVER	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
POOL2_SERVER	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
POOL3_SERVER	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
POOL4_SERVER	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
POOL5_SERVER	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

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

10.11.3 Viewing RME Group Settings

The RME Groups settings are displayed in the following locations:

- ♦ [“RME Groups in the Cluster Report”](#) on page 205
- ♦ [“RME Groups Configuration Page”](#) on page 205

RME Groups in the Cluster Report

The Cluster Report has an RME Groups section that lists the member resources of each group.

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the cluster you want to manage.
- 3 Click *Run Report*, then scroll down to view the *RME Groups*.

The RME Groups section lists only the resources in the cluster that are members in each of the RME Groups.

For example, the following report shows the members for Groups A, B, and D. Group C has no defined members. The `AUTO_POOL_12_SERVER` resource belongs to two RME Groups. It has mutually exclusive relationships with two different groups of resources. These exclusions are managed separately.

Resource Mutual Exclusion (RME) Groups			
Group A	Group B	Group C	Group D
AUTO_POOL_12_SERVER	AUTO_POOL_14_SERVER		AUTO_POOL_10_SERVER
AUTO_POOL_09_SERVER	AUTO_POOL_12_SERVER		AUTO_POOL_07_SERVER
AUTO_POOL_06_SERVER	AUTO_POOL_05_SERVER		Master_IP_Address_Resource
AUTO_POOL_02_SERVER			

RME Groups Configuration Page

The RME Groups configuration page displays all possible resources, and indicates membership by the check boxes that are selected in the same column under the RME Group Name.

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the cluster that you want to manage.
If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, “Setting Up a Personalized List of Clusters to Manage,”](#) on page 117.
- 3 Select the *Cluster Options* tab.
- 4 Click *Properties*, then click the *RME Groups* tab.
The RME Groups configuration page lists all of the resources in the cluster. The selected check boxes in the same group column indicate their memberships in the four possible RME Groups.

10.12 Controlling Resource Monitoring

You can use the `cluster monitor` command to start, stop, or view the status of monitoring for a specified cluster resource. Issue the command as the `root` user from the node where the resource is currently loaded:

```
cluster monitor <resourcename> {start | stop | status}
```

Monitoring must be enabled for the resource. For instructions, see [Section 10.7, “Enabling Monitoring and Configuring the Monitor Script,”](#) on page 192.

10.13 Changing the IP Address of a Cluster Resource

Typically, IP addresses for cluster resources are assigned for the life of the resource. You can modify the IP address of a cluster resource if desired. For example, you might need to assign a cluster and its resources to a different IP subnet.

Your setup must meet the following prerequisites:

- ♦ The resource must be in the *Offline* state in order to change its IP address.
- ♦ The new IP address must be unique in the network and available. It must satisfy the requirements described in [Section 4.2, “IP Address Requirements,” on page 52](#).
- ♦ Novell eDirectory must be running and working properly when you attempt to modify the IP address for the cluster resource.

Use the procedures in this section to modify the IP address for cluster resources:

- ♦ [Section 10.13.1, “Pool Cluster Resource,” on page 206](#)
- ♦ [Section 10.13.2, “Non-Pool Cluster Resources,” on page 207](#)

10.13.1 Pool Cluster Resource

You can change the IP address of a pool cluster resource by using the Protocols page in the resource’s Properties in the Clusters plug-in for Novell iManager. Ensure that you offline the pool cluster resource before attempting to change its IP address. When you save the change, the load, unload, and monitor scripts are automatically updated with the new IP address. It also automatically updates the resource’s IP address that is stored in its NCP Virtual Server object.

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

2 Select the cluster that you want to manage.

If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, “Setting Up a Personalized List of Clusters to Manage,” on page 117](#).

3 On the Cluster Manager page, select the check box next to the pool cluster resource that you want to manage, then click *Offline*.

Wait until the resource reports an Offline status before continuing.

4 Select the pool cluster resource to open its Properties page.

5 Select the *Protocols* tab.

The screenshot shows the 'POOL_D_SERVER' resource properties in the 'Protocols' tab. The 'Virtual Server Name' is 'CLUS1-POOL-D-SERVER,ou=ncs,o=novell' and the 'CIFS Server Name' is 'CLUS1-POOL-D-W'. The 'IP Address' field is set to '10.10.10.41'. Under 'Advertising Protocols', the checkboxes for 'AFP', 'CIFS', and 'NCP' are all checked. At the bottom are 'OK', 'Cancel', and 'Apply' buttons.

POOL_D_SERVER

Policies Monitoring Preferred Nodes Scripts **Protocols** Business Continuity

View or change the Virtual Server Name, Pool IP Address, Advertising Protocols and, if necessary the CIFS Server Name for this cluster pool resource. Changes other than business continuity changes made to a resource will not take affect until the resource is reloaded.

Virtual Server Name: CLUS1-POOL-D-SERVER,ou=ncs,o=novell
CIFS Server Name: CLUS1-POOL-D-W
IP Address: 10.10.10.41

Advertising Protocols:

☒ AFP
☒ CIFS
☒ NCP

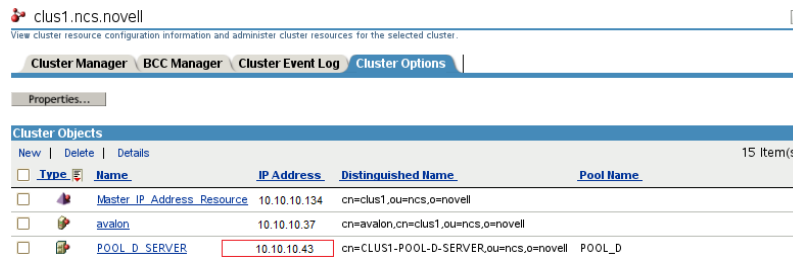
OK Cancel Apply

- 6 In the IP Address field, type the new IP address for the cluster resource in IPv4 dotted decimal format (four octets of the address expressed individually in decimal and separated by periods).
For example, type:

10.10.10.43

- 7 Click OK to save and apply the changes.

The Cluster Options page opens. The new IP address is displayed for the pool cluster resource.



- 8 (Optional) Verify that the IP address has been changed in the pool cluster resource's load, unload, and monitor scripts.
 - 8a On the Cluster Options page, select the cluster resource to open its Properties page, then select the *Scripts* tab.
 - 8b Select the Load, Unload, and Monitor options in turn, and review the scripts to verify that the resource IP address has been changed in all of the relevant places.
- 9 Select the *Cluster Manager* tab.
- 10 On the Cluster Manager page, select the check box next to the pool cluster resource, then click *Online*.

Wait until the resource reports an Online status before continuing to ensure that the resource is working properly.

10.13.2 Non-Pool Cluster Resources

The properties for non-pool cluster resources do not include a Protocols page. You can change the IP address of a non-pool cluster resource by changing the IP address value for the RESOURCE_IP variable in its load, unload, and monitor scripts.

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the cluster that you want to manage.
If the cluster does not appear in your list, add the cluster to your list as described in [Section 8.2, "Setting Up a Personalized List of Clusters to Manage,"](#) on page 117.
- 3 On the Cluster Manager page, select the check box next to the non-pool cluster resource that you want to manage, then click *Offline*.
Wait until the resource reports an Offline status before continuing.
- 4 Select the cluster resource to open its Properties page.
- 5 Change the IP address value for the RESOURCE_IP variable in the load, unload, and monitor scripts:
 - 5a Select the *Scripts* tab.
It automatically opens to the Load Script page.

- 5b** On the Load Script page, type the new IP address in the RESOURCE_IP variable, then click *Apply*.
- 5c** Click *Unload Script*, type the new IP address in the RESOURCE_IP variable, then click *Apply*.
- 5d** Click *Monitor Script*, type the new IP address in the RESOURCE_IP variable, then click *Apply*.
- 5e** Click *OK* to return to the Cluster Options page.
- 6** Select the *Cluster Manager* tab.
- 7** On the Cluster Manager page, select the check box next to the cluster resource, then click *Online*.
Wait until the resource reports an Online status before continuing to ensure that the resource is working properly.

10.14 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
```


10.15 Deleting Cluster Resources, or Disabling Clustering for a Pool, LVM Volume Group, or Service

You can delete a cluster resource for any of the following reasons:

- ♦ You want to re-create a Cluster Resource object. You want to remove the related objects in eDirectory and get the shared storage ready so that the resource can be re-created cleanly.

To delete a resource and create a new one with the same name, you must wait to create the new one until eDirectory synchronizes all of the objects in the tree related to the deleted resource.

- ♦ You want to disable clustering for the shared storage managed by the resource. Afterwards, you will assign the SAN device to only a single node in the cluster. The storage will be available via the node's IP address instead of the cluster resource IP address.
- ♦ You want to disable clustering for the shared service. The service will be available via the node's IP address instead of the cluster resource IP address.
- ♦ You want to delete the shared storage managed by the resource. Afterwards, data is no longer available on any node in the cluster.

We strongly recommend that you delete cluster resource objects only from the master node in the cluster. Ensure that you take the cluster resource offline before you attempt to delete the cluster resource or to delete cluster-enabled shared storage managed by the resource.

WARNING: If you attempt to delete a storage cluster resource without first taking it offline, deletion errors occur, and the data associated with the clustered storage is not recoverable.

All resource configuration must happen from the master node. On the Cluster Options page for iManager, you are automatically connected to the Cluster object, which is associated with the master node. On the *Storage > Pools* page for iManager, connect to the Cluster object, not to the individual servers. Run NSSMU only on the master node.

Use the following procedure to delete a cluster resource:

- 1 If the resource is on a non-master node in the cluster, migrate it to the master node.

As the *root* user, open a terminal console, then enter

```
cluster migrate <resource_name> <master_node_name>
```

The master node must be in the resource's preferred nodes list. To view or modify the list, see [Section 10.9, "Configuring Preferred Nodes and Failover Order for a Resource,"](#) on page 198.

- 2 If the cluster resource is online, take it offline by using one of the following methods:

- ♦ Enter the following at the command prompt as the *root* user:

```
cluster offline <resource_name>
```

Use the `cluster status` command to verify that the resource has a status of *Offline* before you continue.

- ♦ In iManager, go to *Clusters > My Clusters*, then select the cluster. On the Cluster Manager page, select the check box next to the cluster resource, then click *Offline*.

Refresh the Cluster Manager page to verify that the resource has a status of *Offline* before you continue.

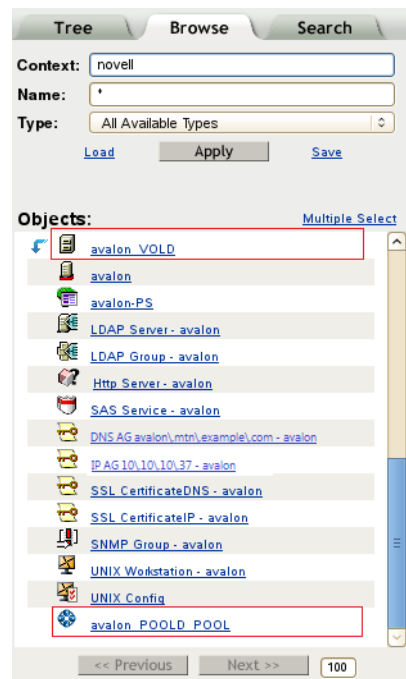
- 3 In iManager, use the Clusters plug-in to delete the cluster resource.

3a Select *Clusters > My Clusters*, then select the cluster.

3b Select the *Cluster Options* tab.

- 3c** Select the check box next to the resource, then click *Delete*.
- This deletes the Cluster Resource object. It does not delete the storage or service represented by the object.
- For a pool cluster resource, this also deletes the related NCP virtual server object, Pool object, and Volume object.
- For an NCP-enabled LVM cluster resource, this also deletes the related NCP virtual server object and Volume object.
- 3d** When you are prompted to confirm the deletion, click *OK* to continue, or click *Cancel* to abort the deletion.
- 3e** In the *Tree View* in iManager, browse to verify that the Cluster Resource objects and related objects were removed from eDirectory.
- If necessary, you can manually delete the objects. In iManager, go to *Directory Administration > Delete Objects*, select the objects, then click *OK*.
- 4** If the deleted resource was a pool cluster resource, use the *Update eDirectory* function to re-create Storage objects for the pool and its volumes.
- 4a** In iManager, select *Storage > Pools*, then select the master node if you plan to re-create the storage object, or select the node where you want the pool to reside as a locally available pool.
- 4b** Select the pool, then click *Activate*.
- 4c** Select the pool, then click *Update eDirectory*.
- This creates a Pool object in eDirectory with a name format of `<server_name>_<pool_name>_POOL`.
- 4d** Select *Storage > Volumes*. The server should still be selected.
- 4e** Select the volume, then click *Mount*.
- 4f** Select the volume, then click *Update eDirectory*.
- This creates a Volume object in eDirectory with a name format of `<server_name>_<volume_name>`.
- 4g** Repeat [Step 4d](#) through [Step 4f](#) for each volume in the pool.
- 4h** In the *Tree View*, browse to verify that the Pool object and Volume object were created.

For example, the following Tree View shows the Pool object `avalon_POOLD_POOL` and the Volume object `avalon_VOLD` were created for server `avalon`.



5 Do one of the following:

- ♦ **Re-create the cluster resource:** Use the Clusters plug-in to cluster-enable the storage area.
 - ♦ To cluster-enable a pool, see [Section 12.5, “Cluster-Enabling an Existing NSS Pool and Its Volumes,”](#) on page 233.
 - ♦ To cluster-enable a Linux LVM volume group and logical volume, see [Section 13.4.2, “Creating a Generic File System Cluster Resource for an LVM Volume Group,”](#) on page 335 and [Section 13.5, “Creating a Virtual Server Object for an LVM Volume Group Cluster Resource,”](#) on page 343

To re-create the cluster resource with the same name, you must wait to create the new one until eDirectory synchronizes all of the objects in the tree related to the deleted resource.

- ♦ **Use the non-clustered pool:** Unshare the devices that contribute space to the pool, then use a third-party SAN management tool to assign the devices to only the desired server.
 1. In iManager, go to *Storage > Devices*, then select the desired server (the one you specified in [Step 4](#)).
 2. Select the device.
 3. Deselect the *Shareable for Clustering* check box, then click *Apply*.
 Unsharing a device fails if the device contains a cluster-enabled pool or split-brain detector (SBD) partition. This is unlikely to be an issue if you used a dedicated device (or devices) for the pool where you have disabled clustering.
 4. Repeat these steps for each device that contributes space to the pool.
 5. Use a third-party SAN management tool to assign the devices to only the desired server.
 6. Provide the node’s IP address to users.

- ♦ **Use the non-clustered Linux LVM volume:** The LVM volume group uses the entire device. Clustered LVM (cLVM) still recognizes the multiple-node assignments from the SAN.
 1. Use a third-party SAN management tool to assign the device to only the desired server.
 2. Provide the node's IP address to users.
- ♦ **Use the non-clustered service:** Modify information for your users so that they access the local node's IP address instead of a clustered service IP address.
- ♦ **Delete the pool or LVM volume group:** If you do not want to keep the data, delete the shared storage area.

WARNING: Deleting a pool or a Linux LVM volume group destroys all data on it.

- ♦ For shared NSS pools and volumes, use NSSMU, the Storage plug-in to iManager, or the `nlvm delete pool <pool_name>` command. Deleting the pool automatically deletes the volumes on it.
- ♦ For Linux LVM volume groups, use NSSMU or the `nlvm delete linux volume <volume_name>` command. These tools automatically delete the LVM logical volume and logical volume group. If the volume was NCP-enabled, it also deletes the related NCP volume.

10.16 Additional Information for Creating Cluster Resources

Cluster resources can be configured for storage, services, or virtual machines.

- ♦ [Section 10.16.1, “Creating Storage Cluster Resources,” on page 212](#)
- ♦ [Section 10.16.2, “Creating Service Cluster Resources,” on page 213](#)
- ♦ [Section 10.16.3, “Creating Virtual Machine Cluster Resources,” on page 213](#)

10.16.1 Creating Storage Cluster Resources

For information about creating cluster resources for shared storage on Linux, see the following:

Table 10-3 *Cluster-Enabling Shared Storage*

Shared Storage	Refer to
NSS pools and volumes	Chapter 12, “Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes,” on page 217
Linux POSIX volumes	Chapter 14, “Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers,” on page 371
NCP volumes	“Configuring NCP Volumes with Novell Cluster Services” in the <i>OES 11 SP1: 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 Volume Pairs with Novell Cluster Services” in the <i>OES 11 SP1: Dynamic Storage Technology Administration Guide</i>

10.16.2 Creating Service Cluster Resources

For information about creating cluster resources for various services on your OES 11 SP1 server, see [Chapter 11, “Creating and Managing Service Cluster Resources,”](#) on page 215.

10.16.3 Creating Virtual Machine Cluster Resources

If you install Novell Cluster Services at the host level of an OES 11 SP1 virtualized server, you can create cluster resources for the virtual machines. For information, see [Section 15.2, “Virtual Machines as Cluster Resources,”](#) on page 392.

11 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 SP1 server. See [Table 11-1](#) for links to those resources. You can also find this list by going to “Clustering OES Services” on the OES 11 SP1 Documentation Web site (http://www.novell.com/documentation/oes11/oes11_toc/data/index-az.html#av-c).

Table 11-1 *Clustering OES 11 SP1 Services with Novell Cluster Services*

OES 11 SP1 Service	For information, see
Apache Web Server	“Apache HTTP Server” in the OES 11 SP1: Novell Cluster Services 2.1 NetWare to Linux Conversion Guide
Archive and Version Services	“Configuring Archive and Version Service for Novell Cluster Services” in the OES 11 SP1: 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 SP1: Novell Cluster Services 2.1 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 SP1: 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 SP1: 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 OES 11 SP1: Novell AFP for Linux Administration Guide
File, CIFS (Windows File Services)	“Configuring CIFS with Novell Cluster Services for an NSS File System” in the OES 11 SP1: Novell CIFS for Linux Administration Guide

OES 11 SP1 Service	For information, see
File, FTP	“Cluster Enabling Pure-FTPd in an OES 11 SP1 Environment” in the <i>OES 11 SP1: Planning and Implementation Guide</i>.
File, NetStorage	“Configuring NetStorage with Novell Cluster Services” in the <i>OES 11 SP1: NetStorage Administration Guide for Linux</i>.
File, Samba	“Configuring Samba for LVM Volume Groups and Novell Cluster Services” in the <i>OES 11 SP1: Novell Samba Administration Guide</i>
iFolder 3.9	“Clustering iFolder Servers with Novell Cluster Services for Linux” in the <i>Novell iFolder 3.9.1 Administration Guide</i>
iPrint	“Configuring iPrint with Novell Cluster Services” in the <i>OES 11 SP1: iPrint Linux Administration Guide</i>
MySQL	A MySQL template is available that uses a shared LVM volume group that you have already created. For information, see “Configuring MySQL with Novell Cluster Services” in the <i>OES 11 SP1: Web Services and Applications Guide</i>.
QuickFinder (Server Synchronization Feature)	“Configuring QuickFinder Server for Novell Cluster Services” in the <i>OES 11 SP1: Novell QuickFinder Server 5.0 Administration Guide</i>
Storage, DST shadow volumes built with NSS volumes	“Configuring DST Shadow Volume Pairs with Novell Cluster Services” in the <i>OES 11 SP1: Dynamic Storage Technology Administration Guide</i>
Storage, Linux POSIX volumes on legacy CSM containers	Chapter 14, “Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers,” on page 371.
Storage, LVM volume groups and logical volumes	Chapter 13, “Configuring and Managing Cluster Resources for Shared LVM Volume Groups,” on page 305.
Storage, NCP volumes	Section 13.6, “Enabling NCP File Access for a Clustered LVM Volume,” on page 348 “Configuring NCP Volumes with Novell Cluster Services” in the <i>OES 11 SP1: NCP Server for Linux Administration Guide</i> “Clustering LVM Volume Groups with Novell Cluster Services” in the <i>OES 11 SP1: Linux POSIX Volume Administration Guide</i> “Create Linux Volume” in the <i>OES 11 SP1: NLVM Reference</i>
Storage, NSS pools and volumes	Chapter 12, “Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes,” on page 217
Xen virtual machines	Section 15.2, “Virtual Machines as Cluster Resources,” on page 392

12 Configuring and Managing 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 12.1, “Requirements for Creating Pool Cluster Resources,” on page 218](#)
- ♦ [Section 12.2, “Guidelines for Using Pool Cluster Resources,” on page 222](#)
- ♦ [Section 12.3, “Initializing and Sharing a SAN Device,” on page 223](#)
- ♦ [Section 12.4, “Creating Cluster-Enabled Pools and Volumes,” on page 227](#)
- ♦ [Section 12.5, “Cluster-Enabling an Existing NSS Pool and Its Volumes,” on page 233](#)
- ♦ [Section 12.6, “Configuring a Load Script for the Shared NSS Pool,” on page 240](#)
- ♦ [Section 12.7, “Configuring an Unload Script for the Shared NSS Pool,” on page 242](#)
- ♦ [Section 12.8, “Configuring a Monitor Script for the Shared NSS Pool,” on page 244](#)
- ♦ [Section 12.9, “Adding Advertising Protocols for NSS Pool Cluster Resources,” on page 245](#)
- ♦ [Section 12.10, “Adding NFS Export for a Clustered Pool Resource,” on page 247](#)
- ♦ [Section 12.11, “Mirroring and Cluster-Enabling Shared NSS Pools and Volumes,” on page 248](#)
- ♦ [Section 12.12, “Renaming a Clustered NSS Pool,” on page 252](#)
- ♦ [Section 12.13, “Renaming a Clustered NSS Volume,” on page 264](#)
- ♦ [Section 12.14, “Renaming the Mount Point Path for a Clustered NSS Volume,” on page 275](#)
- ♦ [Section 12.15, “Changing the Name Space for a Clustered NSS Volume,” on page 281](#)
- ♦ [Section 12.16, “Adding a Volume to a Clustered Pool,” on page 282](#)
- ♦ [Section 12.17, “Changing an Assigned Volume ID,” on page 283](#)
- ♦ [Section 12.18, “Expanding the Size of a Clustered Pool,” on page 284](#)
- ♦ [Section 12.19, “Disabling Clustering for a Pool,” on page 296](#)
- ♦ [Section 12.20, “Deleting and Re-Creating a Pool Cluster Resource \(Disabling and Re-Enabling Clustering for a Pool\),” on page 298](#)
- ♦ [Section 12.21, “Deleting a Clustered Pool,” on page 301](#)

12.1 Requirements for Creating Pool Cluster Resources

Your system must meet the requirements in this section in addition to the cluster requirements described in [Chapter 4, “Planning for Novell Cluster Services,” on page 49](#).

- ♦ [Section 12.1.1, “Novell Cluster Services,” on page 218](#)
- ♦ [Section 12.1.2, “Resource IP Address,” on page 218](#)
- ♦ [Section 12.1.3, “Novell eDirectory,” on page 218](#)
- ♦ [Section 12.1.4, “Novell Storage Services,” on page 218](#)
- ♦ [Section 12.1.5, “Shared Storage,” on page 220](#)
- ♦ [Section 12.1.6, “NCP Server for Linux,” on page 220](#)
- ♦ [Section 12.1.7, “Novell CIFS for Linux,” on page 221](#)
- ♦ [Section 12.1.8, “Novell AFP for Linux,” on page 221](#)
- ♦ [Section 12.1.9, “Novell Samba,” on page 222](#)
- ♦ [Section 12.1.10, “Domain Services for Windows,” on page 222](#)

12.1.1 Novell Cluster Services

Novell Cluster Services must be installed, configured, and running when you share devices and when you create manage the shared NSS pools and volumes. The cluster must be active.

12.1.2 Resource IP Address

Each NSS pool cluster resource requires a unique static IP address. The IP address is used to provide access to and failover capability for the pool cluster resource. Users access the pool by using the resource IP address instead of the server IP address where the pool is active. The IP address you assign to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.

IMPORTANT: The IP address for the virtual server must be in the same IP subnet as the server nodes in the cluster where you plan to use it.

12.1.3 Novell eDirectory

Novell eDirectory must be running and working properly when you cluster-enable a pool and when you manage the pool.

12.1.4 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 SP1: NSS File System Administration Guide for Linux](#). For information about using NLVM commands, see [OES 11 SP1: NLVM Reference](#).

In addition, the following requirements must be met:

- ♦ [“Pool” on page 219](#)
- ♦ [“Volume” on page 219](#)
- ♦ [“Naming Conventions” on page 219](#)

- ♦ [“Pool Cluster Resource Name” on page 219](#)
- ♦ [“Location of Cluster, Server, and Storage Objects” on page 220](#)

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.

For information about using antivirus software with NSS volumes, see [“Antivirus Support for NSS”](#) in the *OES 11 SP1: NSS File System Administration Guide for Linux* and [“McAfee Antivirus Requires Additional Configuration”](#) in the *OES 11 SP1: Planning and Implementation Guide*.

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 10.1.1, “Naming Conventions for Cluster Resources,” on page 178](#).

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 context where the pool was created in order to modify the pool or volume, or to perform other tasks like setting up Distributed File Services junctions or home directories. You receive an eDirectory error if the operation cannot find the information that it needs in the same context.

12.1.5 Shared Storage

You should carefully plan how you want to configure your shared storage prior to installing 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 66](#)
- ♦ [Section 4.9, “SAN Rules for LUN Masking,” on page 69](#)
- ♦ [Section 4.10, “Multipath I/O Configuration Requirements,” on page 70](#)

12.1.6 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 SP1: 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 SP1: NCP Server for Linux Administration Guide*.

You must specify an NCP virtual server name for the cluster-enabled pool. By default, the suggested name is in the form of `<clusterName>-<poolName>-SERVER`. You can accept the suggested name, or specify a customized name for the NCP virtual server name. If you are also enabling CIFS as an

advertising protocol, you might want to choose an NCP virtual server name that can also be used as the CIFS virtual server name. For information about using the NCP virtual server name as the CIFS virtual server name, see [Section 12.1.7, “Novell CIFS for Linux,” on page 221](#).

12.1.7 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 SP1: 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 SP1: Novell CIFS for Linux Administration Guide*.

You must specify a CIFS virtual server name for the cluster enabled pool. You can accept the suggested name, or specify a customized name for the CIFS virtual server name. The name can be up to 15 characters, which is a restriction of the CIFS protocol.

For users to collaborate effectively, all paths for user access should be identical, independent of the access protocol used. This is possible only if the same name is used for the NCP virtual server name and the CIFS virtual server name, and the name can be only up to 15 characters.

By default, the NCP virtual server name is suggested as the CIFS virtual server name. If the name is more than 15 characters, the CIFS virtual server name uses the rightmost 13 characters and adds -w. For example, an NCP virtual server name of CLUSTER1-P_USERS is modified to STER1-P_USERS-w for the CIFS virtual server name. If a default NCP virtual server name was used in the form of `<clusterName>-<poolName>-SERVER` and the name exceeds 15 characters, the CIFS virtual server name uses the rightmost 13 characters of the `<clusterName>-<poolName>` part of the name, and adds -w. For example, an NCP virtual server name of CLUS1-P123-SERVER is modified to CLUS1-P123-w for the CIFS virtual server name.

If an administrator user later changes the NCP virtual server name in NSSMU or iManager, NSSMU automatically applies the name change to the CIFS virtual server name, so that the administrator does not need to make the change twice. To use a different server name for the CIFS virtual server name, you can change the CIFS virtual server name by using the CIFS management tools. This change will not affect the NCP virtual server name.

12.1.8 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 SP1: 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 [OES 11 SP1: Novell AFP for Linux Administration Guide](#).

12.1.9 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 SP1: Novell Samba Administration Guide](#).

Novell Samba requires that users are enabled for Linux with Linux User Management (LUM). For information, see [OES 11 SP1: 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 SP1: NCP Server for Linux Administration Guide](#).

12.1.10 Domain Services for Windows

Cluster-enabled NSS volumes can be used in a Domain Services for Windows environment.

12.2 Guidelines for Using Pool Cluster Resources

Consider the guidelines in this section when working with shared NSS pools and volumes in the cluster:

- ♦ [Section 12.2.1, “Guidelines for Shared Pools,” on page 222](#)
- ♦ [Section 12.2.2, “Guidelines for Shared Volumes,” on page 223](#)

12.2.1 Guidelines for Shared Pools

- ♦ When the pool cluster resource is brought online, the pool is automatically activated by the resource load script. You typically do not activate the pool at the terminal console.
- ♦ The unit of failover is the device (disk or LUN) that contains the pool. If you use multiple devices in the pool, all of those devices must be failed over together.

IMPORTANT: As a best practice, we recommend a single LUN (or device) per shared pool, and one volume per pool. (one LUN=>one pool=>one volume)

- ♦ If you delete a pool cluster resource, Novell Cluster Services automatically removes the Pool Resource object and the virtual server object from Novell eDirectory. After you delete the resource, you can unshare the device and use the pool locally, or you can delete the pool and its volumes.

Ensure that you offline the cluster resource before you attempt to delete the cluster resource or the pool. For information, see [Section 10.15, “Deleting Cluster Resources, or Disabling Clustering for a Pool, LVM Volume Group, or Service,” on page 209](#).

- 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.

12.2.2 Guidelines for Shared Volumes

You must create at least one volume on a shared pool before you migrate it. As a best practice, we recommend that you create only one volume per shared pool.

When you create a volume, commands are added to the pool resource load and unload scripts to automatically mount and dismount the volume when the scripts run. You can modify the load script to comment out the mount command so that you can manually mount the volume on a node in the cluster where the pool resource has been activated.

Volumes on cluster-enabled pools do not appear as independent cluster resources. A cluster-enabled volume does not have an associated load and unload script or an assigned IP address. If you want each volume to be in a separate cluster resource, each volume must have its own pool.

When a server fails, the cluster resources fail over to other servers in the cluster. Because the cluster-enabled pool fails over, all volumes in the pool also fail over, but only the volumes that have been added to the load script are automatically mounted. Any volumes in the pool that are not in the load script must be mounted manually. For this reason, volumes that you do not want to fail over should be in separate pools that are not cluster-enabled.

When you create an encrypted NSS volume in a shared pool, you must mount the volume manually by using NSSMU and enter the password. NSS uses the password to create a key. Instead of storing it in the server memory as it does for non-shared volumes, NSS asks 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.

12.3 Initializing and Sharing a SAN Device

NSS pools can be up to 8 TB in size and use space from one or more devices. For clustered pool resources, we recommend one device per pool. Each of the devices that you want to use for shared pool must be initialized and marked as shareable for clustering.

You must initialize a device (disk or LUN) to set up its device format before you can create an NSS pool on it. You can also initialize a device to wipe its current structure and reconfigure it. Only initialized devices are visible to the NSS management tools when you create a pool.

Initializing a device formats it with an MSDOS or a GPT partitioning scheme. MSDOS supports devices up to 2 TB in size. GPT supports devices of any size. The default is MSDOS. If the device size is greater than 2 TB and the partitioning scheme is not specified, the default partitioning scheme of MSDOS applies, and the device size is truncated to 2 TB with the remainder as unusable space.

WARNING: Initializing a device removes all partitions and data from the device. Do not initialize the device that contains the operating system.

Devices that have never been initialized have a format of *None*. Devices that are being used for a Novell Cluster Services SBD (split brain detector) partition also have a format of *None*; however, you should not use initialization to remove an SBD partition. For information about removing an SBD partition, see [Section 9.16, “Creating or Deleting Cluster SBD Partitions,” on page 157](#). After the SBD partition is removed, you can initialize the device.

You can initialize and share a device by using the Storage plug-in to Novell iManager, the Novell Storage Services (NSS) Management Utility (NSSMU), or the Novell Linux Volume Manager (NLVM) command line interface.

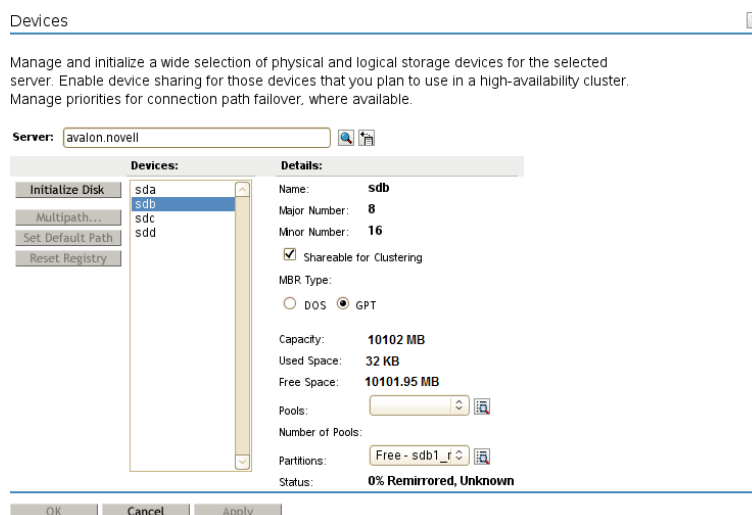
- [Section 12.3.1, “Initializing and Sharing a Device with iManager,” on page 224](#)
- [Section 12.3.2, “Initializing and Sharing a Device with NSSMU,” on page 225](#)
- [Section 12.3.3, “Initializing and Sharing a Device with NLVM,” on page 225](#)

12.3.1 Initializing and Sharing a Device with iManager

- 1 Ensure that the SAN device is attached to all nodes in the cluster.
- 2 Log in to iManager as an administrator user, then select **Storage > Devices**.
- 3 Browse to select the Cluster object (🔗) of the cluster.
Selecting the Cluster object automatically selects the server that is currently the master node in the cluster.
- 4 From the *Devices* list, select the device where you want to create the pool.
- 5 If the device has not been previously initialized or if you want to delete the current partitioning structures on the device, click *Initialize Disk*, then click *OK* to confirm and continue.

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

- 6 If you are prompted, specify whether to use the GPT or DOS partitioning scheme, then click *OK* to complete the initialization.
GPT can be used for any size device. DOS supports devices up to 2 TB in size.
Wait for the page to refresh before continuing.
- 7 In the *Details* area, select the *Shareable for Clustering* check box, then click *OK* to apply the change.



- 8 Repeat [Step 4](#) through [Step 7](#) for each device that you want to use in the pool.

- 9 Exit iManager.
- 10 Continue with [Section 12.4.1, “Creating a Cluster-Enabled Pool and Volume with iManager,”](#) on page 227.

12.3.2 Initializing and Sharing a Device with NSSMU

- 1 Ensure that the SAN device is attached to all nodes in the cluster.
- 2 Log in as the root user to the master node in the cluster, then open a terminal console.
- 3 At the command prompt, enter:

nssmu
- 4 In the NSSMU, select *Devices* and press Enter.
- 5 In the *Devices* list, select the device where you want to create the pool.
- 6 If the device has not been initialized or if you want to delete the current partitioning structures on the device, press F3 to initialize the selected device, then press Y (Yes) to confirm and continue.

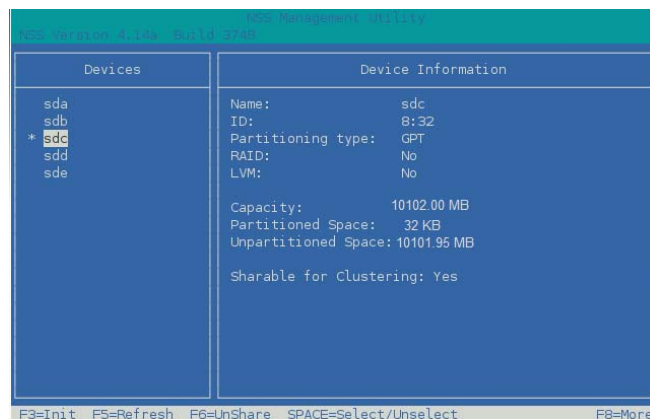
WARNING: Initializing a disk destroys all of the data on it.

- 7 Specify whether to use the GPT or DOS partitioning scheme, then click *OK* to complete the initialization.

GPT can be used for any size device. DOS supports devices up to 2 TB in size.

Wait for the page to refresh before continuing.

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



- 9 Repeat [Step 5](#) through [Step 8](#) for each device that you want to use in the pool.
- 10 Exit NSSMU.
- 11 Continue with [Section 12.4.2, “Creating a Cluster-Enabled Pool and Volume with NSSMU,”](#) on page 230.

12.3.3 Initializing and Sharing a Device with NLVM

- 1 Ensure that the SAN device is attached to all nodes in the cluster.
- 2 Log in as the root user to the master node in the cluster, then open a terminal console.

3 View a list of active devices by entering

```
nlvm list devices
```

4 Identify the device that you want to initialize. If a device contains data, be prepared to lose all data on the device when you initialize it.

In the following example, the `sdd` and `sde` devices are uninitialized and display a format of `None`. The `clus1.sbd` device is a mirrored RAID device that is used for the SBD partition of a cluster named `clus1`. Do not select an SBD device.

```
avalon:~/Desktop # nlvm list devices
sda size=11.00GB free=1.30GB format=MSDOS shared=No RAID=No
sdb size=102.00MB free=0KB format=MSDOS shared=Yes RAID=No
sdc size=102.00MB free=0KB format=MSDOS shared=Yes RAID=No
sdd size=512.00MB free=0KB format=None shared=No RAID=No
sde size=512.00MB free=0KB format=None shared=No RAID=No
clus1.sbd size=99.57MB free=0KB format=None shared=Yes RAID=Yes
```

5 Initialize the device by entering

```
nlvm [--force] [--no-prompt] init <device_name> [format=<gpt|msdos>] shared
```

For command usage information, see “[Init Device](#)” in the [OES 11 SP1: NLVM Reference](#).

You are automatically prompted to confirm the initialize action. Respond Y (Yes) or N (No). Use the `--no-prompt` NLVM option to suppress the confirmation.

Replace *device_name* with the node name of the device to be initialized, such as `sde`. The device name must be the first option after `init`.

Specify `gpt` or `msdos` as the partitioning scheme to use when formatting the device.

The `shared` option marks the device as shareable for clustering after the initialization is complete. A small partition is created on the device to store the shared setting.

For devices that contain data, you can specify the `--force` option to force the initialization if the `init` command cannot delete any pools on the disk.

For example, to initialize a device with the MSDOS partitioning scheme and mark it as shareable for clustering, enter

```
nlvm init sde format=msdos shared
```

6 List details about the device to verify that the device is formatted, and the amount of free space has increased.

```
nlvm list device <device_name>
```

For example, enter

```
nlvm list device sde
```

```
avalon:~/Desktop # nlvm list device sde
sde (8:64) size=512.00MB(1048576) free=511.98MB(1048544) format=MSDOS h:s=255:6
3 shared=Yes RAID=No
```


12.4 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 12.4.1, “Creating a Cluster-Enabled Pool and Volume with iManager,” on page 227](#)
- ♦ [Section 12.4.2, “Creating a Cluster-Enabled Pool and Volume with NSSMU,” on page 230](#)

12.4.1 Creating a Cluster-Enabled Pool and Volume with iManager

- 1 Ensure that the SAN device is attached to all of the nodes in the cluster.
- 2 Log in to iManager as an administrator user.
- 3 If you have not already done so, initialize and share the device you want to use for the clustered pool.

For information, see [Section 12.3, “Initializing and Sharing a SAN Device,” on page 223](#). We recommend one device per pool. If you use multiple devices, each device must be marked as *Shareable for Clustering*.

- 4 Create a clustered pool:

4a In *Roles and Tasks*, select *Storage > Pools*.

4b Browse to select the Cluster object () of the cluster.

Selecting the Cluster object automatically selects the server that is currently the master node in the cluster. The shared device will be assigned to this server and the pool and volume will be created there. After you create the clustered pool and volume, you can modify the resource's preferred node order, then cluster migrate the volume to its most preferred node.

4c Click the *New* link to open the New Pool wizard.

4d Specify the new pool name, then click *Next*.

4e Select the check box next to the shared device where you want to create the pool, then specify the size of the pool. We recommend one that you use one device per pool.

4f Select or deselect the *Mount On Creation* option.

The *Mount On Creation* option determines if the pool you are creating is to be activated (the resource is brought online) as soon as it is created.

The option is enabled by default. This allows you to create volumes immediately after you create the clustered pool. Default settings apply to the pool cluster resource and scripts, but you can modify them later.

If you deselect the option, you can configure the resource and modify the scripts before you bring the resource online for the first time. You must bring the pool cluster resource online before you can create volumes on it.

4g Select the *Cluster Enable on Creation* check box.

This option is selected by default if the device is shared.

If you deselect this option, you are not shown the Cluster Pool Information page. After you create a volume on the non-clustered pool, you can cluster-enable the pool and volume as described in [Section 12.5, “Cluster-Enabling an Existing NSS Pool and Its Volumes,” on page 233](#).

4h On the Cluster Pool Information page, specify the following information:

Parameter	Action
Virtual Server Name	<p>(Optional) The default virtual server name for the resource is the cluster name plus the cluster resource name. For example, if the cluster name is <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> <p>You can use the suggested name, or specify a different Virtual Server name for the cluster resource.</p> <p>You can modify the cluster resource name after the resource has been created by using the <code>cluster rename</code> command. For information, see Section 10.14, "Renaming a Cluster Resource," on page 208. Changing the resource name does not modify the pool name or the virtual server name.</p>
CIFS Server Name	<p>(Optional) If Novell CIFS is installed and running, you can use this field to specify the name of the CIFS virtual server that CIFS clients see when they browse the network.</p> <p>If Novell CIFS is installed and running, but CIFS is disabled as an advertising protocol, this field is not available (dimmed).</p> <p>If Novell CIFS is not installed and running, this field value is <code>NOT_SUPPORTED</code>.</p> <p>CIFS is disabled by default as an advertising protocol. You can select the <i>CIFS</i> check box to enable it. By default, the NCP virtual server name is suggested as the CIFS virtual server name. For more information about how the default name is determined, see Section 12.1.7, "Novell CIFS for Linux," on page 221. You can use the suggested name or specify a customized name for the CIFS virtual server name.</p> <p>If desired, specify a new name for the CIFS virtual server. The name can be up to 15 characters, which is a restriction of the CIFS protocol.</p>
IP Address	<p>Specify an IP address for the pool cluster resource. Tab between the address fields. The address is IPv4 format, such as 10.10.10.243.</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 hosts 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 the Novell CIFS or Novell AFP protocols are not installed and running, selecting the corresponding CIFS or AFP check box has no effect.</p>

Parameter	Action
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 monitor 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 fallback modes, and to configure the preferred nodes.</p>

4i Click *Finish*.

Typically, the pool creation takes less than a minute. 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.

The pool cluster resource should be online and active on the master node if the *Mount on Creation* option was enabled in [Step 4f](#).

5 Create a volume on the clustered pool.

Repeat the following procedure for each cluster volume that you want to create on the shared pool. We recommend using only one volume per shared pool.

5a In iManager, select *Storage*, then select the *Volumes*.

5b Browse to select the Cluster object () of the cluster.

Selecting the Cluster object automatically selects the server that is currently the master node in the cluster.

5c Click *New*.

5d Specify the new volume name, then click *Next*.

Each shared volume in the cluster must have a unique name across all nodes.

5e Select the check box next to the cluster pool where you want to create the volume, select *Allow the volume to grow to the size of the pool*, then click *Next*.

5f Review and change volume attributes by selecting or deselecting the check boxes next to the attributes.

The Backup and Salvage Files attributes are selected by default.

For information about volume attributes, see “[Volume Attributes](#)” in the [OES 11 SP1: NSS File System Administration Guide for Linux](#).

5g Choose whether you want the volume activated and mounted when it is created, then click *Finish*.

Typically, the volume creation takes less than 10 seconds. However, if you have a large tree or the server does not hold an eDirectory replica, the create time can take up to 3 minutes.

6 Verify that the pool cluster resource was created and is online.

6a In iManager, select *Clusters > My Clusters*.

6b Select the cluster.

If the cluster is not in the list, click *Add*, browse to select its Cluster object, then click *OK*.

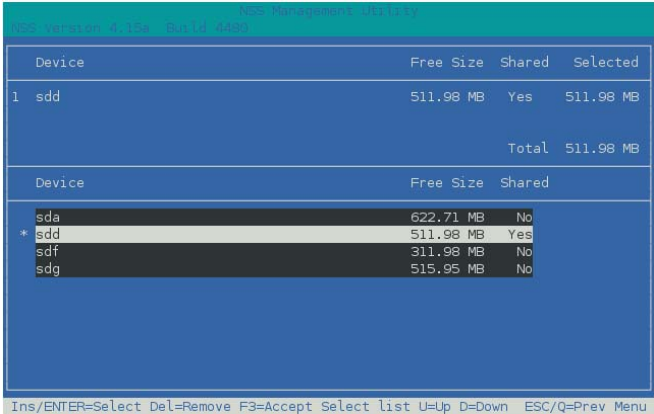
- 6c** On the Cluster Manager page in the Cluster Objects list, view the pool cluster resource, such as POOL1_SERVER.
- If the resource is online, the state is *Running*.
- 7** Continue with [Section 12.6, “Configuring a Load Script for the Shared NSS Pool,”](#) on page 240.

12.4.2 Creating a Cluster-Enabled Pool and Volume with NSSMU

- 1 Ensure that the SAN device is attached to all of the nodes in the cluster.
- 2 On the master node, log in as the root user, then open a terminal console.
- 3 If you have not already done so, initialize and share the device you want to use for the clustered pool.

For information, see [Section 12.3, “Initializing and Sharing a SAN Device,”](#) on page 223. We recommend one device per pool. If you use multiple devices, each device must be marked as *Shareable for Clustering*.
- 4 Start NSSMU by entering `nssmu` at the command prompt.
- 5 Create a clustered pool:
 - 5a** From the NSSMU main menu, select *Pools*.
 - 5b** On the Pools page, press Insert, type a name for the new pool you want to create, then press Enter.
 - 5c** From the list of available devices, select the shared device where you want the pool created (such as `sdc`), then press Enter.

Select the device you marked as shared for this purpose in [Section 12.3, “Initializing and Sharing a SAN Device,”](#) on page 223.
 - 5d** Specify the amount of space (in MB) to use, then press Enter.




Device	Free Size	Shared	Selected
1 sdd	511.98 MB	Yes	511.98 MB
Total			511.98 MB

Device	Free Size	Shared
sda	622.71 MB	No
* sdd	511.98 MB	Yes
sdf	311.98 MB	No
sdg	515.95 MB	No

Ins/ENTER=Select Del=Remove F3=Accept Select list U=Up D=Down ESC/Q=Prev Menu

- 5e** Press F3 to accept the device and size settings.
- The Cluster Pool Information page opens automatically because the selected device is shared.



- 5f** For the *Activate On Creation* option, specify whether you want the pool to be activated when it is created, then continue to the next field by pressing Enter.

The *Activate On Creation* option determines if the pool you are creating is to be activated (the resource is brought online) as soon as it is created.

The option is set to *Yes* (enabled) by default. This allows you to create volumes immediately after you create the clustered pool. Default settings apply to the pool cluster resource and scripts, but you can modify them later.

If you set the value to *No* (disabled), you can configure the resource and modify the scripts before you bring the resource online for the first time. You must bring the pool cluster resource online before you can create volumes on it.

- 5g** Specify *Yes* for the *Cluster Enable on Creation* option.

This option is selected by default if the device is shared.

If you set the value to *No*, the values you set on the Cluster Pool Information page are ignored, and a pool cluster resource is not created. After you create a volume on the non-clustered pool, you can cluster-enable the pool and volume as described in [Section 12.5, “Cluster-Enabling an Existing NSS Pool and Its Volumes,”](#) on page 233.

- 5h** On the Cluster Pool Information page, specify the following information:

Parameter	Action
Virtual Server Name	<p>(Optional) You can use the suggested name or specify a customized name for the NCP Virtual Server name. The default NCP virtual server name for the resource is the cluster name plus the cluster resource name. For example, if the cluster name is <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> <p>To modify the resource’s virtual server name, select the default name, press Enter, specify a new NCP Virtual Server name for the cluster resource, then press Enter.</p> <p>You can modify the cluster resource name after the resource has been created by using the <code>cluster rename</code> command. For information, see Section 10.14, “Renaming a Cluster Resource,” on page 208. Changing the resource name does not modify the pool name or the virtual server name.</p>

Parameter	Action
CIFS Server Name	<p>(Optional) If Novell CIFS is installed and running, you can use this field to specify the name of the CIFS virtual server that CIFS clients see when they browse the network.</p> <p>This field is blank by default if Novell CIFS is installed and running, but CIFS is disabled (set to <i>No</i>) as an advertising protocol.</p> <p>If Novell CIFS is not installed and running, this field value is <code>NOT_SUPPORTED</code>.</p> <p>CIFS is disabled by default as an advertising protocol. To enable CIFS as an advertising protocol, press the down-arrow to go to the CIFS field, then press <i>y</i> (Yes) to enable CIFS. By default, the NCP virtual server name is suggested as the basis for CIFS virtual server name. For more information about how the default name is determined, see Section 12.1.7, “Novell CIFS for Linux,” on page 221. You can use the suggested name or specify a customized name for the CIFS virtual server name.</p> <p>To modify the name, press the up-arrow to go the <i>CIFS Server Name</i> field, press Enter to select the suggested name, specify a new name for the CIFS virtual server, then press Enter. The name can be up to 15 characters, which is a restriction of the CIFS protocol.</p>
IP Address	<p>Press enter to select the field, specify an IP address for the pool cluster resource, then press Enter. Specify the address in IPv4 format, such as 10.10.10.243</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>Press <i>y</i> (Yes) to enable an advertising protocol (AFP, CIFS, NCP) that you want to enable for data requests to this shared pool. Do not disable NCP; NCP is required to support authenticated access to data via the 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 the Novell CIFS or Novell AFP protocols are not installed and running, enabling the corresponding CIFS or AFP option has no effect.</p>

5i Select *Apply* to create and cluster-enable the pool.

Typically, the pool creation takes less than a minute. However, if you have a large tree or the server does not hold an eDirectory replica, the create time can take up to 3 minutes.

6 Press Esc to return to the NSSMU main menu.

7 Create a volume on the clustered pool.

Repeat the following procedure for each cluster volume that you want to create on the shared NSS pool. We recommend using only one volume per shared pool.

7a On the NSSMU main menu, select *Volumes*.

7b On the Volumes page, press Insert, type a name for the new volume you want to create, then press Enter.

Each shared volume in the cluster must have a unique name across all nodes.

7c Specify Y(es) to encrypt volume or N(o) to create a regular volume.

7d From the list of available pools, select the clustered pool where you want the volume to reside, then press Enter.

The new volume appears in the list of volumes.

Typically, the volume creation takes less than 10 seconds. However, if you have a large tree or the server does not hold an eDirectory replica, the create time can take up to 3 minutes.

7e (Optional) From the list of volumes, select the newly created volume, press F8 to view more options, press Enter to open the *Properties* page to review and change volume attributes, then select *Apply* and press Enter to save any changes.

The Backup and Salvage Files attributes are selected by default.

For information about volume attributes, see “[Volume Attributes](#)” in the *OES 11 SP1: NSS File System Administration Guide for Linux*.

8 Exit NSSMU.

9 Verify that the pool cluster resource was created and is online.

9a In iManager, select *Clusters > My Clusters*.

9b Select the cluster.

9c On the Cluster Manager page in the Cluster Objects list, view the pool cluster resource, such as POOL1_SERVER.

If the resource is online, the state is *Running*.

10 Continue with [Section 12.6, “Configuring a Load Script for the Shared NSS Pool,”](#) on page 240.

12.5 Cluster-Enabling an Existing NSS Pool and Its Volumes

Cluster-enabling an existing pool allows it (and its volumes) to be moved or mounted on different servers in the cluster in a manner that supports transparent client reconnect. This makes the data highly available to users.

You might need to cluster enable an existing pool in at least the following situations:

- You are converting stand-alone servers with SAN attached devices to a cluster.
- You are migrating NSS pools between two separate clusters.

In this scenario, you must create a Pool object and Volume objects in the new cluster context by using the Update eDirectory option in NSSMU or the Storage plug-in for iManager.

- You are creating a cluster resource from a storage-based snapshot.

In this scenario, you must create a Pool object and Volume objects for the activated pool snapshot and its volumes by using the Update eDirectory option in NSSMU or the Storage plug-in for iManager.

When you cluster-enable the exiting pool, its pool and volume information will be added automatically to the load, unload, and monitor scripts for the pool cluster resource. In addition, the volume entry is automatically removed from the `/etc/fstab` so that the pool cluster resource load script controls when the volume is mounted.

Before you cluster-enable the pool, we recommend that you deactivate it. The *Online Resource after Creation* option and pool state on the node when you begin determines the resource state and pool state after the resource is created, as described in the following table:

Pool State on the Node	Online Resource After Creation Setting	Resource State after Creation	Pool State after the Resource Is Created
Deactive	Selected	Online	Active as a resource on a preferred node.
Deactive	Deselected	Offline	Deactive. Modify the resource settings and scripts as needed, then bring the resource online on a preferred node.
Active	Deselected	Offline	Active locally on the same node, even if the node is a preferred node. Modify the resource settings and scripts as needed, deactivate the pool locally, then bring the resource online on a preferred node.
Active	Selected	Comatose	Active locally on the same node, even if the node is a preferred node. Take the comatose resource offline. Modify the resource settings and scripts as needed, deactivate the pool locally, then bring the resource online on a preferred node.

The default name of a pool cluster resource is the pool name plus the word “SERVER”, such as `POOL1_SERVER`. You can modify the resource name after the resource has been created by using the `cluster rename` command. For information, see [Section 10.14, “Renaming a Cluster Resource,” on page 208](#). Changing the resource name does not modify the pool name or the virtual server name.

After the pool cluster resource is brought online, users can access the cluster-enabled volume by using the IP address or virtual server name of its pool cluster resource. The default virtual server name of the resource is the cluster name plus the default cluster resource name, such as `CLUSTER1-POOL1-SERVER`. You can specify the name of the resource’s virtual server name when you cluster-enable the pool.

The procedure in this section describes how to enable clustering for an existing pool and its volumes. It assumes the following:

- The pool contains at least one volume.
- The pool name meets the naming conventions supported by Novell Cluster Services. For information, see [“Naming Conventions” on page 219](#).

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 SP1: NSS File System Administration Guide for Linux*.


- The pool and its volumes have Storage objects in the eDirectory tree where you are setting up the cluster.

If the Pool and Volume objects are missing, you can create (or re-create) them in eDirectory by using the *Update eDirectory* option in the Storage plug-in for NSS or in NSSMU. You can do this from any node in the cluster. For information, see [“Updating eDirectory Pool Objects”](#) and [“Updating eDirectory Volume Objects”](#) in the *OES 11 SP1: NSS File System Administration Guide for Linux*.

- The pool resides on a disk that has been enabled as *Sharable for Clustering*. If it is not shared, mark the device as shareable. You can use NSSMU, the `nlvm share` command, or the *Storage > Devices* task in Novell iManager to change the device’s share state.

- ♦ Each pool resource requires its own unique static IP address that is in the same IP subnet as the cluster IP address.
- ♦ To add CIFS as an advertising protocol when you cluster-enable an existing pool, you must install and configure 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.
- ♦ If you use special settings for mounting the volume, you must manually add those settings to the load script before you bring the resource online.

To cluster-enable an existing NSS pool:

- 1 Ensure that the SAN device is attached to all of the nodes in the cluster where you want to run the resource.
- 2 Log in to iManager as a cluster administrator user.
- 3 Deactivate the shared pool on the current server.
 - 3a In *Roles and Tasks*, select *Storage > Pools*.
 - 3b Browse to select the cluster where the pool is currently assigned.
 - 3c In the Pools list, select the pool, then click *Deactivate*.
Wait until the page refreshes and confirms that the pool is deactive.
- 4 Verify that the device is marked as shareable for clustering. If the pool uses multiple devices, you must verify the shared setting of each device.
 - 4a In *Roles and Tasks*, select *Storage > Devices*.
 - 4b Browse to select the cluster where the pool is currently assigned.
 - 4c Select the device to view its Details page.
 - 4d If the *Shareable for Clustering* check box is deselected, select the check box, click *Apply*, read the warning, then click *Continue* to enable sharing.
Marking the device as shareable automatically modifies the Shared attribute on the eDirectory objects for the pools and volumes on the device.
Wait a couple of minutes before continuing to allow the attribute change to synchronize in eDirectory. You can verify the change in the shared state of the pool and its volumes by returning to the *Pools* page and *Volumes* page.
 - 4e Repeat this process for each device that is used by the pool.
- 5 In *Roles and Tasks*, select *Clusters > My Clusters*, select the cluster, then click *Cluster Options*.
- 6 On the Cluster Options page, click the *New* link in the Cluster Objects toolbar.
- 7 On the Resource Type page, specify Pool () as the resource type you want to create by clicking the *Pool* radio button.

My Clusters > clus1.ncs.novell > New Resource

New Resource

Resource Type Select the type of cluster resource to create.

Type

- ☒ Pool
- ☐ Resource
- ☐ Template

<< Back Next >> Cancel

- 8 At the bottom of the Resource Type page, click *Next*.
The Cluster Pool Information page opens.

My Clusters > clus1.ncs.novell > New Resource

New Resource

Cluster Pool Information Cluster-enable a pool.

Pool Name:

Virtual Server Name:

CIFS Server Name:

IP Address:

Advertising Protocols:


- ☐ AFP
- ☐ CIFS
- ☒ NCP

☐ Online Resource after Creation
Only select this option if you have verified that the pool is in an de-active state, otherwise the resource will go comatose during creation.

☒ Define Additional Properties

<< Back Next >> Cancel

- 9 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>clus1</code> and the pool cluster resource name is <code>POOL1_SERVER</code>, then the default virtual server name is <code>CLUS1-POOL1-SERVER</code>.</p>

Parameter	Action
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, which is a restriction of the CIFS protocol.</p> <p>By default, the NCP virtual server name is suggested as the CIFS virtual server name. For more information about how the default name is determined, see Section 12.1.7, “Novell CIFS for Linux,” on page 221. You can use the suggested name or specify a customized name for the CIFS virtual server name.</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>
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 the Novell CIFS or Novell AFP protocols are not installed and running, selecting the corresponding CIFS or AFP check box has no effect.</p>
Online Resource after Creation	<p>The check box is deselected by default.</p> <p>Select this option only if the pool is known to be currently in a deactive state; otherwise the resource goes comatose on creation because it will not automatically deactivate the pool from the local node.</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>

My Clusters > clus1.ncs.novell > New Resource

New Resource ?

Cluster Pool Information Cluster-enable a pool.

Pool Name:

Virtual Server Name:

CIFS Server Name:

IP Address:

Advertising Protocols:

☐ AFP

☐ CIFS

☒ NCP

☐ Online Resource after Creation
Only select this option if you have verified that the pool is in a de-active state, otherwise the resource will go comatose during creation.

☒ Define Additional Properties

<< Back Next >> Cancel

10 At the bottom of the Cluster Pool Information page, click *Next*.

- 11 On the Resource Policies page, configure the resource policies for the start, failover, and failback modes.

For information, see [“Configuring the Start, Failover, and Failback Modes for Cluster Resources” on page 197.](#)

My Clusters > clus1.ncs.novell > New Resource

New Resource ?

Resource Policies 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

<< Back Next >> Cancel

- 12 At the bottom of the Resource Policies page, click *Next*.
- 13 On the Resource Preferred Nodes page, assign the preferred nodes to use for the resource.
- For information, see [“Configuring Preferred Nodes and Failover Order for a Resource” on page 198.](#)

My Clusters > clus1.ncs.novell > New Resource

New Resource ?

Resource Preferred Nodes Assign or unassign nodes and change the node failover order.

Assigned:	Unassigned:
avalon london brighton	newcastle cardiff

<< Back Finish Cancel

- 14 At the bottom of the Resource Preferred Nodes page, click *Finish*.

The pool cluster resource appears in the *Cluster Objects* list on the Cluster Options page, such as POOL1_SERVER.

The load, unload, and monitor scripts are automatically created with information about the pools and volumes. The volume entries are removed from the `/etc/fstab` file so that the resource load script can be used to control when the pools and volumes are mounted. The resource is offline.

- 15 Verify that the pool cluster resource was created and that it is offline.

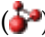
15a In *Roles and Tasks*, select *Clusters > My Clusters*, select the cluster, then click *Cluster Options*.

15b View the pool cluster resource in the Cluster Objects list, such as POOLC1_SERVER.

15c Select *Cluster Manager*, then view the pool cluster resource state as *Offline*.

- 16** Verify the cluster resource scripts for the pool cluster resource, and modify them if desired.
- 16a** In *Roles and Tasks*, select *Clusters > My Clusters*, select the cluster, then click *Cluster Options*.
- 16b** Click the pool cluster resource link to open its Properties page.
- 16c** Configure the following resource settings:

Property	Action
Monitoring	(Optional) Select the <i>Monitoring</i> tab, select <i>Enable Resource Monitoring</i> , and configure monitoring as described in Section 10.7, "Enabling Monitoring and Configuring the Monitor Script," on page 192.
Load script	Select the <i>Scripts</i> tab, select <i>Load Script</i> , then modify the script as described in Section 12.6, "Configuring a Load Script for the Shared NSS Pool," on page 240.
Unload Script	Select the <i>Scripts</i> tab, select <i>Unload Script</i> , then modify the unload script as described in Section 12.7, "Configuring an Unload Script for the Shared NSS Pool," on page 242.
Monitor Script	Select the <i>Scripts</i> tab, select <i>Monitor Script</i> , then modify the unload script as described in Section 10.7, "Enabling Monitoring and Configuring the Monitor Script," on page 192.

- 16d** At the bottom of the page, click *Apply* or *Ok* to save your changes.
- The changes are not applied until the next time the resource is brought online.
- 17** If you did not deactivate the pool before you began, manually deactivate the pool on its currently assigned node.
- You must deactivate the pool from the server before attempting to bring the pool cluster resource online.
- 17a** In *Roles and Tasks*, select *Storage > Pools*.
- 17b** In the *Server* field, browse to select the Cluster object () of the cluster.
- The Cluster object is automatically associated with the node that is currently the master node in the cluster.
- 17c** In the *Pools* list, select the pool, then click *Deactivate*.
- Wait until the page refreshes and confirms that the pool is deactive.
- 18** Bring the resource online:
- 18a** In *Roles and Tasks*, select *Clusters > My Clusters*, select the cluster, then click *Cluster Manager*.
- 18b** Select the check box next to the pool cluster resource, then click *Online* to bring the resource online.
- The pool is activated and its volumes are mounted on the primary preferred node that is configured for the pool cluster resource.
- If the pool is comatose, take the resource offline, then bring it online. If it goes comatose again, check the changes you made to the cluster scripts, then try again.

12.6 Configuring a Load Script for the Shared NSS Pool

A cluster resource load script is automatically generated for the pool when you cluster-enable it. You can modify the script as needed to suit your needs by using iManager. For information about how to access scripts for the cluster resource, see [Section 10.5, “Configuring a Load Script for a Cluster Resource,” on page 190](#).

Newly created pool cluster resources use the following sequence for commands in a load script. The order is reversed in an unload script.

NSS (pool activation and volume mount)

[NFS] (If NFS is used, you manually add the command here.)

Secondary IP address

NCP

[AFP]

[CIFS]

If you modify the load script, you must take the pool cluster resource offline, then bring it online for the changes to take effect.

IMPORTANT: Do not comment out commands that are automatically generated for parameters that define the cluster resource, such as the IP address, pool name, and volume name.

If you need to modify the IP address, administrator credentials, or other attributes of an existing resource, follow the procedure in [Section 8.11, “Moving a Cluster, or Changing the Node IP Addresses, LDAP Servers, or Administrator Credentials for a Cluster,” on page 134](#).

- [Section 12.6.1, “Sample NSS Load Script,” on page 240](#)
- [Section 12.6.2, “Specifying a Customized Mount Point for an NSS Volume,” on page 241](#)
- [Section 12.6.3, “Specifying Customized Mount Options,” on page 241](#)
- [Section 12.6.4, “Loading Novell AFP as an Advertising Protocol,” on page 242](#)
- [Section 12.6.5, “Loading Novell CIFS as an Advertising Protocol,” on page 242](#)

12.6.1 Sample NSS Load Script

The following sample values are used in the NSS load script below:

Variable	Sample Value
Cluster resource’s virtual server name	NCS1-SHPOOL43-SERVER
Resource IP address	10.10.10.43
Pool name	SHPOOL43
Volume name	SHVOL43
Volume ID	252 (valid values are 0 to 254)


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

exit_on_error nss /poolact=SHPOOL43
exit_on_error ncpcon mount SHVOL43=252

exit_on_error add_secondary_ipaddress 10.10.10.43

exit_on_error ncpcon bind --ncpservname=NCS1-SHPOOL43-SERVER
--ipaddress=10.10.10.43

exit 0
```

12.6.2 Specifying a Customized Mount Point for an NSS Volume

The default mount point for NSS volumes is `/media/nss/<volume_name>`, such as `/media/nss/VOL1`. You can specify a customized mount path by adding the `path` option to the `ncpcon mount` command. Ensure that the following requirements are met:

- The *Allow Mount Point to be Renamed* attribute must be enabled on the volume in order for this to work. It is enabled by default.
- The specified mount path must exist on all of the preferred nodes of the pool cluster resource in the cluster. You can create the mount path manually by using the Linux `mkdir` command on each node, or you can add the `mkdir` command to the load script above the `ncpcon mount` command.

To specify a customized mount point, such as `/usr/novell/nss/<volume_name>`, add the `path` option to the `ncpcon mount` command for the volume.

```
# uncomment next line to create the mount path if it does not already exist
# ignore_error mkdir -p </volume_mountpoint>

exit_on_error ncpcon mount <volume_name>=<volume_id>,path=</volume_mntpoint>
```

For example:

```
# create the mount path if it does not already exist
ignore_error mkdir -p </usr/novell/nss/SHVOL43>

exit_on_error ncpcon mount SHVOL43=252,path=/usr/novell/nss/SHVOL43
```

After you modify the load script,

12.6.3 Specifying Customized Mount Options

The load script for an NSS volume assumes that the default name space of LONG is used when mounting the clustered NSS volume. Supported name spaces include LONG, UNIX, DOS, and MAC. For information about name spaces for NSS volumes, see [“Lookup Namespace”](#) in the *OES 11 SP1: NSS File System Administration Guide for Linux*.

You can explicitly specify the desired name space by adding the following switch to the `ncpcon mount` command in the load script:

```
/opt=ns=<long|unix|dos|mac>
```

The options are case-sensitive. In the `ncpcon mount` command, place the `/opt` switch before the volume information.

If you use a name space other than the default for the shared volume, or if you change the name space setting for shared volume by using NSSMU or the Storage plug-in for iManager, you must modify the load script for the pool cluster resource to add the name space to the `ncpcon mount` command for the volume.

For example, to specify the LONG name space, add the `/opt=ns=long` switch as follows:

```
exit_on_error ncpcon mount /opt=ns=long <VOLUMENAME>=<VOLUMEID> /opt=ns=long
```

For example, to specify the Unix name space, add the `/opt=ns=unix` switch as follows:

```
exit_on_error ncpcon mount /opt=ns=unix <VOLUMENAME>=<VOLUMEID>
```

After you modify the load script, you must take the pool cluster resource offline, then bring it online to apply the new name space setting for the volume.

12.6.4 Loading Novell AFP as an Advertising Protocol

If you enable Novell AFP as a advertising protocol for the pool, the following command is automatically added below the `bind` command:

```
exit_on_error cluster_afp.sh add CLUSTER-POOL1-SERVER 10.10.10.43
```

The command uses the NCP virtual server name and the IP address of the pool cluster resource.

12.6.5 Loading Novell CIFS as an Advertising Protocol

If you enable Novell CIFS as an advertising protocol for the pool, the following command is automatically added below the `bind` command:

```
exit_on_error novcifs --add
'--vserver=".cn=CLUSTER-POOL1-SERVER.o=novell.t=AVALON_TREE."'
--ip-addr=10.10.10.43
```

The `--vserver` option uses the fully distinguished name of the NCP virtual server. The `--ip-addr` option uses the IP address of the pool cluster resource.

12.7 Configuring an Unload Script for the Shared NSS Pool

A cluster resource unload script is automatically generated for the pool when you cluster-enable it. You can modify the script as needed to suit your needs by using iManager. For information about how to access the scripts for a cluster resource, see [Section 10.6, “Configuring an Unload Script for a Cluster Resource,” on page 191](#).

Newly created pool cluster resources use the following sequence for commands in an unload script. The order is reversed in a load script.

[CIFS]

[AFP]

NCP

Secondary IP address

[NFS] (If NFS is used, you manually add the command here.)

NSS (pool deactivation, which automatically dismounts the volumes)

If you modify the unload script, you must take the pool cluster resource offline, then bring it online for the changes to take effect.

- ♦ [Section 12.7.1, “Sample NSS Unload Script,” on page 243](#)
- ♦ [Section 12.7.2, “Using Volume Dismount Commands,” on page 243](#)

- [Section 12.7.3, “Unloading Novell AFP as an Advertising Protocol,” on page 243](#)
- [Section 12.7.4, “Unloading Novell CIFS as an Advertising Protocol,” on page 244](#)

12.7.1 Sample NSS Unload Script

The following sample values are used in the NSS unload script below:

Variable	Sample Value
Cluster resource's virtual server name	NCS1-SHPOOL43-SERVER
Resource IP address	10.10.10.43
Pool name	SHPOOL43
Volume name	SHVOL43
Volume ID	252 (valid values are 0 to 254)

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

ignore_error ncpcon unbind --ncpservname=NCS1-SHPOOL43-SERVER --
ipaddress=10.10.10.43

ignore_error del_secondary_ipaddress 10.10.10.43

ignore_error nss /pooldeact=SHPOOL43

exit 0
```

12.7.2 Using Volume Dismount Commands

By default, the pool deactivation automatically dismounts the volumes on the pool. The `ncpcon` dismount commands are not automatically added to the unload script. However, you can add `ncpcon` dismount commands for each volume before the pool deactivation line. For example:

```
ignore_error ncpcon dismount SHVOL43
ignore_error nss /pooldeact=SHPOOL43
```

Adding the `ncpcon` dismount command provides predictable and clean dismount behavior. It allows the volume dismount to occur and be logged independently of the pool deactivation, which can help with troubleshooting dismount issues. However, the `ncpcon` dismount commands are not automatically maintained and updated if you change a volume name or delete a volume. You must modify the unload script yourself.

12.7.3 Unloading Novell AFP as an Advertising Protocol

If Novell AFP is enabled as an advertising protocol for the pool, the following line is automatically added above the unbind command line:

```
ignore_error cluster_afp.sh del CLUSTER-POOL1-SERVER 10.10.10.43
```


12.7.4 Unloading Novell CIFS as an Advertising Protocol

If Novell CIFS is enabled as an advertising protocol for the pool, the following line is automatically added above the unbind command:

```
ignore_error novcifs
--remove '--vserver=".cn=CLUSTER-POOL1-SERVER.o=novell.t=AVALON_TREE."'
--ip-addr=10.10.10.43
```

12.8 Configuring a Monitor Script for the Shared NSS Pool

A cluster resource monitor script is automatically generated for the pool when you cluster-enable it. It is disabled by default. To enable or disable monitoring, see [Section 10.7.2, “Configuring Resource Monitoring,” on page 194](#).

If you modify the unload script, you must take the pool cluster resource offline, then bring it online for the changes to take effect.

- [Section 12.8.1, “Sample NSS Monitor Script,” on page 244](#)
- [Section 12.8.2, “Monitoring Novell AFP as an Advertising Protocol,” on page 245](#)
- [Section 12.8.3, “Monitoring Novell CIFS as an Advertising Protocol,” on page 245](#)

12.8.1 Sample NSS Monitor Script

The following sample values are used in the NSS monitor script below:

Variable	Sample Value
Cluster resource's virtual server name	NCS1-SHPOOL43-SERVER
Resource IP address	10.10.10.43
Pool name	SHPOOL43
Volume name	SHVOL43
Volume ID	252 (valid values are 0 to 254)

The following is a sample monitor script for an NSS pool cluster resource:

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

exit_on_error status_fs /dev/pool/POOL1 /opt/novell/nss/mnt/.pools/POOL1 nsspool
exit_on_error ncpcon volume VOL1
exit_on_error status_secondary_ipaddress 10.10.10.41
exit 0
```

You can add lines to the monitor script that check the status of processes used by clustering. For information, see [Section 10.7.4, “Monitoring Services that Are Critical to Clustering,” on page 196](#).

12.8.2 Monitoring Novell AFP as an Advertising Protocol

If Novell AFP is enabled as an advertising protocol for the pool, the Novell AFP status command (`afpstat`) is automatically added to help keep AFP up and running. It checks to see if the AFP daemon is running.

```
exit_on_error afpstat
```

12.8.3 Monitoring Novell CIFS as an Advertising Protocol

If Novell CIFS is enabled as an advertising protocol for the pool, the Novell CIFS monitoring command is automatically added to help keep CIFS up and running.

```
exit_on_error rcnovell-cifs monitor
```

The command checks to see if the `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:

Pool cluster resources that were created on OES 2 SP3 and earlier clusters might use the CIFS `status` command in the monitor script. It checks the status and immediately invokes the resource's failover settings if the `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".

12.9 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 a cluster administrator user.
- 2 Offline the cluster resource that you want to modify.

For information, see [Section 9.9, "Onlining and Offlining \(Loading and Unloading\) Cluster Resources from a Cluster Node,"](#) on page 152.

- 3 In *Roles and Tasks*, select *Storage*, then go the Volumes and Pools pages to verify that the volumes are not mounted and the pool is deactive.
- 4 In *Roles and Tasks*, select *Clusters > My Clusters*, select the cluster, then select *Cluster Options*.
- 5 On the Cluster Options page, use one of the following methods to open the Cluster Pool Properties page for the pool cluster resource that you want to manage:
 - ♦ Select the check box next to the pool cluster resource, then click *Details*.
 - ♦ Click the *Name* link for the pool cluster resource.

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 SP1: Novell CIFS for Linux Administration Guide*.


- 9 Online the cluster resource. Click *Cluster Manager*, select the check box next to the resource, then click *Online*.

12.10 Adding NFS Export for a Clustered Pool Resource

You can add NFS support for a clustered NSS volume by adding an `exportfs(8)` command to the load script and unload script for a clustered pool resource. Basically, the order in load script should be NSS, NFS, IP, and NCP. The order is reversed in the unload script.

Before you use the `exportfs(8)` command in resource scripts, ensure that you install the `nfs-kernel-server` package on all of the cluster nodes in the resource's preferred nodes list. This contains the support utilities for the kernel `nfsd` daemon. It is not installed by default. You can use the *YaST > Software > Software Management* tool to find and install the package.

Figure 12-1 Using YaST Software Management to Install the `nfs-kernel-server` Package



- ◆ [Section 12.10.1, “Sample NFS Load Script,” on page 247](#)
- ◆ [Section 12.10.2, “Sample NFS Unload Script,” on page 248](#)

12.10.1 Sample NFS Load Script

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

# Activate the NSS pool and mount its volume
exit_on_error nss /poolact=POOL_16
exit_on_error ncpcon mount VOL_161=224

# Start the NFS service
exit_on_error rcnfsserver start

# 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
```


12.10.2 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
```

12.11 Mirroring and Cluster-Enabling Shared NSS Pools and Volumes

- ♦ [Section 12.11.1, “Understanding NSS Mirroring,” on page 248](#)
- ♦ [Section 12.11.2, “Requirements for NSS Mirroring,” on page 250](#)
- ♦ [Section 12.11.3, “Creating and Mirroring NSS Pools on Shared Storage,” on page 250](#)
- ♦ [Section 12.11.4, “Verifying the NSS Mirror Status in the Cluster,” on page 252](#)

12.11.1 Understanding NSS Mirroring

NSS mirroring is a checkpoint-based synchronous mirroring solution. Data blocks are written synchronously to multiple storage devices. It is an alternative to synchronous replication options that are provided by a SAN storage array.

Hardware configuration and placement for NSS mirroring can vary depending on geography and fault tolerance requirements. For example, [Figure 12-2](#) depicts a simple hardware configuration in which one side of the mirrored NSS volume is located in a separate building from the rest of the cluster hardware. If a disaster occurs in one building, data is still safe on the mirrored NSS volume in the other building.

Figure 12-2 *Single Cluster with Mirrored NSS Volumes in Separate Buildings*


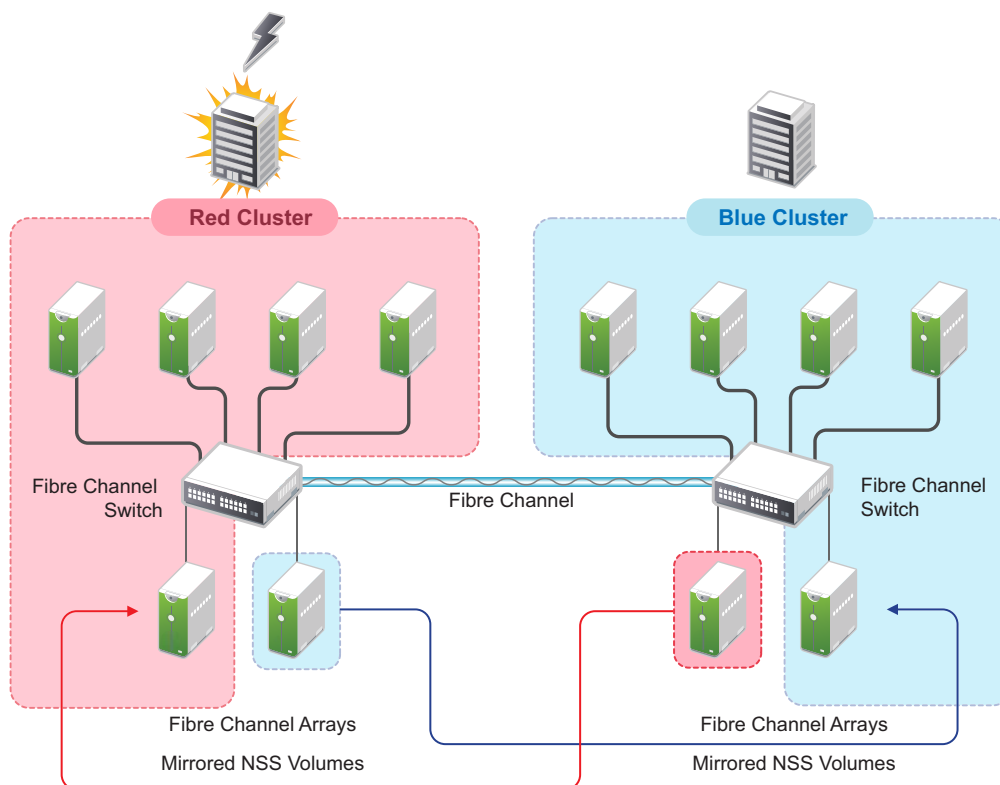


Figure 12-3 depicts a more complex hardware configuration in which two clusters are placed in separate buildings, but each has one side of a mirrored NSS volume in the building with the other cluster. If a disaster occur in one building, data is still safe and immediately accessible on the mirrored NSS volume in the other building.

Figure 12-3 *Two Clusters with Mirrored NSS Volumes in Separate Buildings*



The maximum distance for a single stretch of fibre without optical repeaters to boost distance is about 10 kilometers. There are various devices available that convert Fibre Channel to IP in order to extend SANs to WAN scale.

NSS mirroring is synchronous, and so mirror I/Os must complete before the next I/O can be sent. Performance is a function of time to do an I/O over the longest link.

12.11.2 Requirements for NSS Mirroring

NSS partitions must be mirrored after they are created. If you have an existing partition that you want to mirror, you can either create another partition of equal size on another device to mirror the first partition to, or let the mirroring software automatically create another partition of equal size on another device.

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.

12.11.3 Creating and Mirroring NSS Pools on Shared Storage

Prior to creating and mirroring NSS partitions on shared storage, ensure that you have the following:

- ♦ OES 11 SP1 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 250](#)
- ♦ [“Mirroring the Partition of an Existing NSS Pool” on page 251](#)

Creating a Pool on a Shared Mirrored Device

To create an NSS pool resource on a shared mirrored device:

- 1 Start NSSMU by entering `nssmu` at the command prompt of a cluster server.
- 2 Share the devices.
 - 2a Select *Devices* from the NSSMU main menu.
 - 2b Select the two devices that you want to use, then press F6 to mark the devices as *Sharable for Clustering*.
- 3 Create a mirrored RAID device using the shared devices.
 - 3a Select *RAID Devices* from the NSSMU main menu.
 - 3b Press Insert to create a software RAID, select RAID 1, then press Enter.

- 3c Use the arrow keys to select the two shared devices that you want to contribute space to the RAID, specify the amount of space to use, then press Enter.
The software RAID device automatically inherits the Sharable for Clustering setting.
- 4 Create a clustered NSS pool on the shared RAID device.
 - 4a Select *Pools* from the NSSMU main menu.
 - 4b Press Insert to create a new pool, then create the clustered pool as usual on the shared software RAID device as described in [Section 12.4.2, “Creating a Cluster-Enabled Pool and Volume with NSSMU,”](#) on page 230.

Mirroring the Partition of an Existing NSS Pool

You can use NSSMU to mirror a clustered pool’s partition. The second device must have as much space available as you know the pool’s real allocated size to be, and be marked as Sharable for clustering. To mirror the partition for an existing clustered NSS pool resource:

- 1 Start NSSMU by entering `nssmu` at the command prompt of a cluster server.
- 2 Share the second device.
 - 2a Select *Devices* from the NSSMU main menu.
 - 2b Select the device that you want to use, then press F6 to mark it as *Sharable for Clustering*.
To ensure disaster recovery, the device you select to mirror should be in a different storage array.
- 3 Select *Partitions* from the NSSMU main menu.
- 4 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%.

12.11.4 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.

12.12 Renaming a Clustered NSS Pool

When an NSS pool is cluster-enabled with Novell Cluster Services, special steps must be followed to rename the pool. The pool cluster resource must be taken offline during the renaming process. You can use NSSMU or the Storage plug-in to iManager to rename the shared pool.

When you use an NSS management tool to rename a cluster-enabled pool, it automatically changes the Pool object name in eDirectory, updates the pool name for its Volume objects, and updates the pool name that appears in the load, unload, and monitor scripts for the pool's cluster resource. The rename function does not change the names for the following resource objects:

- ♦ Cluster resource name (such as `POOL1_SERVER`)
You can use the `cluster rename` command to rename the resource. See [Section 10.14, "Renaming a Cluster Resource,"](#) on page 208.
- ♦ Virtual server name (the NCS:NCP Server object, such as `CLUS1-POOL1-SERVER`)
- ♦ CIFS virtual server name (such as `CLUS1-POOL1-W`), if CIFS is enabled as an advertising protocol for the resource

To rename a pool name in a way that also modifies names of resource objects, you can disable clustering for the pool, rename the pool, and then enable clustering for the renamed pool. This method is also described in this section.

Because renaming requires that information is modified in eDirectory, you should rename the pool from the master node. NSS management tools require that the shared pool is active locally on a cluster node that has its NCP Server object in the same context as the Pool object of the pool you want to rename. The server must also be in the resource's *Preferred Nodes* list to migrate to it. To view or modify the list, see [Section 10.9, "Configuring Preferred Nodes and Failover Order for a Resource,"](#) on page 198.

You can use any of the methods in this section to rename a shared pool.

- ♦ [Section 12.12.1, "Renaming a Shared Pool with iManager,"](#) on page 253
- ♦ [Section 12.12.2, "Renaming a Shared Pool with NSSMU,"](#) on page 255
- ♦ [Section 12.12.3, "Renaming a Shared Pool and Its Resource Objects,"](#) on page 258
- ♦ [Section 12.12.4, "Renaming a Shared Pool in a DST Cluster Resource,"](#) on page 262

12.12.1 Renaming a Shared Pool with iManager

- 1 Ensure that the pool cluster resource is running on the master node.

If you need to migrate the resource:

- 1a In iManager, select *Clusters > My Clusters*, then select the cluster that you want to manage.
- 1b On the Cluster Manager page, select the check box next to the pool cluster resource, then click *Migrate*.
- 1c Select the master node, then click *OK*.

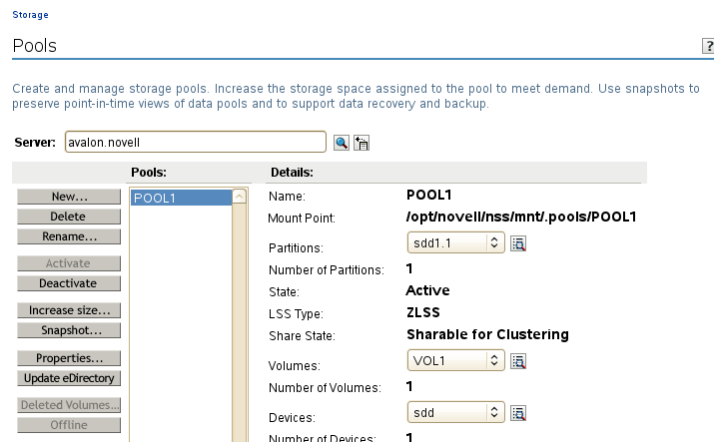
Wait for the resource to report a *Running* state before you continue.

- 2 Take the pool cluster resource offline. On the Cluster Manager page, select the check box next to the pool cluster resource, then click *Offline*.

Wait for the resource to report an *Offline* state before you continue.

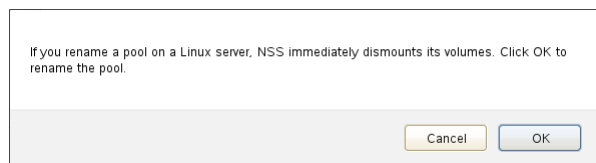
- 3 Activate the pool locally on the cluster node:

- 3a In iManager, select *Storage > Pools*.
- 3b Select the cluster node where you took the resource offline.
- 3c Select the pool, then click *Activate*.

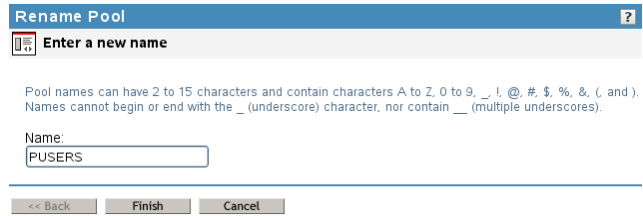


- 4 Rename the shared pool:

- 4a In the *Pools* list, select the pool, then click *Rename*.
- 4b Read the warning message that the pool's volumes will be dismounted, then click *OK* to confirm that you want to rename the pool.

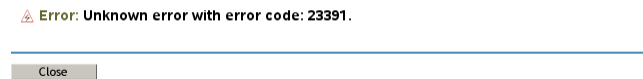


4c In the *Rename Pool* dialog box, type the new name, then click *Finish*.



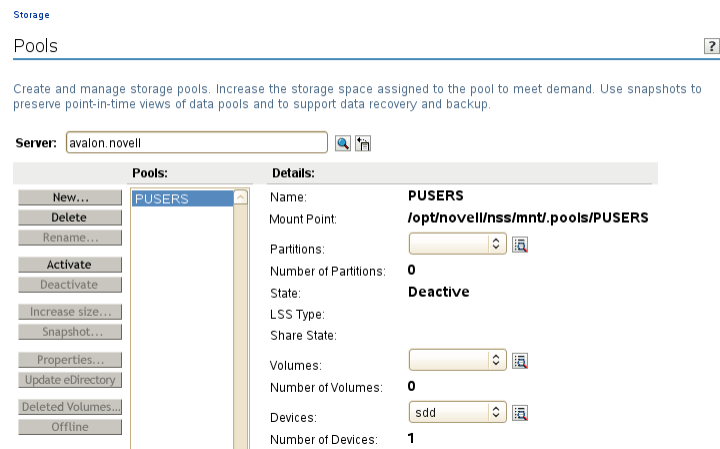
4d Read the error message, then click *OK* to dismiss it.

Error 23391 is `EDIR OBJECT NOT FOUND`. Typically, it means that the old Pool object has been removed successfully and can no longer be found.



4e View the renamed pool in the *Pools* list.

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



5 Verify that the pool name was modified in the pool resource scripts:

5a In iManager, select *Clusters > My Clusters*, then select the cluster that you want to manage.

5b On the Cluster Manager page, click the name link of the pool resource.

The pool resource name has not been modified by the pool rename. Its default name is `<old_pool_name>_SERVER`.

5c On the Resource Properties page, click the *Scripts* tab.

5d Verify that the new pool name is used in the load, unload, and monitor scripts.

Remember that the resource name and virtual server name are not modified by a pool rename. Do not modify those names in scripts.

5d1 On the Load Script page, verify that the pool name was changed in the `nss` command. If you modify the load script, click *Apply*.

```
exit_on_error nss /poolact=<new_pool_name>
```

For example:


```
exit_on_error nss /poolact=PUSERS
```

- 5d2** Click *Unload Script*, then verify that the pool name was changed in the `nss` command. Also check any custom commands where the volume name appears. If you modify the unload script, click *Apply*.

```
exit_on_error nss /pooldeact=<new_pool_name>
```

For example:

```
exit_on_error nss /pooldeact=PUSERS
```

- 5d3** Click *Monitor Script*, then verify that the pool name was changed in the `status_fs` command. Also check any custom commands where the pool name appears. If you modify the monitor script, click *Apply*.

```
exit_on_error status_fs /dev/pool/<new_pool_name> /opt/novell/nss/mnt/.pools/<new_pool_name>
```

For example:

```
exit_on_error status_fs /dev/pool/PUSERS /opt/novell/nss/mnt/.pools/PUSERS
```

- 5e** Click *Cancel* to return to the Cluster Manager page, or click *OK* to apply any changes you made in the scripts.
- 6** Bring the pool cluster resource online:
- 6a** In iManager, select *Clusters > My Clusters*, then select the cluster that you want to manage.
- 6b** On the Cluster Manager page, select the check box next to the pool cluster resource, then click *Online*.
- The resource is brought online on a node in its *Preferred Nodes* list, according to availability. Wait for the resource to report a *Running* state before you continue.
- 7** (Optional) Rename the pool cluster resource. For example, rename `POOL1_SERVER` as `PUSERS_SERVER`.

For instructions, see [Section 10.14, “Renaming a Cluster Resource,”](#) on page 208.

12.12.2 Renaming a Shared Pool with NSSMU

- 1 Log in to the master node as the `root` user, then launch a terminal console.
- 2 Ensure that the pool cluster resource is running on the master node. Enter

```
cluster status
```

If you need to migrate the resource, enter the following at the command prompt:

```
cluster migrate <resource_name> <master_node_name>
```
- 3 Take the pool cluster resource offline. At the command prompt, enter

```
cluster offline <resource_name>
```

You can verify the offline state by entering

```
cluster status
```
- 4 Activate the pool locally on the master node:
 - 4a** At the command prompt, enter

```
nssmu
```


4b Select *Pools* to view a list of pools. The pool state is *Deactive*.

Pools	Pool Information
* POOL1	State: deactive Share State: Sharable for Clustering Snapshot: No Has Snapshot: No

Ins=Create Del=Delete F5=Refresh F7=Activate F9=Deleted Volumes S=Segments

4c Select the pool, then press *F7 Activate*.

Pools	Pool Information
* POOL1	State: active Share State: Sharable for Clustering Snapshot: No Has Snapshot: No Number of Logical Volumes: 1 Capacity: 501.00 MB Free Space: 488.10 MB Used Space: 12.88 MB Purgeable Space: 12 KB Block Size: 4 KB Creation Date: Fri Aug 30 12:40:18 2013 Last Update: Fri Aug 30 12:40:18 2013

Ins=Create Del=Delete F3=Expand F6=Rename F7=Deactivate S=Segments F8=More

5 Rename the pool:

5a On the NSSMU Pools page, select the pool.

5b Press *F6 Rename*.

5c Read the warning message that the pool’s volumes will be dismounted, then click *OK* to confirm that you want to rename the pool.

Pools	Pool Information
* POOL1	State: active Share State: Sharable for Clustering

Warning

This pool contains logical volumes. If you continue this rename operation all of the logical volumes will be dismounted.

Purgeable Space: 12 KB


BL Rename pool?

Cr 0 12:40:18 2013

La 0 12:40:18 2013


Ins=Create Del=Delete F3=Expand F6=Rename F7=Deactivate S=Segments F8=More

5d Specify the new name, then press Enter.




5e Read the error message, then press o (OK) or press Enter to dismiss it.

Error 23391 is EDIR OBJECT NOT FOUND. Typically, it means that the old Pool object has been removed successfully and can no longer be found.



5f After the rename, the pool is in a deactive state.



6 Press Esc twice to exit NISSMU.

7 Verify that the pool name was modified in the pool resource scripts:

7a In iManager, select *Clusters > My Clusters*, then select the cluster that you want to manage.

7b On the Cluster Manager page, click the name link of the pool resource.

The pool resource name has not been modified by the pool rename.

7c On the Resource Properties page, click the *Scripts* tab.

7d Verify that the new pool name is used in the load, unload, and monitor scripts.

Remember that the resource name and virtual server name are not modified by a pool rename. Do not modify those names in scripts.

7d1 On the Load Script page, verify that the pool name was changed in the `nss` command. If you modify the load script, click *Apply*.

```
exit_on_error nss /poolact=<new_pool_name>
```


For example:

```
exit_on_error nss /poolact=PUSERS
```

- 7d2** Click *Unload Script*, then verify that the pool name was changed in the `nss` command. Also check any custom commands where the volume name appears. If you modify the unload script, click *Apply*.

```
exit_on_error nss /pooldeact=<new_pool_name>
```

For example:

```
exit_on_error nss /pooldeact=PUSERS
```

- 7d3** Click *Monitor Script*, then verify that the pool name was changed in the `status_fs` command. Also check any custom commands where the pool name appears. If you modify the monitor script, click *Apply*.

```
exit_on_error status_fs /dev/pool/<new_pool_name> /opt/novell/nss/mnt/.pools/<new_pool_name>
```

For example:

```
exit_on_error status_fs /dev/pool/PUSERS /opt/novell/nss/mnt/.pools/PUSERS
```

- 7e** Click *Cancel* to return to the Cluster Manager page, or click *OK* to apply any changes you made in the scripts.
- 8** Bring the pool cluster resource online. At the command prompt, enter

```
cluster online <resource_name>
```

You can verify the online state by entering

```
cluster status
```

- 9** (Optional) Rename the pool cluster resource. For example, rename `POOL1_SERVER` as `PUSERS_SERVER`.

For instructions, see [Section 10.14, “Renaming a Cluster Resource,”](#) on page 208.

12.12.3 Renaming a Shared Pool and Its Resource Objects

To modify the pool name as well as its related resource objects, you can disable clustering for the pool, rename the pool, and then enable clustering for the renamed pool. Ensure that the pool is on the master node when you disable clustering. The Clusters plug-in in iManager must be used to disable and enable clustering for the pool. The rename procedure can be performed in iManager or NSSMU.

- 1** Disable clustering for the shared pool:

- 1a** Ensure that the pool cluster resource is running on the master node.

If you need to migrate the resource:

- 1a1** In iManager, select *Clusters > My Clusters*, then select the cluster that you want to manage.

- 1a2** On the Cluster Manager page, select the check box next to the pool cluster resource, then click *Migrate*.

- 1a3** Select the master node, then click *OK*.

Wait for the resource to report a *Running* state before you continue.

- 1b** Take the pool cluster resource offline. On the Cluster Manager page, select the check box next to the pool cluster resource, then click *Offline*.

Wait for the resource to report an *Offline* state before you continue.

- 1c** In iManager, use the Clusters plug-in to delete the cluster resource for the shared pool:

1c1 Select *Clusters > My Clusters*, then select the cluster.

1c2 Select the *Cluster Options* tab.

1c3 Select the check box next to the resource, then click *Delete*.

1c4 When you are prompted to confirm the deletion, click *OK* to continue, or click *Cancel* to abort the deletion.

1c5 In the *Tree View* in iManager, browse to verify that the Cluster Resource objects and related objects were removed from eDirectory.

If necessary, you can manually delete the objects. In iManager, go to *Directory Administration > Delete Objects*, select the objects, then click *OK*.

- 1d** In iManager, re-create Storage objects for the pool and its volumes:

These objects are needed by NSS when you rename the pool. They are also used by Novell Cluster Services when you re-create the resource.

1d1 In iManager, select *Storage > Pools*, then select the cluster node where you want to rename the pool.

1d2 In the *Pools* list, select the pool, then click *Activate*.

1d3 Select the pool, then click *Update eDirectory*.

This creates a Pool object in eDirectory with a name format of
<server_name>_<pool_name>_POOL.

1d4 Select *Storage > Volumes*. The server should still be selected.

1d5 In the *Volumes* list, select the volume, then click *Mount*.

1d6 Select the volume, then click *Update eDirectory*.

This creates a Volume object in eDirectory with a name format of
<server_name>_<volume_name>.

1d7 Repeat [Step 4e](#) through [Step 4f](#) for each volume in the pool.

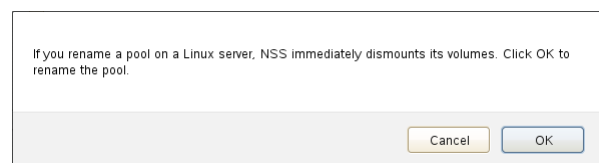
2 Rename the pool.

The following rename procedure uses iManager. You can alternatively use NSSMU to rename the pool.

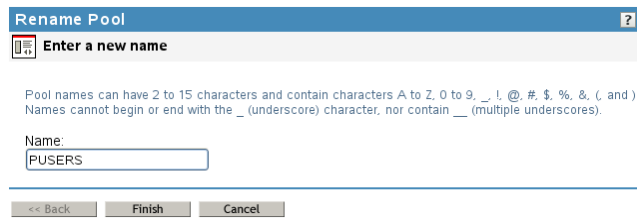
2a In iManager, select *Storage > Pools*. The server should still be selected.

2b In the *Pools* list, select the pool, then click *Rename*.

2c Read the warning message that the pool's volumes will be dismounted, then click *OK* to confirm that you want to rename the pool.



2d In the *Rename Pool* dialog box, type the new name, then click *Finish*.



Rename Pool ?

Enter a new name

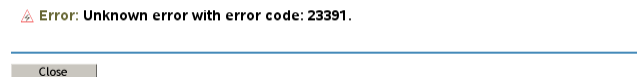
Pool names can have 2 to 15 characters and contain characters A to Z, 0 to 9, _ (underscore), !, @, #, \$, %, &, (, and). Names cannot begin or end with the _ (underscore) character, nor contain __ (multiple underscores).

Name:

<< Back Finish Cancel

2e Read the error message, then click *OK* to dismiss it.

Error 23391 is EDIR OBJECT NOT FOUND. Typically, it means that the old Pool object has been removed successfully and can no longer be found.



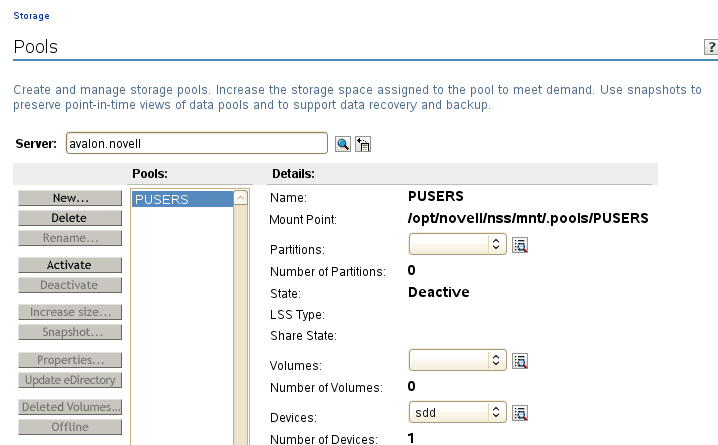
Error: Unknown error with error code: 23391.

Close

2f View the renamed pool in the *Pools* list.

After the rename, the pool is in a deactive state. Do not activate it on any nodes.

If you activate it to verify that it is working as desired with the new name, deactivate it again before you continue.



Storage

Pools ?

Create and manage storage pools. Increase the storage space assigned to the pool to meet demand. Use snapshots to preserve point-in-time views of data pools and to support data recovery and backup.

Server:

Pools:	Details:
<div> New... Delete Rename... Activate Deactivate Increase size... Snapshot... Properties... Update eDirectory Deleted Volumes... Offline </div> <div> PUSERS </div>	Name: PUSERS Mount Point: /opt/novell/nss/mnt/pools/PUSERS Partitions: <input type="text" value=""/> Number of Partitions: 0 State: Deactive LSS Type: Share State: Volumes: <input type="text" value=""/> Number of Volumes: 0 Devices: <input type="text" value="sdd"/> Number of Devices: 1

3 Use the Clusters plug-in to cluster-enable the pool and create new resource objects with the new pool name.

For detailed instructions, see [Step 5](#) thru [Step 18](#) in [Section 12.5, "Cluster-Enabling an Existing NSS Pool and Its Volumes,"](#) on page 233.

3a In *Roles and Tasks*, select *Clusters > My Clusters*, and then select the cluster.

3b Select the *Cluster Options* tab.

3c On the Cluster Options page, click the *New* link in the Cluster Objects toolbar.

3d On the Resource Type page, select the *Pool* radio button, then click *Next*.

The Cluster Pool Information page opens.

My Clusters > clus1.ncs.novell > New Resource

New Resource ?

Cluster Pool Information Cluster-enable a pool.

Pool Name:

Virtual Server Name:

CIFS Server Name:

IP Address:

Advertising Protocols:

☐ AFP

☐ CIFS

☒ NCP

☐ Online Resource after Creation
Only select this option if you have verified that the pool is in an de-active state, otherwise the resource will go comatose during creation.

☒ Define Additional Properties

<< Back Next >> Cancel

3e On the Cluster Pool Information page, specify the following resource setup information, then click *Next*:

- ◆ Pool name

Browse to select the Pool Object of the pool.

- ◆ Virtual server name

A default name is suggested based on the pool name. You can specify a different name for the virtual server.

- ◆ IP address

This is the IP address to use for the resource's virtual server.

- ◆ Advertising protocols

- ◆ NCP (default, mandatory)
- ◆ AFP
- ◆ CIFS

If you enable CIFS, a default name is suggested based on the NCP virtual server name. You can specify a different CIFS Server name.

- ◆ Deselect the *Online Resource after Creation* check box.

- ◆ Select the *Define Additional Properties* check box.

My Clusters > clus1.ncs.novell > New Resource

New Resource ?

Cluster Pool Information Cluster-enable a pool.

Pool Name:

Virtual Server Name:

CIFS Server Name:

IP Address:

Advertising Protocols:

☐ AFP

☐ CIFS

☒ NCP

☐ Online Resource after Creation
Only select this option if you have verified that the pool is in an de-active state, otherwise the resource will go comatose during creation.

☒ Define Additional Properties

<< Back Next >> Cancel

- 3f** On the Resource Policies page, configure the resource policies for the start, failover, and failback modes, then click *Next*.
See [“Configuring the Start, Failover, and Failback Modes for Cluster Resources” on page 197](#).
- 3g** On the Resource Preferred Nodes page, assign the preferred nodes to use for the resource, then click *Finish*.
See [“Configuring Preferred Nodes and Failover Order for a Resource” on page 198](#).
The pool cluster resource appears in the *Cluster Objects* list on the Cluster Options page, such as POOL1_SERVER.
- 3h** (Optional) View and modify the resource properties.
- 3i** (Optional) Enable monitoring and configure the monitoring script.
- 3j** Bring the resource online. Select the Cluster Manager tab, select the check box next to the resource, then click *Online*.
The pool is activated and its volumes are mounted on the primary preferred node that is configured for the pool cluster resource.
If the pool goes comatose, take it offline, check that the pool is deactivated on the local server, then try again.

12.12.4 Renaming a Shared Pool in a DST Cluster Resource

The DST cluster resource is a pool cluster resource for the primary pool that has been modified to manage both the primary pool and secondary pool. When you use an NSS management tool to rename a cluster-enabled pool, it automatically changes the Pool object name in eDirectory, updates the pool name for its Volume objects, and updates the pool name that appears in the load, unload, and monitor scripts for the pool's cluster resource. However, renaming a secondary pool cannot also update its name in the scripts for the primary pool's resource because the secondary pool's default Pool object and Volume object do not contain information about the primary pool. You must manually modify the scripts for the DST cluster resource after you rename a secondary pool.

Renaming the shared primary pool or secondary pool in the DST cluster resource does not affect the settings that define the DST shadow relationship.

The resource objects for a DST cluster resource are typically named when you cluster-enable the primary pool. Renaming a primary pool does not modify the names of the resource objects. To rename a primary pool name in a way that also modifies names of resource objects, you can disable clustering for the primary pool, rename the pool, and then enable clustering for the renamed pool. This process is described in [“Renaming a Shared Pool and Its Resource Objects” \(http://www.novell.com/documentation/oes11/clus_admin_lx/data/rename_cluspool.html#b16vptoo\)](http://www.novell.com/documentation/oes11/clus_admin_lx/data/rename_cluspool.html#b16vptoo). Afterwards, reconfigure the primary pool cluster resource to include the necessary commands for the secondary pool and the DST shadow volume pair, just as you did when you first set up the DST cluster resource. Because the primary volume name and secondary volume name are not affected by the pool rename, you do not need to modify the related commands in the `/etc/opt/novell/ncp2nss.conf` file and `/etc/opt/novell/ncpserv.conf` file.

To rename a shared pool in a DST cluster resource:

- 1** Log in to the master node as the `root` user, then launch a terminal console.
- 2** Ensure that the pool cluster resource is running on the master node. Enter

```
cluster status
```

If you need to migrate the resource, enter the following at the command prompt:


```
cluster migrate <resource_name> <master_node_name>
```

- 3** Take the pool cluster resource offline. At the command prompt, enter

```
cluster offline <resource_name>
```

You can verify the offline state by entering

```
cluster status
```

- 4** Activate the pool to be renamed locally on the master node:

- 4a** At the command prompt, enter

```
nssmu
```

- 4b** Select *Pools* to view a list of pools. The pool state is *Deactive*.

- 4c** Select the pool, then press F7 *Activate*.

- 5** Rename the pool:

- 5a** On the NSSMU Pools page, select the pool.

- 5b** Press F6 *Rename*.

- 5c** Read the warning message that the pool's volumes will be dismounted, then click *OK* to confirm that you want to rename the pool.

- 5d** Specify the new name, then press Enter.

- 5e** Read the error message, then press o (OK) or press Enter to dismiss it.

Error 23391 is EDIR OBJECT NOT FOUND. Typically, it means that the old Pool object has been removed successfully and can no longer be found.

- 5f** After the rename, the pool is in a deactive state.

- 6** Press Esc twice to exit NSSMU.

- 7** Verify that the pool name was modified in the pool resource scripts:

- 7a** In iManager, select *Clusters > My Clusters*, then select the cluster that you want to manage.

- 7b** On the Cluster Manager page, click the name link of the pool resource.

Remember that the pool resource name has not been modified by the pool rename.

- 7c** On the Resource Properties page, click the *Scripts* tab.

- 7d** Verify that the new pool name is used in the load, unload, and monitor scripts.

If you modified the secondary pool name, you must manually modify that line in the script.

Remember that the resource name and virtual server name are not modified by a pool rename. Do not modify those names in scripts.

- 7d1** On the Load Script page, verify that the pool name was changed in the *nss* command. If you modify the load script, click *Apply*.

```
exit_on_error nss /poolact=<new_pool_name>
```

For example:

```
exit_on_error nss /poolact=PUSERS
```

- 7d2** Click *Unload Script*, then verify that the pool name was changed in the *nss* command. Also check any custom commands where the volume name appears. If you modify the unload script, click *Apply*.

```
exit_on_error nss /pooldeact=<new_pool_name>
```

For example:


```
exit_on_error nss /pooldeact=PUSERS
```

- 7d3** Click *Monitor Script*, then verify that the pool name was changed in the `status_fs` command. Also check any custom commands where the pool name appears. If you modify the monitor script, click *Apply*.

```
exit_on_error status_fs /dev/pool/<new_pool_name> /opt/novell/nss/mnt/.pools/<new_pool_name>
```

For example:

```
exit_on_error status_fs /dev/pool/PUSERS /opt/novell/nss/mnt/.pools/PUSERS
```

- 7e** Click *Cancel* to return to the Cluster Manager page, or click *OK* to apply any changes you made in the scripts.
- 8** Bring the pool cluster resource online. At the command prompt, enter

```
cluster online <resource_name>
```

You can verify the online state by entering

```
cluster status
```

12.13 Renaming a Clustered NSS Volume

When an NSS volume is in a cluster-enabled pool, special steps must be followed to rename it. Renaming the volume automatically changes the Volume object name, updates the volume information for the pool, and updates the volume name in the scripts for the pool cluster resource. Custom mount options are preserved in the `ncpcon` mount command. After the volume is renamed, the default mount point changes from `/media/nss/<old_volume_name>` to `/media/nss/<new_volume_name>`.

Because renaming requires that information is modified in eDirectory, you should rename the volume from the master node. NSS management tools require that the shared pool and volume are active locally on a cluster node that has its NCP Server object in the same context as the Pool object and Volume object of the volume you want to rename. The server must also be in the resource's *Preferred Nodes* list to migrate to it. To view or modify the list, see [Section 10.9, "Configuring Preferred Nodes and Failover Order for a Resource,"](#) on page 198.

Before you rename the volume, check the volume properties to ensure that the attribute *Allow mount point to be renamed* has been enabled. If the attribute is disabled, the volume name will be modified and the mount point path is unchanged on the initial remount, making it appear that you have retained the old mount point path. However, for the shared volume, the default mount path of `/media/nss/<volume_name>` is enforced with the new name on subsequent mounts. If your intent is to use a custom mount point path for a shared NSS volume, see [Section 12.14, "Renaming the Mount Point Path for a Clustered NSS Volume,"](#) on page 275.

If the pool cluster resource manages multiple volumes, we recommend that you rename only one volume at a time. Verify that the renamed volume and pool cluster resource are working as expected before you attempt to rename another volume in the shared pool.

The NCP Server, Novell AFP, and Novell CIFS services are integrated with NSS and Cluster Services. The NCP, AFP, and CIFS file access continues to work after the volume is renamed. However, you must ensure that your users have the proper path information for the volume after it is renamed. If you have other protocols and services configured for the volume, you might need to re-configure them to work with the new volume name.

After you rename the NSS volume, you must manually edit some lines in the resource scripts if you use a custom mount point path for the NSS volume. For information about configuring scripts to use a custom mount point path for an NSS volume, see [Section 12.14, “Renaming the Mount Point Path for a Clustered NSS Volume,”](#) on page 275.

If the volume is used as the primary volume or secondary volume in a Dynamic Storage Technology resource script, special steps are required to remove references to the old names as a shadow volume pair. Some information must be removed manually on each node before you bring the modified resource online.

You can use the procedures in the following sections to rename shared volumes:

- ♦ [Section 12.13.1, “Renaming a Shared NSS Volume with iManager,”](#) on page 265
- ♦ [Section 12.13.2, “Renaming a Shared NSS Volume with NSSMU,”](#) on page 269
- ♦ [Section 12.13.3, “Renaming a Shared NSS Volume in a DST Cluster Resource,”](#) on page 273

12.13.1 Renaming a Shared NSS Volume with iManager

- 1 Ensure that the pool cluster resource is running on the master node.

If you need to migrate the resource:

- 1a In iManager, select *Clusters > My Clusters*, then select the cluster that you want to manage.
- 1b On the Cluster Manager page, select the check box next to the pool cluster resource, then click *Migrate*.
- 1c Select the master node, then click *OK*.

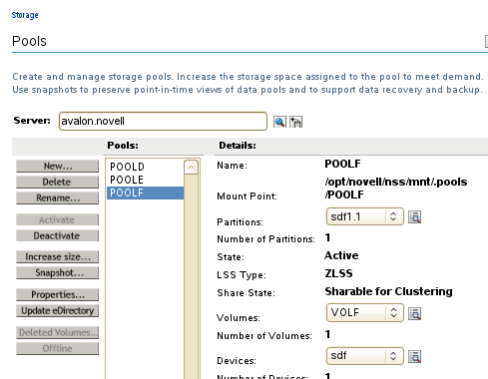
Wait for the resource to report a *Running* state before you continue.

- 2 Take the pool cluster resource offline. On the Cluster Manager page, select the check box next to the pool cluster resource, then click *Offline*.

Wait for the resource to report an *Offline* state before you continue.

- 3 Activate the pool locally on the master node:

- 3a In iManager, select *Storage > Pools*.
- 3b Browse to select the server where you want to activate the pool.
- 3c In the *Pools* list, select the shared pool that contains the volume, then click *Activate*.

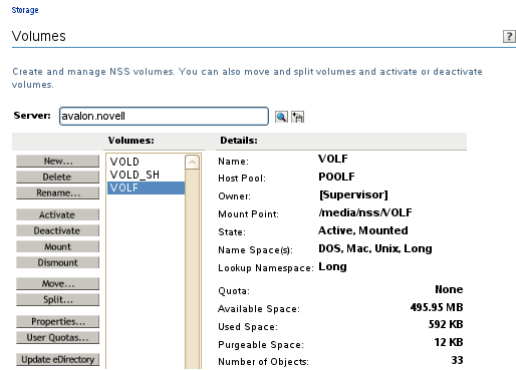


- 4 Mount the volume locally on the master node:

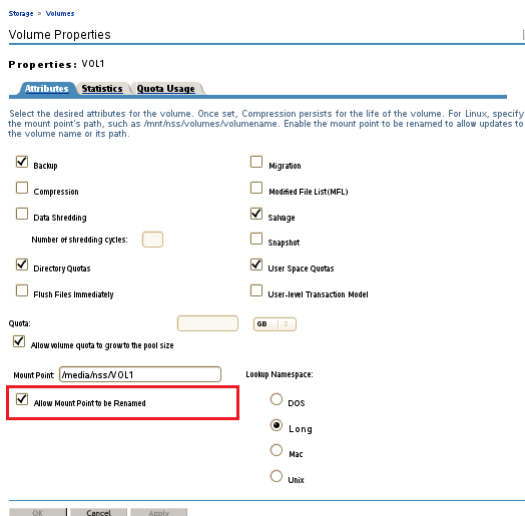
- 4a In iManager, select *Storage > Volumes*.

The server field is automatically populated with the server you chose in the previous step.

4b In the *Volumes* list, select the volume, then click *Mount*.




5 View the volume properties to ensure that the attribute *Allow Mount Point to be Renamed* has been enabled. If it is disabled, select it, then click *OK*.




6 Rename the volume:

6a In the *Volumes* list, select the volume, then click *Rename*.

6b In the *Rename Volume* dialog box, specify the new name, then click *OK*.



6c Wait for the volume to be automatically dismounted and remounted. When the page refreshes, the volume is mounted in the default mount point path `/media/nss/<new_volume_name>`.



If you try to force the page to refresh before the eDirectory changes have been synchronized, you might get an Error 601 eDirectory error. The error condition is temporary and should not prevent the transaction from completing successfully.


Error: -601 - eDirectory Error - Object does not exist.

Close

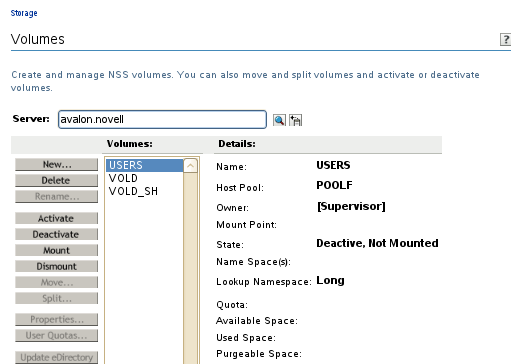
7 Dismount the volume:

7a In the *Volumes* list, select the volume, then click *Dismount*.

7b Read the warning message that open files will be closed by the action, then click *OK* to confirm the dismount.



7c Wait for the volume state to change to *Deactive, Not Mounted*, then continue.




8 Deactivate the pool:


8a In iManager, select *Storage > Pools*.

8b In the *Pools* list, select the pool, then click *Deactivate*.

8c Read the warning message that the pool's volumes will be dismounted, then click *OK* to continue.

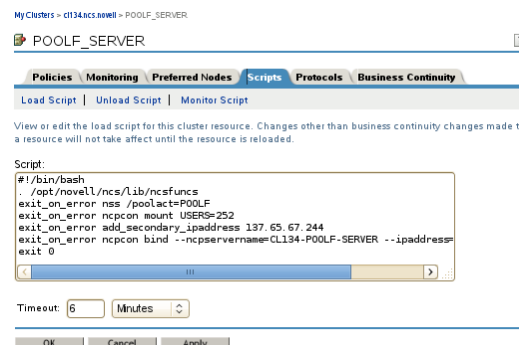


8d When the page refreshes, verify that the pool state is *Deactive*.



9 Verify that the volume name was updated in the scripts for its pool cluster resource:

- 9a** In iManager, select *Clusters > My Clusters*, then select the cluster that you want to manage.
- 9b** On the Cluster Manager page or Cluster Options page, click the name of the pool cluster resource to open the resource properties.
- 9c** Click the *Scripts* tab.



9d On the Load Script page, verify that the volume name was changed in the `ncpcn mount` command. Also check any custom commands where the volume name appears. If you modify the load script, click *Apply*.

```
exit_on_error ncpcn mount <new_volume_name>=<volume_id>
```

For example:

```
exit_on_error ncpcn mount USERS=252
```

9e Click *Unload Script*, then check any custom commands where the volume name appears. If you modify the unload script, click *Apply*.

The unload script does not contain any default volume commands.

- 9f** Click *Monitor Script*, then verify that the volume name was changed in the `ncpcon volume` command. Also check any custom commands where the volume name appears. If you modify the monitor script, click *Apply*.

```
exit_on_error ncpcon volume <new_volume_name>
```

For example:

```
exit_on_error ncpcon volume USERS
```

- 9g** Click *Cancel* to return to the Cluster Manager page, or click *OK* to apply any changes you made in the scripts.
- 10** Bring the pool cluster resource online:
- 10a** In iManager, select *Clusters > My Clusters*, then select the cluster that you want to manage.
- 10b** On the Cluster Manager page, select the check box next to the pool cluster resource, then click *Online*.
- The resource is brought online on a node in its *Preferred Nodes* list, according to availability. Wait for the resource to report a Running state before you continue.

12.13.2 Renaming a Shared NSS Volume with NSSMU

- 1** Log in to the master node as the `root` user, then launch a terminal console.
- 2** Ensure that the pool cluster resource is running on the master node. Enter

```
cluster status
```

If you need to migrate the resource, enter the following at the command prompt:

```
cluster migrate <resource_name> <master_node_name>
```

- 3** Take the pool cluster resource offline. At the command prompt, enter

```
cluster offline <resource_name>
```

You can verify the offline state by entering

```
cluster status
```

- 4** Activate the pool locally on the master node:

- 4a** At the command prompt, enter

```
nssmu
```

- 4b** On the NSSMU *Main Menu*, select *Pools*, then press Enter.

4c On the Pools page, select the pool, then press F7 *Activate*.

Pools	Pool Information
* PUSERS	State: active Share State: Sharable for Clustering Snapshot: No Has Snapshot: No Number of Logical Volumes: 1 Capacity: 501.00 MB Free Space: 488.10 MB Used Space: 12.88 MB Purgeable Space: 12 KB Block Size: 4 KB Creation Date: Fri Aug 30 12:41:24 2013 Last Update: Fri Aug 30 12:41:24 2013

Ins=Create Del=Delete F3=Expand F6=Rename F7=Deactivate S=Segments F8=More

4d Press Esc to return to the NSSMU *Main Menu*.

5 Mount the volume locally on the master node:

5a On the NSSMU *Main Menu*, select *Volumes*, then press Enter.

5b On the Volumes page, select the volume, then press F7 *Mount*.

Logical Volumes	Logical Volume Information
* VOL1	Host Pool: PUSERS Mount Point: /media/nss/VOL1 State: Mounted Supported Name Space(s): Dos Mac Unix Long Name Space: Long Quota: None Free Space: 488.10 MB Used Space: 600 KB Purgeable Space: 12 KB Number of objects: 22 Number of Files: 21 Creation Date: Fri Aug 30 17:33:03 2013 Last Update: Fri Aug 30 17:33:03 2013 Last Archive: Never

Ins=Create Del=Delete F3=Rename F5=Refresh F7=Dismount F8=More

6 On the Volumes page, view the volume properties to ensure that the attribute *Allow renaming mount point* has been enabled. If it is disabled, select the attribute value, press y (Yes) to enable it, then select *Apply* and press Enter.

Volume Name: VOL1	
Backup:	Yes
Compression:	No
Data Shredding:	No
Times to shred data:	1
Directory quotas:	No
Modified File List (MFL):	No
Salvage Files:	Yes
Snapshot - File Level:	No
User Space Restrictions:	No
Flush Files Immediately:	No
Migration:	No
Allow renaming mount point:	Yes
Read Ahead Count in Blocks:	16
Block Size = 4 KB	
Data area for PUSERS is:	0 MB
Volume Quota in (GB):	0 GB
0 = Grow to pool size	
Cancel	Apply

7 Rename the volume:

7a On the Volumes page, select the mounted NSS volume that you want to rename.

7b Press F3 *Rename*, specify the new name, then press Enter.

Logical Volumes	Logical Volume Information
* VOL1	Host Pool: PUSERS Mount Point: /media/nss/VOL1 State: Mounted Supported Name Space(s): Dos Mac Unix Long Name Space: Long Enter new volume name: HOME Used Space: 600 KB Purgeable Space: 12 KB Number of objects: 23 Number of Files: 22 Creation Date: Fri Aug 30 17:47:27 2013 Last Update: Fri Aug 30 17:47:27 2013 Last Archive: Never
Ins=Create Del=Delete F3=Rename F5=Refresh F7=Dismount F8=More	

- 7c** Wait for the volume to be automatically dismounted and remounted. When the page refreshes, the volume is mounted in the default mount point path `/media/nss/<new_volume_name>`.

Logical Volumes	Logical Volume Information
* HOME	Host Pool: PUSERS Mount Point: /media/nss/HOME State: Mounted Supported Name Space(s): Dos Mac Unix Long Name Space: Long Quota: None Free Space: 488.09 MB Used Space: 604 KB Purgeable Space: 12 KB Number of objects: 24 Number of Files: 23 Creation Date: Fri Aug 30 17:47:27 2013 Last Update: Fri Aug 30 17:47:59 2013 Last Archive: Never
Ins=Create Del=Delete F3=Rename F5=Refresh F7=Dismount F8=More	

If you try to force the page to refresh before the eDirectory changes have been synchronized, you might get an eDirectory error. The error condition is temporary and should not prevent the transaction from completing successfully.

8 Dismount the volume:

- 8a** On the NSSMU Volumes page, select the volume, press *F7 Dismount*.
- 8b** Read the warning message that open files will be closed by the action, then press *y (Yes)* to confirm the dismount.

Logical Volumes	Logical Volume Information
* HOME	Host Pool: PUSERS Mount Point: /media/nss/HOME State: Mounted Supported Name Space(s): Dos Mac Unix Long Name Space: Long Used Space: 604 KB Creation Date: Fri Aug 30 17:47:27 2013 Last Update: Fri Aug 30 17:47:59 2013 Last Archive: Never
<div style="background-color: black; color: white; padding: 10px; text-align: center;"> Warning This operation will cause all open files on volume to be closed. Continue dismounting volume? (y/n) y </div>	
Ins=Create Del=Delete F3=Rename F5=Refresh F7=Dismount F8=More	

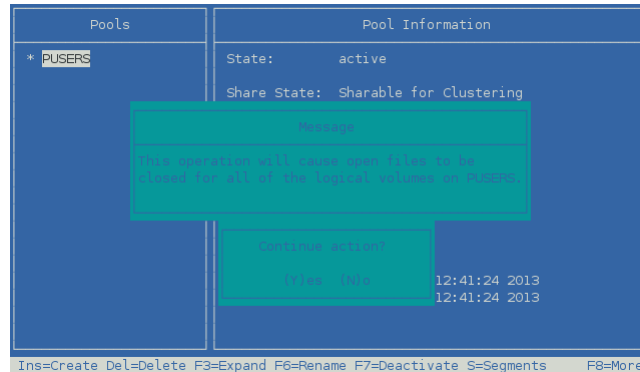
- 8c** Press *Esc* to return to the NSSMU *Main Menu*.

9 Deactivate the pool on the cluster node:

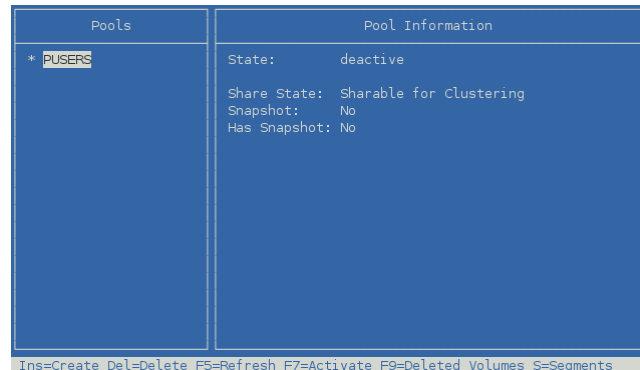
9a On the NSSMU *Main Menu*, select *Pools*, then press Enter.

9b On the Pools page, select the pool, then press F7 *Deactivate*.

9c Read the warning message that the pool's volumes will be dismounted, then press y (Yes) to confirm the deactivation.



9d When the page refreshes, verify that the pool state is *Deactive*.



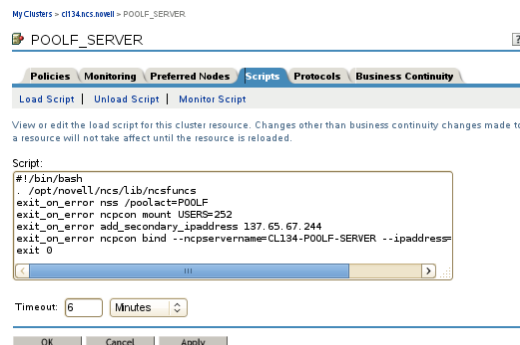
9e Press Esc twice to exit NSSMU.

10 Verify that the volume name was updated in the scripts for its pool cluster resource:

10a In iManager, select *Clusters > My Clusters*, then select the cluster that you want to manage.

10b On the Cluster Manager page or Cluster Options page, click the name of the pool cluster resource to open the resource properties.

10c Click the *Scripts* tab.



- 10d** On the Load Script page, verify that the volume name was changed in the `ncpcon mount` command. Also check any custom commands where the volume name appears. If you modify the load script, click *Apply*.

```
exit_on_error ncpcon mount <new_volume_name>=<volume_id>
```

For example:

```
exit_on_error ncpcon mount USERS=252
```

- 10e** Click *Unload Script*, then check any custom commands where the volume name appears. If you modify the unload script, click *Apply*.

The unload script does not contain any default volume commands.

- 10f** Click *Monitor Script*, then verify that the volume name was changed in the `ncpcon volume` command. Also check any custom commands where the volume name appears. If you modify the monitor script, click *Apply*.

```
exit_on_error ncpcon volume <new_volume_name>
```

For example:

```
exit_on_error ncpcon volume USERS
```

- 10g** Click *Cancel* to return to the Cluster Manager page, or click *OK* to apply any changes you made in the scripts.

- 11** Bring the pool cluster resource online. At the command prompt, enter

```
cluster online <resource_name>
```

You can verify the online state by entering

```
cluster status
```

12.13.3 Renaming a Shared NSS Volume in a DST Cluster Resource

A Dynamic Storage Technology (DST) shadow volume pair consists of two NSS volumes that are managed together. Users see a single merged view of the data stored on the two volumes. This section assumes some understanding of how clustered DST volumes are configured. For more information, see “[Configuring DST Shadow Volume Pairs with Novell Cluster Services](#)” in the *OES 11 SP1: Dynamic Storage Technology Administration Guide*.

The DST cluster resource is a pool cluster resource for the primary pool that has been modified to manage both the primary pool and secondary pool. When you use an NSS management tool to rename a cluster-enabled pool, it automatically changes the Pool object name in eDirectory, updates the pool name for its Volume objects, and updates the pool name that appears in the load, unload, and monitor scripts for the pool’s cluster resource. However, renaming a secondary pool cannot also update its name in the scripts for the primary pool’s resource because the secondary pool’s default Pool object and Volume object do not contain information about the primary pool. You must manually modify the scripts for the DST cluster resource after you rename a secondary pool.

Renaming the shared primary pool or secondary pool in the DST cluster resource does not affect the settings that define the DST shadow relationship.

The resource objects for a DST cluster resource are typically named when you cluster-enable the primary pool. Renaming a primary pool does not modify the names of the resource objects. To rename a primary pool name in a way that also modifies names of resource objects, you can disable clustering for the primary pool, rename the pool, and then enable clustering for the renamed pool. This process is described in “[Renaming a Shared Pool and Its Resource Objects](http://www.novell.com/documentation/oes11/clus_admin_lx/data/rename_cluspool.html#b16vptoo)” (http://www.novell.com/documentation/oes11/clus_admin_lx/data/rename_cluspool.html#b16vptoo).

Afterwards, reconfigure the primary pool cluster resource to include the necessary commands for the secondary pool and the DST shadow volume pair, just as you did when you first set up the DST cluster resource. Because the primary volume name and secondary volume name are not affected by the pool rename, you do not need to modify the related commands in the `/etc/opt/novell/ncp2nss.conf` file and `/etc/opt/novell/ncpserv.conf` file.

To rename a shared pool in a DST cluster resource:

- 1 Log in to the master node as the `root` user, then launch a terminal console.

- 2 Ensure that the pool cluster resource is running on the master node. Enter

```
cluster status
```

If you need to migrate the resource, enter the following at the command prompt:

```
cluster migrate <resource_name> <master_node_name>
```

- 3 Take the pool cluster resource offline. At the command prompt, enter

```
cluster offline <resource_name>
```

You can verify the offline state by entering

```
cluster status
```

- 4 Activate the pool to be renamed locally on the master node:

- 4a At the command prompt, enter

```
nssmu
```

- 4b Select *Pools* to view a list of pools. The pool state is *Deactive*.

- 4c Select the pool, then press F7 *Activate*.

- 5 Rename the pool:

- 5a On the NSSMU Pools page, select the pool.

- 5b Press F6 *Rename*.

- 5c Read the warning message that the pool's volumes will be dismounted, then click *OK* to confirm that you want to rename the pool.

- 5d Specify the new name, then press Enter.

- 5e Read the error message, then press *o* (OK) or press Enter to dismiss it.

Error 23391 is EDIR OBJECT NOT FOUND. Typically, it means that the old Pool object has been removed successfully and can no longer be found.

- 5f After the rename, the pool is in a deactive state.

- 6 Press Esc twice to exit NSSMU.

- 7 Verify that the pool name was modified in the pool resource scripts:

- 7a In iManager, select *Clusters > My Clusters*, then select the cluster that you want to manage.

- 7b On the Cluster Manager page, click the name link of the pool resource.

Remember that the pool resource name has not been modified by the pool rename.

- 7c On the Resource Properties page, click the *Scripts* tab.

- 7d Verify that the new pool name is used in the load, unload, and monitor scripts.

If you modified the secondary pool name, you must manually modify that line in the script.

Remember that the resource name and virtual server name are not modified by a pool rename. Do not modify those names in scripts.

- 7d1** On the Load Script page, verify that the pool name was changed in the `nss` command. If you modify the load script, click *Apply*.

```
exit_on_error nss /poolact=<new_pool_name>
```

For example:

```
exit_on_error nss /poolact=PUSERS
```

- 7d2** Click *Unload Script*, then verify that the pool name was changed in the `nss` command. Also check any custom commands where the volume name appears. If you modify the unload script, click *Apply*.

```
exit_on_error nss /pooldeact=<new_pool_name>
```

For example:

```
exit_on_error nss /pooldeact=PUSERS
```

- 7d3** Click *Monitor Script*, then verify that the pool name was changed in the `status_fs` command. Also check any custom commands where the pool name appears. If you modify the monitor script, click *Apply*.

```
exit_on_error status_fs /dev/pool/<new_pool_name> /opt/novell/nss/mnt/.pools/<new_pool_name>
```

For example:

```
exit_on_error status_fs /dev/pool/PUSERS /opt/novell/nss/mnt/.pools/PUSERS
```

- 7e** Click *Cancel* to return to the Cluster Manager page, or click *OK* to apply any changes you made in the scripts.

- 8** Bring the pool cluster resource online. At the command prompt, enter

```
cluster online <resource_name>
```

You can verify the online state by entering

```
cluster status
```

12.14 Renaming the Mount Point Path for a Clustered NSS Volume

This section describes how to modify the load and unload scripts to specify a non-standard mount point for a clustered NSS volume. The NSS volume attribute *Allow renaming mount point* must be enabled.

For local NSS volumes, the Linux mount point for the device that contains the volume is stored in the `/etc/fstab` file, and the NCP mount point for the volume is stored in the `/etc/opt/novell/ncpserv.conf` file. By design, clustered NSS volumes do not have entries in these files. The mounting and dismounting of clustered volumes is controlled by the load and unload scripts. The mount location is assumed to be the default `/media/nss/<volume_name>`.

In an NSS pool cluster resource's load script, the pool activation command automatically mounts the NSS volume in Linux at the default mount point `/media/nss/<volume_name>`. The `ncpcon mount` command creates an NCP volume instance for the NSS volume and also assumes the default mount point.

If you enable the NSS volume attribute *Allow renaming mount point*, you can modify the load and unload scripts to mount and dismount the NSS volume and NCP instance of the volume in a non-standard location. However, you use a Linux command to dismount the volume from the default location, then mount it in the new location. Use the `path` option in the `ncpcon mount` command to specify the preferred mount point path. Before you bring the resource online by using the modified load script, the customized mount point path must be created manually on all nodes where the resource is allowed to load. You can alternatively add the Linux `mkdir -p` command before the Linux `mount` command to create the path.

After you modify the load and unload scripts to use a non-standard mount location, you should use the `cluster online` command to mount the volume and use the `cluster offline` command to dismount the volume. If you need to mount the volume locally, do not use NSSMU or the Storage plug-in in iManager to mount or dismount the volume. These tools rely on information in `/etc/fstab` file to provide the preferred mount point location, but the clustered volume has no entry there.

- ♦ To mount the volume locally, take the resource offline, activate the pool, then use the Linux `mount` command to mount it in Linux and the `ncpcon mount` command to mount it for NCP. Ensure that you specify the non-standard mount point location in the commands as you do in the load script.
- ♦ To dismount the volume locally, use the `ncpcon dismount` command to dismount the volume from NCP, then use the Linux `umount` command to dismount the volume from Linux, then deactivate the pool. Bring the resource online to mount the volume for the cluster.

The following sample values are used in the NSS pool cluster resource scripts in the following procedure. Replace the values with the actual values for your own clustered pool and volume.

Variable	Sample Value
Cluster resource's virtual server name	CLUS1 - POOL-D-SERVER
Resource IP address	10.10.10.44
Device	/dev/sdd
Pool name	POOL_D
Volume name	VOL_D
Volume ID	252 (valid values are 0 to 254)
Advertising protocols enabled for the resource	NCP (You can add others.)
Default mount point location	/media/nss/VOL_D
New mount point location	/usr/novell/nss/VOL_D

To rename the mount point path for a clustered NSS volume to use a non-standard mount point location:

- 1 Verify that the NSS volume attribute *Allow renaming mount point* is enabled to allow the mount point to be renamed. The attribute is disabled by default.
 - 1a Log in as the `root` user to the server where the NSS pool cluster resource is online.
 - 1b Open a terminal console, then launch NSSMU:


```
nssmu
```
 - 1c In the NSSMU main menu, select *Volumes*, then press Enter.
 - 1d Select the volume in the Volumes list, then press Enter to open its Volume Properties page.

- 1e Use the Down arrow to navigate to and select the value for the *Allow renaming mount point* attribute.

```

Volume Name:  VOL_D

Backup:                               Yes
Compression:                           No
Data Shredding:                         No
    Times to shred data:                 1
Directory quotas:                       No
Modified File List (MFL):               No
Salvage Files:                          Yes
SnapShot - File Level:                  No
User Space Restrictions:                 No
Flush Files Immediately:                 No
Migration:                              No
Allow renaming mount point:              Yes
Read Ahead Count in Blocks:             16
    Block Size = 4 KB
Data area for POOL_D is:                 0 MB
Volume Quota in (GB):                   0 GB
    0 = Grow to pool size

Cancel  Apply
  
```

- 1f If the value is set to no, press y (Yes) to enable it.
 - 1g Use the arrow keys to navigate to and select *Apply*, then press Enter.
 - 1h Press Esc twice to exit NSSMU.
 - 2 Open Novell iManager in a Web browser, then log in as a cluster administrator user.
 - 3 In *Roles and Tasks*, select *Clusters > My Clusters*, select the cluster, then click the *Cluster Manager* tab.
 - 4 On the Cluster Manager page, select the resource's name link to open its Cluster Properties page, then click the *Scripts* tab.
- The Scripts tab automatically displays the load script.
- 5 On the Load Scripts page, view the default load script for the NSS pool cluster resource.

The following sample load script is enabled for NCP as an advertising protocol. The pool name is POOL_D and the volume name is VOL_D.

```

#!/bin/bash

. /opt/novell/ncs/lib/ncsfuns

exit_on_error nss /poolact=POOL_D

exit_on_error ncpcon mount VOL_D=252

exit_on_error add_secondary_ipaddress 10.10.10.44

exit_on_error ncpcon bind --ncpservname=CLUS1-POOL-D-SERVER --ipaddress=10.10.10.44

exit 0
  
```

- 6 Modify the load script to specify the new mount point path:
 - 6a After the pool activation command, add a Linux `umount` command to dismount the NSS volume from Linux.

```

# unmount the NSS volume from Linux (currently mounted at /media/nss/volname)
ignore_error umount -t nssvol VOL_D
  
```


The pool activation automatically mounts the volume on Linux at /media/nss/<volume_name>. You must dismount the volume from Linux before you can remount it for Linux at the new location.

- 6b** After the `umount` command, add a Linux `mkdir` command to create the new mount point path on nodes if it does not exist.

```
# create the non-standard mount point path if it does not exist on the node
ignore_error mkdir -p /usr/novell/nss/VOL_D
```

You can alternatively create the new path by using the `mkdir` command in a terminal console on each node where the resource is allowed to load. If the master node is not the most preferred node, ensure that you make the path before you bring the resource online.

- 6c** After the `mkdir` command, add a Linux `mount` command to mount the NSS volume at the new mount point path.

```
# mount the NSS volume using the Linux mount command
exit_on_error mount -t nssvol VOL_D /usr/novell/nss/VOL_D -o name=VOL_D
```

- 6d** Add the `path` option to the `ncpcon mount` command, and specify the custom mount point path.

```
# add the path command to create the NCP instance at the new location
exit_on_error ncpcon mount VOL_D=252,path=/usr/novell/nss/VOL_D
```

- 6e** Review the modified load script to ensure the correct order of commands:

```
#!/bin/bash

. /opt/novell/ncs/lib/ncsfncs

exit_on_error nss /poolact=POOL_D

# unmount the NSS volume from Linux (currently mounted at /media/nss/volname)
ignore_error umount -t nssvol VOL_D

# create the non-standard mount point path if it does not exist on the node
ignore_error mkdir -p /usr/novell/nss/VOL_D

# mount the NSS volume using the Linux mount command
exit_on_error mount -t nssvol VOL_D /usr/novell/nss/VOL_D -o name=VOL_D

# add the path to create the NCP instance at the new location
exit_on_error ncpcon mount VOL_D=252,path=/usr/novell/nss/VOL_D

exit_on_error add_secondary_ipaddress 10.10.10.44

exit_on_error ncpcon bind --ncpservname=CLUS1-POOL-D-SERVER --ipaddress=10.10.10.44

exit 0
```

- 6f** Click *Apply*.

Wait for the page to refresh before you continue.

- 7** On the Scripts tab, click *Unload Script*.

- 8** On the Unload Script page, view the default unload script for the NSS pool cluster resource.

The following sample unload script is enabled for NCP as an advertising protocol. The pool name is POOL_D and the volume name is VOL_D.

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

ignore_error ncpcon unbind --ncpservername=CLUS1-POOL-D-SERVER --ipaddress=10.10.10.44

ignore_error del_secondary_ipaddress 10.10.10.44

ignore_error nss /pooldeact=POOL_D

exit 0
```

- 9** Modify the unload script to explicitly dismount the volume and NCP instance from the non-standard mount point location:

- 9a** After the `del_secondary_ipaddress` command, add an `ncpcon dismount` command to dismount the NCP instance from the new mount point location.

```
# dismount the NSS volume from NCP
ignore_error ncpcon dismount VOL_D
```

- 9b** After the `ncpcon dismount` command, add a Linux `umount` command to unmount the NSS volume from Linux.

```
# unmount the NSS volume from Linux
ignore_error umount -t nssvol VOL_D
```

- 9c** Review the modified unload script to ensure the correct order of commands:

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

ignore_error ncpcon unbind --ncpservername=CLUS1-POOL-D-SERVER --ipaddress=10.10.10.44

ignore_error del_secondary_ipaddress 10.10.10.44

# dismount the NSS volume from NCP
ignore_error ncpcon dismount VOL_D

# unmount the NSS volume from Linux
ignore_error umount -t nssvol VOL_D

ignore_error nss /pooldeact=POOL_D

exit 0
```

- 9d** Click *Apply*.

Wait for the page to refresh before you continue.

- 10** (Optional) On the Scripts page, click *Monitor Script* and view the monitor script.

No changes are required for the monitor script.

- 11** At the bottom of the *Scripts* page, click *OK* to save your changes and close the Properties page.

The changes do not take effect until the resource is taken offline and brought online.

- 12** Take the resource offline, and bring the resource online in order for the script changes to take effect.

- 12a** On the Cluster Manager page, select the check box next to the resource, then click *Offline*. Wait for the status to report that it is offline, then continue.

- 12b** Select the check box next to the resource, then click *Online*.

12c Verify that the resource comes online and reports a *Running* state.

If the resource goes into a Comatose state, it is probably because you made a mistake in the lines you added or modified in the script. Take the resource offline, then go back to correct the scripts, and try to bring it online again.

13 In NSSMU, verify that the new mount point is used when the pool cluster resource is brought online:

13a Log in as the `root` user to the node that is hosting the resource, then start NSSMU by entering:

```
nssmu
```

13b From the NSSMU main menu, select *Volumes*, then press Enter.

13c In the *Volumes* list, select the clustered NSS volume.

13d View the volume details to verify that the mount point has changed.

For example, the mount point in this example should display `/usr/novell/nss/VOL_D`.

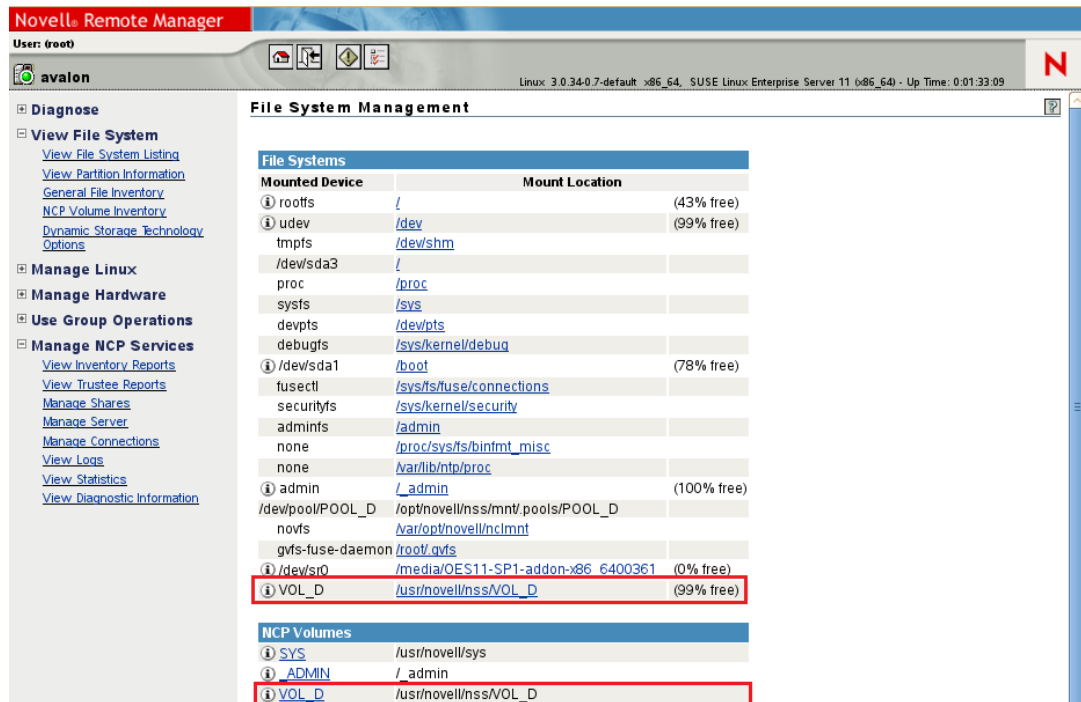
Logical Volumes	Logical Volume Information
* VOL_D	Host Pool: POOL_D Mount Point: /usr/novell/nss/VOL_D State: Mounted Supported Name Space(s): Dos Mac Unix Long Name Space: Long Quota: None Free Space: 486.14 MB Used Space: 604 KB Purgeable Space: 8 KB Number of objects: 24 Number of Files: 23 Creation Date: Thu Aug 23 13:28:50 2012 Last Update: Thu Aug 23 13:32:59 2012 Last Archive: Never
Ins=Create Del=Delete F3=Rename F5=Refresh F7=Dismount F8=More	

13e Press Escape twice to exit NSSMU.

14 In Novell Remote Manager, verify that the new mount point is used for the NSS volume and its NCP instance when the pool cluster resource is brought online:

14a In a Web browser, launch Novell Remote Manager for the node that is hosting the resource, then log in as the `root` user.

14b On the File System Management page, view the mount information for the NSS volume under *File Systems* and its NCP instance under *NCP Volumes*. Note that both volumes are mounted at the new mount point location `/usr/novell/nss/VOL_D`.



14c Exit Novell Remote Manager.

- Open a Linux terminal console as the root user, then list the files in the new mount point location.

```
11 /usr/novell/nss/VOL_D
```

For example, the following is the list for a new NSS volume. You should see the default `~DFSINFO.8-P` file and the hidden `._NETWARE` directory. If AFP is enabled as an advertising protocol, the `DESKTOP.AFP` directory is also present. For an existing volume that contained user data, you should also see additional directories listed.

```
avalon:~/Desktop # ll /usr/novell/nss/VOL_D
total 12
drwxrwxrwx 1 root root 0 Aug 22 22:16 DESKTOP.AFP
-r--r--r-- 1 root root 24 Aug 22 22:16 ~DFSINFO.8-P
d--x--x--x 1 root root 0 Aug 23 00:00 ._NETWARE
```

12.15 Changing the Name Space for a Clustered NSS Volume

The load script for an NSS volume assumes that the default name space of LONG is used when mounting the clustered NSS volume. If you change the name space setting for shared volume by using NSSMU or the Storage plug-in for iManager, you must modify the load script for the pool cluster resource to add the name space to the `ncpcon mount` command for the volume. For information, see [Section 12.6.3, “Specifying Customized Mount Options,” on page 241](#).

- Open Novell iManager in a Web browser, then log in as a cluster administrator user.
- In *Roles and Tasks*, select *Clusters > My Clusters*, select the cluster, then click the *Cluster Manager* tab.
- Modify the mount point path value in the load scripts for the pool cluster resource:
 - On the Cluster Manager page, select the resource’s name link to open its Cluster Properties page, then click the *Scripts* tab.

The Scripts tab automatically displays the load script.

3b Modify the load script:

- 3b1** In the load script, add the `/opt=ns=<long|unix|dos|mac>` option for the `ncpcon mount` command and specify the custom name space to use for the volume. For example:

```
exit_on_error ncpcon mount USERS=254 /opt=ns=unix
```

- 3b2** Click *Apply*, then click *OK* to acknowledge the confirmation message.

3c At the bottom of the page, click *OK* to close the Properties page and save your changes.

The changes do not take effect until the resource is taken offline and brought online.

4 Take the resource offline, and bring the resource online in order for the script changes to take effect.

- 4a** On the Cluster Manager page, select the check box next to the resource, then click *Offline*. Wait for the status to report that it is offline, then continue.

- 4b** Select the check box next to the resource, then click *Online*.

- 4c** Verify that the resource comes online and reports a *Running* state.

If the resource goes into a Comatose state, it is probably because you made a mistake in the lines you added or modified in the script. Take the resource offline, then go back to correct the scripts, and try to bring it online again.

12.16 Adding a Volume to a Clustered Pool

When you use NSSMU, the Storage plug-in to Novell iManager, or the NLVM command line interface to add an NSS volume to a clustered pool, a `mount` command is automatically added to the load script for the pool cluster resource. Offline and online the primary pool cluster resource to apply the modified load script.

If you add a volume to the primary pool for a clustered DST shadow volume, the `mount` command is added twice in the primary pool's cluster load script, once after the primary pool's activation command and once after the secondary pool's activation command. You must manually delete the instance that occurs after the secondary pool's activation, then offline and online the primary pool cluster resource to apply the modified load script.

- 1** In iManager under *Roles and Tasks*, select *Storage > Volumes*.

- 2** Browse to select the Cluster object of the cluster where the pool is currently online.

You can also select the server where the pool cluster resource is currently mounted.

- 3** Click *New* to open the *New Volume* wizard, then create the new volume on the shared pool.

For information, see [“Creating Unencrypted NSS Volumes”](#) in the *OES 11 SP1: NSS File System Administration Guide for Linux*.

- 4** In *Roles and Tasks*, select *Clusters > My Clusters*, then select the cluster you want to manage.

If the cluster you want to manage is not in the list, you can add it. Click *Add*, browse to locate and select the Cluster object of the cluster, then click *OK*.

- 5** If you added a volume to the primary pool for a DST shadow volume, modify the load script to remove the instance of the `mount` command that immediately follows the pool activation line for the secondary pool.

- 5a** Select *Cluster Options*.

- 5b** Click the name link of the pool cluster resource to open its Properties page, then click the *Scripts* tab.

- 5c In the load script, delete the `mount` command for the new volume that follows the pool activation command for the secondary pool.
- 5d Click OK to save your changes.
- 5e Click *Cluster Manager* to return to the Cluster Manager page.
- 6 On the Cluster Manager page in the *Cluster Resource* list, select the check box next to the pool cluster resource, then click *Offline* to take the resource offline.
Wait for the resource to report itself as Offline before you continue.
- 7 In the *Cluster Resource* list, select the check box next to the pool cluster resource, then click *Online* to bring the resource online.
This applies the modified load script, which mounts the newly created volume.

12.17 Changing an Assigned Volume ID

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 onlined to apply the script changes. If you modify the volume ID for a volume in the cluster resource scripts, ensure that you do the following:

- ☐ Volume IDs that you manually assign must be unique across every volume on all servers in cluster.
- ☐ After the value is changed, you must offline and online the cluster resource for the volume in order to mount the volume with its new volume ID.
- ☐ After the volume is mounted with its new ID, the clients must log out and log in to the volume in order to reconnect to the volume with its new volume ID. Automatic reconnection after cluster resource failovers or migrations occurs properly after this one-time reset.

Some clients might cache the volume IDs. To reset the cached value, the client must be rebooted and reconnected to the volume.

- ☐ After the volume is mounted with its new ID, if the backup software is running on a client connection, you might need to restart the backup to reconnect to the volume with its new volume ID. Automatic reconnection after cluster resource failovers or migrations occurs properly after this one-time reset.

12.18 Expanding the Size of a Clustered Pool

NSS pools are made up of segments of space from one or more devices up to a maximum total size of 8 TB (terabytes). Each device can be up to 2 TB in size if a DOS partitioning scheme is used, or up to 8 TB in size if a GPT partitioning scheme is used. You can expand a pool by adding more space from the same device or from a different shared device. If you extend the size of a LUN device that is currently being used by the pool, you must make the extended size visible to all nodes in the cluster before you add the newly available space to the clustered pool. If you add a new shared LUN device, you must make the device visible to all nodes in the cluster before you add space from the device to the clustered pool.

- ◆ [Section 12.18.1, “Planning to Expand a Clustered Pool,” on page 284](#)
- ◆ [Section 12.18.2, “Increasing the Size of a Clustered Pool by Extending the Size of Its LUN,” on page 285](#)
- ◆ [Section 12.18.3, “Increasing the Size of a Clustered Pool by Adding Space from a Different Shared LUN,” on page 291](#)
- ◆ [Section 12.18.4, “Verifying the Expanded Pool Cluster Resource on All Nodes,” on page 295](#)

12.18.1 Planning to Expand a Clustered Pool

Consider the guidelines in this section when you plan to expand the size of a clustered NSS pool.

- ◆ [“Free Space on a Shared Device” on page 284](#)
- ◆ [“Extending a LUN Might Require a Service Outage” on page 284](#)
- ◆ [“Adding a LUN Might Require a Service Outage” on page 285](#)
- ◆ [“Multipath I/O Devices” on page 285](#)

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 `mpathhb`) or device ID.
- ♦ Steps in this section that pertain only to multipath I/O devices are preceded by the keyword "MPIO".

12.18.2 Increasing the Size of a Clustered Pool by Extending the Size of Its LUN

One way to increase the size of a clustered pool is to add space to its LUN device on the SAN storage array, and then increase the size of the pool to use the newly available space. Before you extend the size of a LUN, ensure that you understand the cautions in ["Extending a LUN Might Require a Service Outage" on page 284](#).

- ♦ ["Extending the Size of a LUN" on page 286](#)
- ♦ ["Increasing the Pool Size by Using Free Space on the Same LUN" on page 289](#)

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 `mpathhb`.

- 2 (Offline LUN extension only) If the SAN storage array or vendor-specific storage driver does not support online LUN extension, take the pool cluster resource offline. Open a terminal as the root user, then at the command prompt, enter:

```
cluster offline <resource_name>
```

- 3 On the SAN storage array, use the third-party vendor tools to extend the size of the existing LUN.

Devices can be up to 2 TB in size with the DOS partitioning scheme. Devices can be up to 8 TB in size with the GPT partitioning scheme.

For example, increase the size of the existing LUN from 100 GB to 150 GB.

- 4 Beginning with the active node, perform the following steps on each node in turn in order to update the node's storage map with the new LUN size:

- 4a Log in to the node as the root user, then open a terminal console.

- 4b Scan for devices to recognize the new size for the LUN. At the command prompt, enter

```
/bin/rescan-scsi-bus.sh -forcerescan [--luns=XX]
```

WARNING: In EMC PowerPath environments, do not use the `rescan-scsi-bus.sh` utility provided with the operating system or the HBA vendor scripts for scanning the SCSI buses. To avoid potential file system corruption, EMC requires that you follow the procedure provided in the vendor documentation for EMC PowerPath for Linux.

Use the `--luns=XX` option if you want to scan only for the extended LUN. You can use the `lsscsi(8)` command to locate the LUN number for the device. Information for each device is output on a single line. The first entry is four numbers separated by colons that represent the `Host:Bus:Target:LUN` of the SCSI device.

For information about other `rescan-scsi-bus.sh` options, see “[Scanning for New Devices without Rebooting](http://www.suse.com/documentation/sles11/stor_admin/data/scandev.html)” (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.

The screenshot shows the iManager interface for the 'server1.clusters.context' server. The 'Devices' tab is active, displaying a list of devices: `cciss/c0d0`, `mpatha`, `mpathb` (selected), `mpathc`, `mpathd`, `mpathe`, `mpathf`, `mpathg`, and `mpathh`. The 'Details' pane for the selected `mpathb` device shows the following information:

- Name: `mpathb`
- Major Number: 253
- Minor Number: 1
- ☒ Shareable for Clustering
- MBR Type: ☐ DOS ☒ GPT
- Capacity: 150.00 GB (highlighted with a red box)
- Used Space: 100.00 GB
- Free Space: 50.00 GB
- Pools: POOL1
- Number of Pools: 1
- Partitions: Free - mpathb
- Status: 0% Remirrored, Unknown

At the bottom of the window are buttons for 'OK', 'Cancel', and 'Apply'.

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 289.

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.

Pools:	Details:
<div>New... Delete Rename... Activate Deactivate Increase size... Snapshot... Properties... Update eDirectory Deleted Volumes... Offline</div>	<div><div>POOL1</div><div>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 <div>Total Space: 100.00 GB</div><div>Free Space: 160.00 KB</div><div>Used Space: 100.00 GB</div><div>Purgeable Space: 0.00 GB</div><div>Other in-use space: 100.00 GB</div><div>Block Size: 4.00 KB</div><div>Creation Date: Sep 16, 2010 2:58:27 pm</div><div>Last Update: Sep 16, 2010 2:58:27 pm</div></div></div>

- 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.

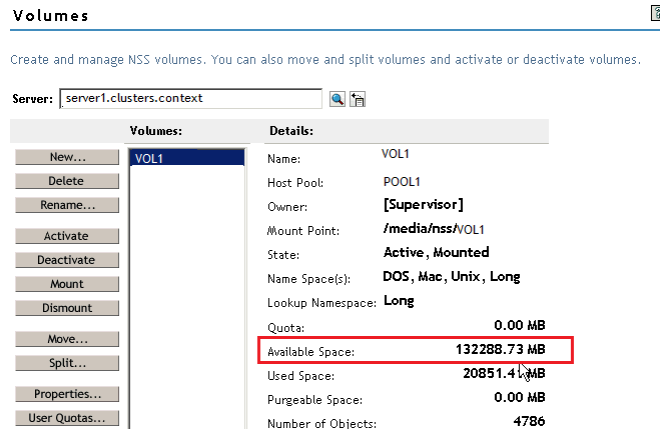
For example, the partitions mpathb1.1 and mpathb1.2 are allocated to POOL1.

Server:

Partitions						
<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 295.](#)

12.18.3 Increasing the Size of a Clustered Pool by Adding Space from a Different Shared LUN

You can increase the size of a clustered pool by adding space to the pool from a different shared LUN device that can be failed over with the resource. Before you add a LUN, ensure that you understand the cautions in [“Adding a LUN Might Require a Service Outage” on page 285.](#)

- ♦ [“Creating and Sharing a New LUN Device” on page 291](#)
- ♦ [“Increasing the Pool Size by Adding Free Space from a Different LUN” on page 294](#)

Creating and Sharing a New LUN Device

If you do not have a shared device that can be failed over with the resource, you can create a new LUN and share it with all nodes in the cluster. If the LUN has multiple I/O paths, special handling is required to add the device to the multipath map on each node. These steps are marked with the keyword “MPIO”.

- 1 Use third-party tools for the SAN storage array to create a new shared LUN device, and assign it to each cluster node.
- 2 Beginning with the active node, perform the following steps on each node in turn in order to update the node’s storage map and multipath map with the new device:
 - 2a Log in to the node as the root user, then open a terminal console.
 - 2b Scan for devices to help the OS recognize the new LUN. At the command prompt, enter:

```
/bin/rescan-scsi-bus.sh -forcerescan
```

WARNING: In EMC PowerPath environments, do not use the `rescan-scsi-bus.sh` utility provided with the operating system or the HBA vendor scripts for scanning the SCSI buses. To avoid potential file system corruption, EMC requires that you follow the procedure provided in the vendor documentation for EMC PowerPath for Linux.

For information about other `rescan-scsi-bus.sh` options, see “Scanning for New Devices without Rebooting” (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	Sharable 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'. The 'Devices' list on the left includes 'cciss/c0d0', 'mpatha', 'mpathb', 'mpathc', 'mpathd', 'mpathe', 'mpathf', 'mpathg', 'mpathh', and 'mpathj', which is currently selected. The 'Details' pane on the right shows the following information for 'mpathj':
Name: mpathj
Major Number: 253
Minor Number: 40
Shareable for Clustering: ☒
MBR Type: ☐ DOS ☒ GPT
Capacity: 25.00 GB
Used Space: 16.00 KB
Free Space: 25.00 GB
Pools: [Empty dropdown]
Number of Pools: [Empty dropdown]
Partitions: Free - mpathj [Empty dropdown]
Status: 0% Remirrored, Unknown
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 294.

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:

Pools:	Details:
<div>New... Delete Rename... Activate Deactivate Increase size... Snapshot... Properties... Update eDirectory Deleted Volumes... Offline</div>	<div><div>POOL1 POOL2 POOL3 POOL4 POOL5 POOL6 POOL7</div><div>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</div></div>

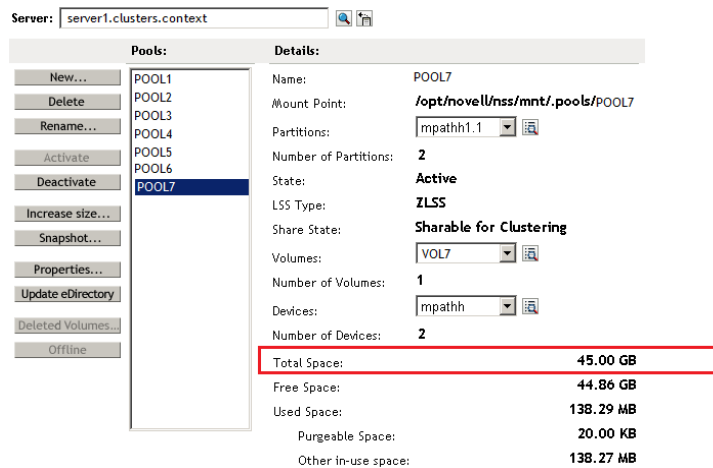
- 3 Increase the size of the pool by adding space from a different LUN:
 - 3a On the Pool page, select the pool, then click *Increase Size*.
The Expand a Pool wizard opens and presents all devices with available free space.
 - 3b In the list of devices, select the check box for the new LUN, such as mpathj, 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"/>	0	mpatha	1015
<input type="checkbox"/>	0	mpathb	2
<input type="checkbox"/>	0	mpathc	1
<input checked="" type="checkbox"/>	25599	mpathj	25599

Pool Size: 46077

- 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, POOL7 has increased in size from 20 GB to about 45 GB.



- 5 For each volume on the pool, if the volume quota is set to *No Quota*, the volume can grow to the size of the pool. Go to *Storage > Volumes*, select the volume, then verify that a volume's *Available Space* field has increased by the amount of space that was added to the pool.
- 6 After the pool size is increased, continue with [Section 12.18.4, "Verifying the Expanded Pool Cluster Resource on All Nodes,"](#) on page 295.

12.18.4 Verifying the Expanded Pool Cluster Resource on All Nodes

Verify that the expanded pool size is reported correctly on each node in the cluster.

- 1 Log in to a node as the root user, then open a terminal console.
- 2 Cluster migrate the pool cluster resource to each node in turn, and use NSSMU to verify that the devices, partitions, pool, and volumes show the correct information on that server:
 - 2a Cluster migrate the pool cluster resource.


```
cluster migrate resource_name node_name
```
 - 2b Launch NSSMU by entering


```
nssmu
```
 - 2c On the Devices page, select the device and verify that the new size is reported.
 - 2d On the Partitions page, select each partition for the device, then verify that the partitions are shown for the expanded pool.
 - 2e On the Pools page, select the pool, then verify the total space available.
 - 2f On the Volumes page, select each volume that has no quota set, then verify that the available space has increased by the amount of space added to its pool.
 - 2g If the information is incorrect, wait a few minutes to allow time for the system to recognize the new partitions, then try again.
- 3 Repeat the previous validation for each node in the cluster.

12.19 Disabling Clustering for a Pool

You can disable clustering for a pool by deleting the resource-related objects, and then re-creating the Pool and Volume objects on a desired node and unsharing the device. We strongly recommend that you delete the resource for a pool from the master node in the cluster.

The Cluster Options page of the Clusters plug-in for iManager provides a *Delete* option that automatically deletes the resource-related objects in eDirectory and updates the cluster information:

- ♦ Cluster-named Pool object
- ♦ Cluster-named Volume object for each of the pool's volumes
- ♦ Cluster Resource object for the pool
- ♦ Virtual server for the cluster resource (NCS:NCP Server object)

Ensure that you offline the cluster resource before attempting to delete either the cluster resource or the clustered pool. For example, if you want to unshare a pool, offline the cluster resource for the pool before you mark the pool or the device as Not Shareable for Clustering, then you can delete the eDirectory object for the cluster resource.

WARNING: If you attempt to delete a cluster resource without first offlining it, deletion errors occur, and the data associated with the clustered pool is not recoverable.

All resource configuration must happen from the master node. On the Cluster Options page for iManager, you are automatically connected to the Cluster object, which is associated with the master node. On the *Storage > Pools* page for iManager, connect to the Cluster object, not to the individual servers. Run NSSMU only on the master node.

This deletes the resource-related objects, but not the storage area they represent.

Use the following procedure to disable clustering for a pool:

- 1 If the resource is on a non-master node in the cluster, migrate it to the master node.

As the *root* user, open a terminal console, then enter

```
cluster migrate <resource_name> <master_node_name>
```

The master node must be in the resource's preferred nodes list. To view or modify the list, see [Section 10.9, "Configuring Preferred Nodes and Failover Order for a Resource,"](#) on page 198.

- 2 Use the `cluster status` command to check the resource status. If the resource is online or comatose, take it offline by using one of the following methods:

- ♦ Enter the following at the command prompt as the *root* user:

```
cluster offline <resource_name>
```

Use the `cluster status` command to verify that the resource has a status of *Offline* before you continue.

- ♦ In iManager, go to *Clusters > My Clusters*, then select the cluster. On the Cluster Manager page, select the check box next to the cluster resource, then click *Offline*.

Refresh the Cluster Manager page to verify that the resource has a status of *Offline* before you continue.

- 3 In iManager, use the Clusters plug-in to delete the cluster resource.

3a Select *Clusters > My Clusters*, then select the cluster.

3b Select the *Cluster Options* tab.

- 3c** Select the check box next to the resource, then click *Delete*.
- 3d** When you are prompted to confirm the deletion, click *OK* to continue, or click *Cancel* to abort the deletion.
- 3e** In the *Tree View* in iManager, browse to verify that the Cluster Resource object and related objects were removed from eDirectory.
- If necessary, you can manually delete the objects. In iManager, go to *Directory Administration* > *Delete Objects*, select the objects, then click *OK*.
- 4** Use the *Update eDirectory* function to re-create Storage objects for the pool and its volumes.
- 4a** In iManager, select *Storage* > *Pools*, select the node where you want the pool to reside as a locally available pool.
- 4b** Select the pool, then click *Activate*.
- 4c** Select the pool, then click *Update eDirectory*.
- This creates a Pool object in eDirectory with a name format of `<server_name>_<pool_name>_POOL`.
- 4d** Select *Storage* > *Volumes*. The server should still be selected.
- 4e** Select the volume, then click *Mount*.
- 4f** Select the volume, then click *Update eDirectory*.
- This creates a Volume object in eDirectory with a name format of `<server_name>_<volume_name>`.
- 4g** Repeat [Step 4e](#) through [Step 4f](#) for each volume in the pool.
- 4h** In the *Tree View* in iManager, browse to verify that the server-named Pool object and Volume object were created.
- For example, the following Tree View shows the Pool object `avalon_POOLD_POOL` and the Volume object `avalon_VOLD` were created for server `avalon`.



- 5 Unshare the devices that contribute space to the pool, then use a third-party SAN management tool to assign the devices to only the desired server.
 - 5a In iManager, go to *Storage > Devices*, then select the desired server (the one you specified in [Step 4](#)).
 - 5b Select the device.
 - 5c Deselect the *Shareable for Clustering* check box, then click Apply.

Unsharing a device fails if the device contains a cluster-enabled pool or split-brain detector (SBD) partition. This is unlikely to be an issue if you used a dedicated device (or devices) for the pool where you have disabled clustering.
 - 5d Repeat these steps for each device that contributes space to the pool.
 - 5e Use a third-party SAN management tool to assign the devices to only the desired server.
- 6 Provide the node's IP address to users who access volumes on the non-clustered pool.

12.20 Deleting and Re-Creating a Pool Cluster Resource (Disabling and Re-Enabling Clustering for a Pool)

You can re-create the resource objects for a pool by deleting the resource-related objects, creating new Pool and Volume objects, and then cluster-enabling the existing pool. This is the same as disabling and re-enabling clustering for a pool. We strongly recommend that you delete the resource for a pool from the master node in the cluster.

The Cluster Options page of the Clusters plug-in for iManager provides a *Delete* option that automatically deletes the resource-related objects in eDirectory and updates the cluster information:

- ♦ Cluster-named Pool object
- ♦ Cluster-named Volume object for each of the pool's volumes
- ♦ Cluster Resource object for the pool
- ♦ Virtual server for the cluster resource (NCS:NCP Server object)

This deletes the resource-related objects, but not the storage area they represent. To delete a resource and create a new one with the same name, you must wait to create the new one until eDirectory synchronizes all of the objects in the tree related to the deleted resource.

Ensure that you offline the cluster resource before attempting to delete either the cluster resource or the clustered pool. For example, if you want to unshare a pool, offline the cluster resource for the pool before you mark the pool or the device as Not Shareable for Clustering, then you can delete the eDirectory object for the cluster resource.

WARNING: If you attempt to delete a cluster resource without first offlining it, deletion errors occur, and the data associated with the clustered pool is not recoverable.

All resource configuration must happen from the master node. On the Cluster Options page for iManager, you are automatically connected to the Cluster object, which is associated with the master node. On the *Storage > Pools* page for iManager, connect to the Cluster object, not to the individual servers. Run NSSMU only on the master node.

Use the following procedure to delete and re-create a cluster resource for a pool:

- 1 If the resource is on a non-master node in the cluster, migrate it to the master node.

As the *root* user, open a terminal console, then enter


```
cluster migrate <resource_name> <master_node_name>
```

The master node must be in the resource's preferred nodes list. To view or modify the list, see [Section 10.9, "Configuring Preferred Nodes and Failover Order for a Resource,"](#) on page 198.

- 2 Use the `cluster status` command to check the resource status. If the resource is online or comatose, take it offline by using one of the following methods:

- ♦ Enter the following at the command prompt as the root user:

```
cluster offline <resource_name>
```

Use the `cluster status` command to verify that the resource has a status of *Offline* before you continue.

- ♦ In iManager, go to *Clusters > My Clusters*, then select the cluster. On the Cluster Manager page, select the check box next to the cluster resource, then click *Offline*.

Refresh the Cluster Manager page to verify that the resource has a status of *Offline* before you continue.

- 3 In iManager, use the Clusters plug-in to delete the cluster resource.

3a Select *Clusters > My Clusters*, then select the cluster.

3b Select the *Cluster Options* tab.

3c Select the check box next to the resource, then click *Delete*.

3d When you are prompted to confirm the deletion, click *OK* to continue, or click *Cancel* to abort the deletion.

3e In the *Tree View* in iManager, browse to verify that the Cluster Resource objects and related objects were removed from eDirectory.

If necessary, you can manually delete the objects. In iManager, go to *Directory Administration > Delete Objects*, select the objects, then click *OK*.

- 4 Use the *Update eDirectory* function to re-create Storage objects for the pool and its volumes.

These objects are needed by Novell Cluster Services when you re-create the resource.

4a In iManager, select *Storage > Pools*, then select the master node if you plan to re-create the storage object, or select the node where you want the pool to reside as a locally available pool.

4b Select the pool, then click *Activate*.

4c Select the pool, then click *Update eDirectory*.

This creates a Pool object in eDirectory with a name format of `<server_name>_<pool_name>_POOL`.

4d Select *Storage > Volumes*. The server should still be selected.

4e Select the volume, then click *Mount*.

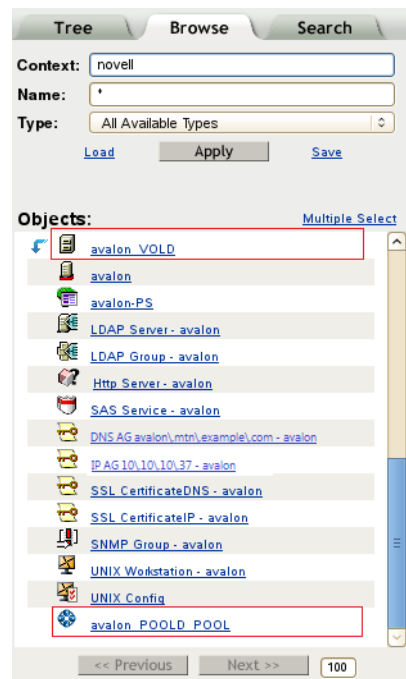
4f Select the volume, then click *Update eDirectory*.

This creates a Volume object in eDirectory with a name format of `<server_name>_<volume_name>`.

4g Repeat [Step 4e](#) through [Step 4f](#) for each volume in the pool.

4h In the *Tree View* in iManager, browse to verify that the Pool object and Volume object were created.

For example, the following Tree View shows the Pool object `avalon_POOLD_POOL` and the Volume object `avalon_VOLD` were created for server `avalon`.



5 Use the Clusters plug-in to cluster-enable the pool:

For detailed instructions, see [Step 5](#) thru [Step 18](#) in [Section 12.5, “Cluster-Enabling an Existing NSS Pool and Its Volumes,”](#) on page 233.

To re-create the cluster resource with the same name, you must wait to create the new one until eDirectory synchronizes all of the objects in the tree related to the deleted resource.

5a In *Roles and Tasks*, select *Clusters > My Clusters*, select the cluster.

5b Select the *Cluster Options* tab.

5c On the Cluster Options page, click the *New* link in the Cluster Objects toolbar.

5d On the Resource Type page, select the *Pool* radio button, then click *Next*.

The Cluster Pool Information page opens.

My Clusters > clust1.ncs.novell > New Resource

New Resource ?

Cluster Pool Information Cluster-enable a pool.

Pool Name:

Virtual Server Name:

CIFS Server Name:

IP Address:

Advertising Protocols:

☐ AFP

☐ CIFS

☒ NCP

☐ Online Resource after Creation
Only select this option if you have verified that the pool is in an de-active state, otherwise the resource will go comatose during creation.

☒ Define Additional Properties

<< Back Next >> Cancel

5e On the Cluster Pool Information page, specify the following information, then click *Next*:

- ♦ Pool name
- ♦ Virtual server name
- ♦ IP address
- ♦ Advertising protocols (NCP, AFP, CIFS)
- ♦ If you enable CIFS, specify the CIFS Server name.
- ♦ Deselect the *Online Resource after Creation* check box.
- ♦ Select the *Define Additional Properties* check box.

My Clusters > clus1.ncs.novell > New Resource

New Resource ?

Cluster Pool Information Cluster-enable a pool.

Pool Name:

Virtual Server Name:

CIFS Server Name:

IP Address:

Advertising Protocols:

☐ AFP

☐ CIFS

☒ NCP

☐ Online Resource after Creation
Only select this option if you have verified that the pool is in a de-active state, otherwise the resource will go comatose during creation.

☒ Define Additional Properties

<< Back Next >> Cancel

5f On the Resource Policies page, configure the resource policies for the start, failover, and failback modes, then click *Next*.

For information, see [“Configuring the Start, Failover, and Failback Modes for Cluster Resources” on page 197](#).

5g On the Resource Preferred Nodes page, assign the preferred nodes to use for the resource, then click *Finish*.

For information, see [“Configuring Preferred Nodes and Failover Order for a Resource” on page 198](#).

The pool cluster resource appears in the *Cluster Objects* list on the Cluster Options page, such as POOL1_SERVER.

5h (Optional) View the resource properties, and enable monitoring and configure the monitoring script.

5i Bring the resource online. Select the Cluster Manager tab, select the check box next to the resource, then click *Online*.

The pool is activated and its volumes are mounted on the primary preferred node that is configured for the pool cluster resource.

If the pool goes comatose, take it offline, check that the pool is deactivated on the local server, then try again.

12.21 Deleting a Clustered Pool

We strongly recommend that you delete a cluster-enabled pool only from the master node in the cluster. This allows the cluster information to be automatically updated.

WARNING: Deleting a pool destroys all data on it.

NSS management tools delete the cluster-enabled pool as well as the related objects in eDirectory:

- ♦ pool and its volumes from the file system and from NCP
- ♦ Cluster-named Pool object
- ♦ Cluster-named Volume object for each of the pool's volumes
- ♦ Cluster Resource object for the pool
- ♦ Virtual server for the cluster resource (NCS:NCP Server object)

When the pool resides on the master node, the cluster information is automatically updated.

When the pool resides on a non-master node, additional steps are required to update the cluster information. A cluster restart might be needed to force the information to be updated.

Use the following procedure to delete a cluster-enabled pool:

- ♦ [Section 12.21.1, “Deleting a Clustered Pool on the Master Node,” on page 302](#)
- ♦ [Section 12.21.2, “Deleting a Clustered Pool on a Non-Master Node,” on page 303](#)

12.21.1 Deleting a Clustered Pool on the Master Node

- 1 If the pool cluster resource is on a non-master node in the cluster, migrate it to the master node. As the root user, open a terminal console, then enter

```
cluster migrate <resource_name> <master_node_name>
```

To migrate the resource, the master node must be in the resource's preferred nodes list. To view or modify the list, see [Section 10.9, “Configuring Preferred Nodes and Failover Order for a Resource,” on page 198](#).

- 2 Use the `cluster status` command to check the resource status. If the resource is online or comatose, take it offline by using one of the following methods:

- ♦ As the root user, enter

```
cluster offline <resource_name>
```

Use the `cluster status` command to verify that the resource has a status of *Offline* before you continue.

- ♦ In iManager, go to *Clusters > My Clusters*, then select the cluster. On the Cluster Manager page, select the check box next to the cluster resource, then click *Offline*.

Refresh the Cluster Manager page to verify that the resource has a status of *Offline* before you continue.

- 3 Delete the pool on the master node by using NSSMU.

You can alternatively use the Storage plug-in in iManager or the `nlvm delete pool <pool_name>` command.

3a In NSSMU, select *Pools*, then press Enter.

3b Select the deactive pool, then press Delete.

3c Select OK to confirm, then press Enter.

- 4 In the *Tree View* in iManager, browse the objects to verify that the following objects were deleted:

- ♦ Pool object
- ♦ Volume object (for each volume in the pool)

- ♦ Pool cluster resource object (from the Cluster container)
 - ♦ Virtual server for the resource (NCS:NCP Server object)
- 5 (Optional) Unshare the device:
- 5a In iManager, go to *Storage > Devices*, then select the node where you want the unshared device to reside.
 - 5b Select the device.
 - 5c Deselect the *Shareable for Clustering* check box, then click *Apply*.
Unsharing a device fails if the device contains a cluster-enabled pool or split-brain detector (SBD) partition. This is unlikely to be an issue if you used a dedicated device (or devices) for the cluster-enabled pool you deleted.
 - 5d Repeat [Step 5b](#) to [Step 5c](#) for each device that contributes space to the pool.
 - 5e Use a third-party SAN management tool to assign the devices to only the desired server.

12.21.2 Deleting a Clustered Pool on a Non-Master Node

- 1 Log in as the root user to the non-master node where the cluster resource currently resides, then open a terminal console.
- 2 Use the `cluster status` command to check the resource status. If the resource is online or comatose, take it offline by using one of the following methods:


```
cluster offline <resource_name>
```


Use the `cluster status` command to verify that the resource has a status of *Offline* before you continue.
- 3 At the command prompt on the non-master node, enter


```
/opt/novell/ncs/bin/ncs-configd.py -init
```
- 4 Look at the file `/var/opt/novell/ncs/resource-priority.conf` to verify that it has the same information (REVISION and NUMRESOURCES) as the file on the master node.
- 5 Delete the pool on the non-master node by using NSSMU.
You can alternatively use the Storage plug-in in iManager or the `nlvm delete pool <pool_name>` command.
 - 5a In NSSMU, select *Pools*, then press Enter.
 - 5b Select the pool, then press Delete.
 - 5c Select OK to confirm, then press Enter.
- 6 In the *Tree View* in iManager, browse the objects to verify that the following objects were deleted:
 - ♦ Pool object
 - ♦ Volume object (for each volume in the pool)
 - ♦ Pool cluster resource object (from the Cluster container)
 - ♦ Virtual server for the resource (NCS:NCP Server object)
- 7 On the master node, log in as the root user, open a terminal console, then enter


```
/opt/novell/ncs/bin/ncs-configd.py -init
```
- 8 Look at the file `/var/opt/novell/ncs/resource-priority.conf` to verify that it has the same information (REVISION and NUMRESOURCES) as that of the non-master node where you deleted the cluster resource.
- 9 In iManager, select *Clusters > My Clusters*, select the cluster, then select the *Cluster Options* tab.

10 Click *Properties*, select the *Priorities* tab, then click *Apply* on the *Priorities* page.

11 At a command prompt, enter

```
cluster view
```

The cluster view should be consistent.

12 Look at the file `/var/opt/novell/ncs/resource-priority.conf` on the master node to verify that the revision number increased.

If the revision number increased, skip [Step 13](#).

If the deleted resource is the only one in the cluster, the priority won't force the update. A phantom resource might appear in the interface. You need to restart Cluster Services to force the update, which also removes the phantom resource.

13 If the revision number did not automatically update in the previous steps, restart 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
```

14 (Optional) Unshare the device:

14a In iManager, go to *Storage > Devices*, then select the node where you want the unshared device to reside.

14b Select the device.

14c Deselect the *Shareable for Clustering* check box, then click *Apply*.

Unsharing a device fails if the device contains a cluster-enabled pool or split-brain detector (SBD) partition. This is unlikely to be an issue if you used a dedicated device (or devices) for the cluster-enabled pool you deleted.

14d Repeat [Step 14b](#) to [Step 14c](#) for each device that contributes space to the pool.

14e Use a third-party SAN management tool to assign the devices to only the desired server.

13 Configuring and Managing Cluster Resources for Shared LVM Volume Groups

After you have installed and configured 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 Btrfs, Ext2, Ext3, ReiserFS, and XFS. This section describes how to configure and manage an LVM volume group cluster resource.

- ♦ [Section 13.1, “Requirements for Creating LVM Cluster Resources,” on page 305](#)
- ♦ [Section 13.2, “Initializing a SAN Device,” on page 308](#)
- ♦ [Section 13.3, “Configuring an LVM Volume Group Cluster Resource with NSS Management Tools,” on page 312](#)
- ♦ [Section 13.4, “Configuring an LVM Cluster Resource with LVM Commands and the Generic File System Template,” on page 330](#)
- ♦ [Section 13.5, “Creating a Virtual Server Object for an LVM Volume Group Cluster Resource,” on page 343](#)
- ♦ [Section 13.6, “Enabling NCP File Access for a Clustered LVM Volume,” on page 348](#)
- ♦ [Section 13.7, “Modifying the LVM Resource Scripts for Other Services,” on page 355](#)
- ♦ [Section 13.8, “Renaming the Mount Point Path for a Clustered LVM Volume,” on page 356](#)
- ♦ [Section 13.9, “Disabling Clustering for an LVM Volume Group and Logical Volume,” on page 357](#)
- ♦ [Section 13.10, “Deleting a Clustered LVM Volume Group and Logical Volume,” on page 362](#)
- ♦ [Section 13.11, “Deleting a Clustered LVM Volume \(Created in NSSMU or NLVM\),” on page 366](#)
- ♦ [Section 13.12, “Linux LVM Management Tools,” on page 368](#)

13.1 Requirements for Creating LVM Cluster Resources

Your system must meet the requirements in this section in addition to the cluster requirements described in [Chapter 4, “Planning for Novell Cluster Services,” on page 49](#).

- ♦ [Section 13.1.1, “Novell Cluster Services,” on page 306](#)
- ♦ [Section 13.1.2, “Linux Logical Volume Manager 2 \(LVM2\),” on page 306](#)
- ♦ [Section 13.1.3, “Clustered Logical Volume Manager Daemon \(CLVMD\),” on page 306](#)
- ♦ [Section 13.1.4, “Resource IP Address,” on page 306](#)
- ♦ [Section 13.1.5, “Shared Storage Devices,” on page 306](#)
- ♦ [Section 13.1.6, “All Nodes Must Be Present,” on page 307](#)

- ♦ [Section 13.1.7, “Working in Mixed Node OES Clusters,” on page 307](#)
- ♦ [Section 13.1.8, “NCP File Access with Novell NCP Server,” on page 307](#)
- ♦ [Section 13.1.9, “SMB/CIFS File Access with Novell Samba,” on page 308](#)
- ♦ [Section 13.1.10, “Linux File Access Protocols,” on page 308](#)

13.1.1 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.

13.1.2 Linux Logical Volume Manager 2 (LVM2)

The Linux Logical Volume Manager (LVM) 2 software supports LVM volume groups and logical volumes. LVM2 must be installed and running on each node in the cluster. LVM2 runs automatically on OES 11 and later servers; no separate installation or setup is required.

13.1.3 Clustered Logical Volume Manager Daemon (CLVMD)

The Linux Clustered Volume Manager Daemon (CLVMD, `clvmd`) software allows you to exclusively mount a shared volume group on one node at a time in a cluster. It distributes the LVM metadata updates around a cluster. CLVM must be installed and running on each node in the cluster. CLVMD runs automatically on OES 11 and later servers; no separate installation or setup is required.

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

13.1.4 Resource IP Address

Each cluster resource requires a unique static IP address that is in the same subnet as the IP addresses that are used for the cluster and cluster nodes. The IP address is used to provide access and failover capability for the cluster-enabled volume.

13.1.5 Shared Storage Devices

The shared SAN storage device that you use for an LVM volume group cluster resource must be initialized and have no partitions on it. When the device is used in a cluster resource, LVM uses the entire device for the volume group. Ensure that you size your LUNs accordingly. Use the SAN management tools to assign the LUN to all nodes in the cluster.

IMPORTANT: If you use NSS management tools to manage devices, do not enable the Shareable for Clustering option. Doing so adds a 4 KB partition to the device, which makes it unavailable to LVM.

13.1.6 All Nodes Must Be Present

LVM requires the presence of all the nodes in the cluster to modify the metadata on shared storage. This allows LVM to get the exclusive locks it needs to perform actions on shared storage.

Before you attempt to create or modify LVM volume group cluster resources:

- ♦ All of the nodes must be joined in the cluster and running properly.
- ♦ The `clvmd` daemon must be running on all nodes.

13.1.7 Working in Mixed Node OES Clusters

LVM volume group cluster resources are not supported in mixed-node OES clusters when you upgrade from OES 2 SP3 to OES 11 (or later). Complete the upgrade before you create new cluster resources. For information, see [Section 7.3, “Requirements and Guidelines for Upgrading Clusters from OES 2 SP3,” on page 108](#).

13.1.8 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 and later servers. Its NCP volumes feature can be used to provide NCP access to files on an LVM volume group cluster resource. NCP Server must be installed, configured, and running on each node in the cluster.

- ♦ [“Naming Conventions for NCP Volumes” on page 307](#)
- ♦ [“Creating an NCP Volume for a New Clustered LVM Volume” on page 307](#)
- ♦ [“Creating an NCP Volume on an Existing Clustered LVM Volume” on page 308](#)
- ♦ [“Using Antivirus Software with NCP Volumes” on page 308](#)

Naming Conventions for NCP Volumes

NCP volume names can be up to 14 alphanumeric characters, using uppercase letters A through Z and numbers 0 through 9. Underscores (`_`) are allowed.

If you NCP enable a Linux volume as you create it with NSSMU or the `nlvm create linux volume` command, the NCP volume name uses the specified Linux volume name, but all letters are uppercase. NCP treats the Linux volume name as case insensitive. Ensure that the specified Linux volume name does not exceed 14 characters, does not use special characters, and is unique across all nodes in the cluster for both Linux and NCP.

Creating an NCP Volume for a New Clustered LVM Volume

You can configure NCP file access for an LVM volume group cluster resource when you create the resource by using NSSMU or the `nlvm create linux volume` command. With the NCP option enabled, these tools automatically add commands to the resource scripts that mount, dismount, and monitor an NCP volume. The NCP volume is named the same as the LVM logical volume name, and all letters in the name are uppercase. The tools automatically create an NCP Virtual Server object for the volume group cluster resource.

Creating an NCP Volume on an Existing Clustered LVM Volume

You can create an NCP Virtual Server object for the LVM cluster resource to make the resource visible in the Novell eDirectory tree. The virtual server alone does not provide NCP file access.

You can add NCP file access support to an existing LVM cluster resource by creating an NCP Volume object and adding commands to the scripts to manage the volume. For details about setting up NCP volumes on an existing clustered Linux volume, see [“Configuring NCP Volumes with Novell Cluster Services”](#) in the *OES 11 SP1: NCP Server for Linux Administration Guide*.

Using Antivirus Software with NCP Volumes

For information about using antivirus software with NCP volumes, see [“Security”](#) in the *OES 11 SP1: Planning and Implementation Guide*.

13.1.9 SMB/CIFS File Access with Novell Samba

Samba is an open source software suite that lets Linux and other non-Windows servers provide file and print services to clients that support the Microsoft SMB (Server Message Block) and CIFS (Common Internet File System) protocols. Novell Samba is the Linux Samba software that has been modified to work with Novell eDirectory.

Novell Samba can be used to provide SMB/CIFS access to files on an LVM volume group cluster resource for Novell eDirectory users who are enabled for Linux User Management (LUM). For information about LUM-enabling your eDirectory users, see the *OES 11 SP1: Novell Linux User Management Administration Guide*.

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 SP1: Novell Samba Administration Guide*.

IMPORTANT: If you enable both NCP and Novell Samba file access for users, we recommend that you enable the Cross-Protocol Lock (CPL) parameter for NCP Server. CPL helps prevent potential data corruption when files are accessed by non-NCP file access protocols and by other applications that directly access the files with POSIX APIs. CPL is enabled by default. For information, see [“Configuring Cross-Protocol File Locks for NCP Server”](#) in the *OES 11 SP1: NCP Server for Linux Administration Guide*.

13.1.10 Linux File Access Protocols

You can provide native Linux file access to files on an LVM volume group cluster resource for Novell eDirectory users who are enabled for Linux User Management (LUM). The Linux file access protocols must also be LUM enabled on each node of the cluster. For information about LUM-enabling your eDirectory users and native Linux file access protocols, see the *OES 11 SP1: Novell Linux User Management Administration Guide*.

13.2 Initializing a SAN Device

Before you begin, you must initialize the SAN device to set up its device format. You can also initialize a device to wipe its current structure and reconfigure it.

WARNING: Initializing a device removes all partitions and data from the device. Do not initialize the device that contains the operating system.

Devices that you want to use for a clustered Linux LVM volume group should contain no partitions and be in an unshared state. When you initialize the SAN device by using NSS management tools, ensure that the *Shareable for Clustering* option is disabled. CLVM handles the sharing for clustered LVM volumes.

IMPORTANT: NLVM and NSSMU do not support using Linux software RAID devices and NSS software RAID devices with Linux POSIX file systems. You can use a hardware RAID device to achieve device fault tolerance for Linux POSIX volumes.

Initializing a device formats it with an MSDOS or a GPT partitioning scheme. MSDOS supports devices up to 2 TB in size. GPT supports devices of any size. The default is MSDOS. If the device size is greater than 2 TB and the partitioning scheme is not specified, the default partitioning scheme of MSDOS applies, and the device size is truncated to 2 TB with the remainder as unusable space.

Devices that have never been initialized have a format of *None* (NLVM commands) or *Uninitialized* (in NSSMU). Devices that are being used for a Novell Cluster Services SBD (split brain detector) partition also have a format of *None*. You should not use the `nlvm init` command to remove an SBD partition. For information about removing an SBD partition, see [Section 9.16, “Creating or Deleting Cluster SBD Partitions,”](#) on page 157.

Use either procedure in this section to initialize the SAN device that you want to use for the LVM volume group. Do not mark it as shareable for clustering.

- ♦ [Section 13.2.1, “Using NSSMU to Initialize a Device,”](#) on page 309
- ♦ [Section 13.2.2, “Using NLVM Commands to Initialize a Device,”](#) on page 311

13.2.1 Using NSSMU to Initialize a Device

- 1 Ensure that the SAN device is attached to all of the nodes in the cluster.
- 2 Log in as the `root` user to the master node of the cluster, then open a terminal console.
- 3 Launch NSSMU:

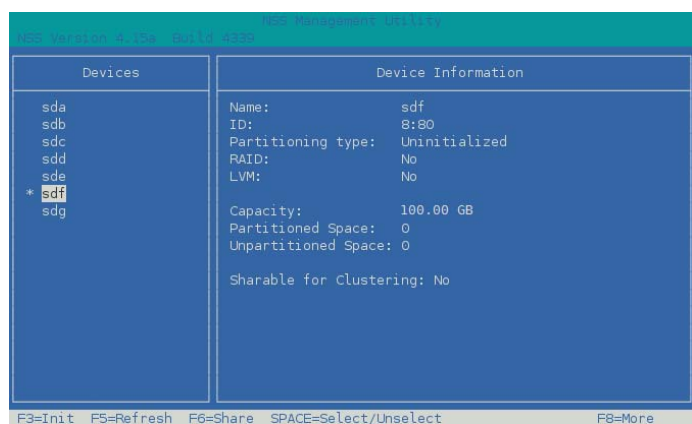
```
nssmu
```

- 4 In the NSSMU main menu, select *Devices*, then press Enter.



- 5 In the *Devices* list, select the SAN device (such as `sdf`), then view information about it.

A device that has never been initialized reports a partition type of *Uninitialized*. If the device contains partitions or data, be prepared to lose all data on the device when it is initialized. The clustered volume group requires the entire device.



6 Press F3 to initialize the selected device.

You are prompted to confirm the initialization. If you continue, any existing partitions on the device are deleted, and a new partition table is created.




7 Read the advisory message, then do one of the following:

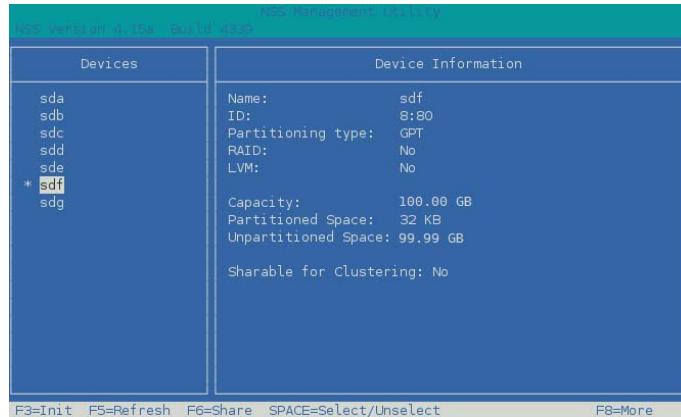
- ◆ Press y (Yes) to confirm that you want to initialize the device, and continue to [Step 8](#).
- ◆ Press n (No) (or press Esc) to cancel. Return to [Step 5](#) and choose a different device.

8 Specify the partitioning scheme to use as DOS or GPT, then press Enter.

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



- 9 Select the device and view its details to verify that the device is initialized and unshared (that is, *Shareable for Clustering* is set to No.) If *Shareable for Clustering* is set to Yes, press F6 to unshare the device. This is necessary to make the entire device available to Clustered LVM (CLVM).



- 10 Press Esc twice to exit NLSMU.

13.2.2 Using NLVM Commands to Initialize a Device

- 1 Ensure that the SAN device is attached to all of the nodes in the cluster.
- 2 Log in as the root user to the master node of the cluster, then open a terminal console.
- 3 At the command prompt, enter

```
nlvm [--force] [--no-prompt] init <device_name> format=<gpt|msdos> unshared
```

Replace *device_name* with the device node name of the device to be initialized, such as `sdf`. This must be the first command option to follow `init`.

Specify the partitioning scheme as `gpt` or `msdos`. The default is `msdos`.

Specify the `unshared` option to ensure that the device is not marked as shareable for clustering.

You can specify the `--force` NLVM option to force the initialization.

You are automatically prompted to confirm the initialize action. Respond by pressing `y` (Yes) or `n` (No). You can use the `--no-prompt` NLVM option to suppress the confirmation.

For example, enter

```
nlvm init sdf format=gpt unshared
```

- 4 Confirm that the device was initialized and is unshared. At the command prompt, enter

```
nlvm list device <device_name>
```

Replace *device_name* with the device node name of the device that you initialized, such as `sdf`.

The command returns information about the device, such as its major:minor number, device size, partitioning format, and whether the device is shared.

For example, enter

```
nlvm list device sde
```

```
avalon:~/Desktop # nlvm list device sde
sde (8:64) size=512.00MB(1048576) free=511.98MB(1048544) format=MSDOS h:s=255:6
3 shared=No RAID=No
```


13.3 Configuring an LVM Volume Group Cluster Resource with NSS Management Tools

This section describes how to use the NSSMU utility and NLVM commands to create a clustered LVM volume group and logical volume on a single SAN device that is assigned to all nodes in the cluster. The volume is exclusively mounted on only one node at a time. Clustered LVM manages the locks for the exclusive mount.

After you create the resource, you should view and modify the cluster resource settings and scripts before you bring the resource online. You can also add lines to its load script, unload script, and monitor script to customize the resource for other uses.

If you enable NCP when you create the LVM volume group in NSSMU or with `nlvm` commands, commands are automatically added to the resource scripts to mount, dismount, and monitor an NCP volume. The NCP volume is named the same as the LVM logical volume name, and all letters in the name are uppercase. An NCP Virtual Server object is created for the resource and commands are added to the scripts to bind or unbind it with the resource IP address.

- ♦ [Section 13.3.1, “Sample Values,” on page 312](#)
- ♦ [Section 13.3.2, “Creating an LVM Volume Group Cluster Resource with NSSMU,” on page 314](#)
- ♦ [Section 13.3.3, “Creating an LVM Volume Group Cluster Resource with NLVM Commands,” on page 318](#)
- ♦ [Section 13.3.4, “Configuring the LVM Cluster Resource Settings,” on page 321](#)
- ♦ [Section 13.3.5, “Viewing or Modifying the LVM Resource Scripts,” on page 324](#)
- ♦ [Section 13.3.6, “Sample LVM Resource Scripts,” on page 326](#)

13.3.1 Sample Values

The configuration scenarios use following sample values. Ensure that you replace the sample values with information for your configuration.

Parameter	Sample Value
Device name for the shared SAN device	<code>/dev/sdf</code> The device is initialized and contains no partitions. It is not enabled as shareable for clustering.
Volume group name	<code>vol44</code> By default, NSSMU uses the logical volume name as the LVM volume group name. If you use the <code>NLVM 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> .
Volume name	<code>vol44</code> If you enable NCP for the volume, this name must comply with the “Naming Conventions for NCP Volumes” on page 307 . Lowercase letters are automatically changed to uppercase for the NCP volume name.

Parameter	Sample Value
NCP volume name	VOL44 The NCP volume is created only if you enable NCP as you create the clustered LVM volume by using NSSMU or the NLVM <code>create linux volume</code> command. The mount point path for the NCP volume is the mount point for the LVM volume. This means that the NCP share is created at the root of the mounted LVM volume when the LVM resource is online.
Linux POSIX file system type	ext3 Valid values are btrfs (requires the <code>btrfsprogs</code> package), ext2, ext3, reiserfs, and xfs.
Make options for the file system	None (do not specify a value). Press Enter to continue. For a list of the supported file system options for the file system type you are making, see the <code>mkfs(8)</code> man page and the man page for the specific file system: <code>mkfs.btrfs(8)</code> , <code>mkfs.ext2(8)</code> , <code>mkfs.ext3(8)</code> , <code>mkfs.reiserfs(8)</code> , or <code>mkfs.xfs(8)</code> . IMPORTANT: The file system creation fails if you specify a make option that is not supported by the file system type.
Mount options	rw The Read/Write (rw) option is specified by default. For a list of available options that work with the file system type you are using, see the <code>mount(8)</code> man page.
Volume size	100 GB A 100 GB LUN is prepared in the shared storage subsystem. It is attached to the nodes in the cluster. The device must be initialized and contain no partitions. It should not be marked as Shareable for clustering. You are not prompted to enter a volume size. The clustered LVM volume group and logical volume use the entire device. When you select the device, all of the device's free available space is displayed in the <i>Free Size</i> field.
Resource IP address	10.10.10.44 This is the IP address of the virtual server for the cluster resource. The address must be unique and in the same subnet as the cluster's IP address. Specify the IP address in IPv4 format.
Mount device	/dev/vol44/vol44 The mount device path format is <code>/dev/<volume_group_name>/<logical_volume_name></code> If you specify a different name for the volume group, such as <code>vg44</code> , by using the NLVM <code>create linux volume</code> command, the mount device path is <code>/dev/vg44/vol44</code> .


Parameter	Sample Value
Mount point path	/mnt/vol44 You can use any valid Linux path as the mount point. The default mount point location for LVM logical volumes created by NSSMU and NLVM is /usr/novell/<lx_volume_name>. NSSMU automatically creates the mount point path if it does not exist on this node. However, you must manually create the mount point on each of the other nodes in the cluster. If the path does not exist on a node when you fail over the resource to it, the resource goes comatose You can alternatively add the following line to the load script before the Linux file system mount command to create the path on a node if it does not exist: ignore_error mkdir -p \$MOUNT_PATH

13.3.2 Creating an LVM Volume Group Cluster Resource with NSSMU

This section describes how to use NSSMU to create and cluster-enable an LVM volume group. NSSMU automatically uses the same script format as the Generic File System template (Generic_FS_Template) to create an LVM volume group cluster resource.


- 1 Ensure that the SAN device is attached to all of the nodes in the cluster.
- 2 Log in as the root user to the master node of the cluster, then open a terminal console.
- 3 If you have not already done so, initialize the SAN device that you want to use for the LVM volume group.
LVM allows only one device per volume group. For information, see [Section 13.2, “Initializing a SAN Device,” on page 308](#).
- 4 Launch NSSMU:

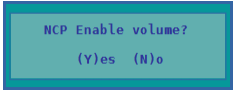


nssmu
- 5 In the NSSMU main menu, select *Linux Volumes*, then press Enter.



- 6 On the Linux Volumes page, press *Insert* to launch the volume creation wizard, then enter the following information as you are prompted for it:

Parameter	Action
Select LVM type	Select <i>Cluster Enabled LVM2 Volume</i> , then press Enter.



Parameter	Action
<i>NCP Enable volume?</i>	<p>Specify whether to enable the LVM logical volume for NCP file access. Press y (yes) to allow NCP file access, or press n (No) to disallow NCP file access at this time.</p>  <p>If you enable NCP, NSSMU automatically adds commands to the resource scripts that mount, dismount, and monitor an NCP volume. The NCP volume is named the same as the LVM logical volume name, and all letters in the name are uppercase. The tools automatically create an NCP Virtual Server object for the resource.</p> <p>If you do not enable NCP at this time, you can manually create an NCP virtual server, create an NCP Volume object, and add NCP commands to the scripts. For information, see Section 13.5, “Creating a Virtual Server Object for an LVM Volume Group Cluster Resource,” on page 343.</p>
<i>Enter volume name</i>	<p>Type the name for the LVM logical volume (such as vol44), then press Enter.</p>  <p>The specified name is also used for the LVM volume group, which is used in the name of the LVM volume group cluster resource.</p> <p>If you enable NCP for the volume, the specified name is also applied to the NCP volume. An NCP volume name can be up to 14 alphanumeric characters. Underscores are allowed. All letters in the LVM volume name are uppercase for the NCP volume name. For example, if the LVM volume name is vol44, the NCP volume name is VOL44.</p>
<i>Enter volume IP address</i>	<p>Type the IP address to use for the LVM cluster resource in IPv4 format (such as 10.10.10.44), then press Enter.</p> 
<i>Select Volume type</i>	<p>Select one of the following the Linux POSIX file systems, then press Enter:</p> <ul style="list-style-type: none"> ◆ btrfs <p>This option is displayed only if the <code>btrfsprogs</code> package is installed.</p> ◆ ext2 ◆ ext3 ◆ resiserfs ◆ xfs

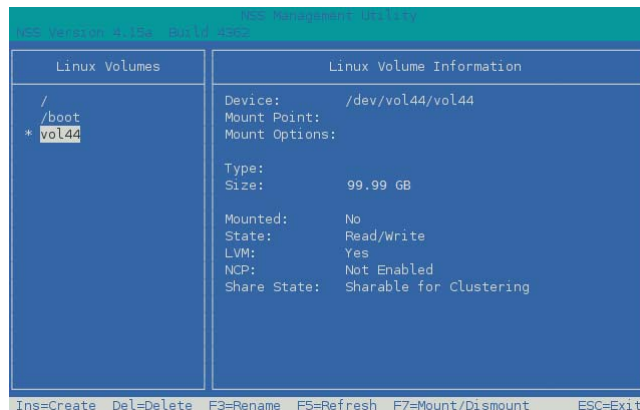
Parameter	Action																								
Enter full mount point path	<p>Type the full mount point path for the LVM logical volume (such as /mnt/vol44), then press Enter.</p> <div>Enter full mountpoint path: /usr/novell/vol44</div> <p>The default path is /usr/novell/<lx_volume_name>, such as /usr/novell/vol44.</p> <p>If NCP is enabled, the specified path is also used as the mount point path for the NCP volume.</p>																								
Enter any make options	<p>Press Enter to continue without specifying options, or specify the desired make options for the file system type you are making, then press Enter.</p> <div>Enter any make options:</div> <p>For a list of the supported file system options for the file system type you are making, see the mkfs(8) man page and the man page for the specific file system: mkfs.btrfs(8), mkfs.ext2(8), mkfs.ext3(8), mkfs.reiserfs(8), or mkfs.xfs(8).</p> <p>IMPORTANT: The file system creation fails if you specify a make option that is not supported by the specified file system type.</p>																								
Enter any mount options	<p>Press Enter to accept the default read and write options (rw).</p> <div>Enter any mount options: rw</div> <p>The Read/Write (rw) option is specified by default. You can add other mount options, or modify this option. For a list of available options that work with the specified file system type, see the mount(8) man page.</p>																								
Device	<p>From the list of available devices, select the device that you initialized in Section 13.2, “Initializing a SAN Device,” on page 308, such as sdf, then press Insert or Enter to select the device.</p> <p>The LVM volume group requires the entire device. You cannot specify the amount of space to use. The device’s free available space is displayed in the Free Size field.</p> <div><div>Novell Management Utility</div><div>Novell Version 4.10a, Build 4325</div><table><tr><th>Device</th><th>Free Size</th><th>Shared</th><th>Selected</th></tr><tr><td>1 sdf</td><td>511.95 MB</td><td>No</td><td>99.99 GB</td></tr><tr><td colspan="3">Total</td><td>99.99 GB</td></tr></table><table><tr><th>Device</th><th>Free Size</th><th>Shared</th></tr><tr><td>sde</td><td>511.98 MB</td><td>No</td></tr><tr><td>+ sdf</td><td>99.99 GB</td><td>No</td></tr><tr><td>sdg</td><td>715.95 MB</td><td>No</td></tr></table><div>Ins/ENTER=Select Del=Remove F3=Accept Select list U=Up D=Down ESC/Q=Prev Menu</div></div>	Device	Free Size	Shared	Selected	1 sdf	511.95 MB	No	99.99 GB	Total			99.99 GB	Device	Free Size	Shared	sde	511.98 MB	No	+ sdf	99.99 GB	No	sdg	715.95 MB	No
Device	Free Size	Shared	Selected																						
1 sdf	511.95 MB	No	99.99 GB																						
Total			99.99 GB																						
Device	Free Size	Shared																							
sde	511.98 MB	No																							
+ sdf	99.99 GB	No																							
sdg	715.95 MB	No																							

- 7 Press F3 to accept the setup you entered for the volume group cluster resource.

The resource is created and brought online on the node where you created it. The resource is named `<volume_group_name>_resource`. In the example, the name of the volume group is the same as the logical volume, so the resource name is `vol144_resource`.

Typically, the volume creation takes less than 10 seconds. However, if you have a large tree or the server does not hold an eDirectory replica, the create time can take up to 3 minutes.

- 8 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: <code>/dev/vol144/vol144</code>
Mount Point	When the resource is online, this specifies the path on the root file system where this volume is mounted. Example: <code>/mnt/vol144</code>
Mount options	When the resource is online, this specifies the mount options that are applied whenever this volume is automatically mounted after a reboot. Example: <code>rw</code>
Type	When the resource is online, this specifies the file system type. Examples: <code>btrfs</code> , <code>ext2</code> , <code>ext3</code> , <code>reiserfs</code> , <code>xfs</code>
Size	Specifies the amount of space reserved for this volume. Example: <code>99.58 GB</code>
Mounted	Specifies whether the volume is mounted or unmounted. When the resource is brought online, the load script mounts the logical volume. Value: Yes or No
State	Specifies the availability for the file system. Example: <code>Read/Write</code>
LVM	Specifies whether the specified volume is an LVM logical volume. Value: Yes

Parameter	Description
NCP	Specifies whether the LVM logical volume is enabled for NCP. The NCP state cannot be determined for a clustered LVM volume. You can use the Clusters plug-in for Novell iManager to determine if there are NCP commands in the load script. On the Cluster Options page, select the resource to view its properties, then click the <i>Scripts</i> tab.
Share State	Specifies whether the LVM logical volume is cluster enabled for a Novell Cluster Services cluster. Value: Shareable for Clustering When an LVM logical volume group is clustered, CLVM manages the share state, not the device. NSSMU reports the device as Not Shareable for Clustering. NSSMU reports the clustered LVM volume is Shareable for Clustering.

9 Press Escape twice to exit NSSMU.

10 Continue with [Section 13.3.4, “Configuring the LVM Cluster Resource Settings,”](#) on page 321.

13.3.3 Creating an LVM Volume Group Cluster Resource with NLVM Commands

This section describes how to use NLVM commands to create and cluster-enable an LVM volume group. NLVM automatically uses the same script format as the Generic File System template (`Generic_FS_Template`) to create an LVM volume group cluster resource. The NLVM command allows you to specify a group name that is different than the volume name.

- 1 Ensure that the SAN device is attached to all of the nodes in the cluster.
- 2 Log in to the master node of the cluster as the `root` user, then open a terminal console.
- 3 If you have not already done so, initialize the SAN device that you want to use for the LVM volume group.

For information, see [Section 13.2, “Initializing a SAN Device,”](#) on page 308.

- 4 Create a clustered LVM volume group and logical volume. At the command prompt, enter the following (all on the same line, of course):

```
nlvm [nlvm_options] create linux volume
type=<btrfs|ext2|ext3|reiserfs|xfs>
device=<device_name>
[mp=</mount_path>]
[mkopt=<option1[,option2,...]>]
[mntopt=<option1[option2[]...]>]
lvm
name=<lvm_volume_name>
[group=<lvm_group_name>]
shared
ip=<IP_address_for_LVM_volgroup_cluster_resource>
[ncp]
```

For details about using this command, see “[Create Linux Volume](#)” in the [OES 11 SP1: NLVM Reference](#).

Parameters and Options

type=<btrfs|ext2|ext3|reiserfs|xfs>

You must specify the type of file system to use for the volume. Btrfs requires that the `btrfsprogs` package is installed on all nodes. Supported file systems are btrfs, ext2, ext3, reiserfs, and xfs.

device=<device_name>

You must specify the device to use. Replace *device_name* with the device node name of the device that you want to use for the volume, such as `sdf`. The entire device is used for the LVM volume group. You cannot specify a size to use. The device must already be initialized, contain no partitions, and must not be marked as shareable for clustering.

[mp=</mount_path>]

If a mount point path is not specified, the utility assigns a default mount path of `/usr/novell/<volume_name>`. Use the `mp` option to specify a custom mount point. Replace *mount_path* with the full Linux path of the mount point where the volume is to be mounted. The final directory name in the path can be the same or different than the specified volume name. If the path does not currently exist, it is automatically created on that node. You must manually create the path on other nodes in the cluster.

[mkopt=<option1[,option2,...]>]

You can use the `mkopt` option to specify the options to use when running `mkfs`. For a list of available options, see the `mkfs(8)` man page. No default option is specified. For a list of the supported file system options for the file system type you are making, see the `mkfs(8)` man page and the man page for the specific file system: `mkfs.btrfs(8)`, `mkfs.ext2(8)`, `mkfs.ext3(8)`, `mkfs.reiserfs(8)`, or `mkfs.xfs(8)`.

IMPORTANT: The file system creation fails if you specify a make option that is not supported by the file system type.

[mntopt=<option1[option2[...]]>]

You can use the `mntopt` option to specify the options to use when mounting the volume. For a list of available options, see the `mount(8)` man page. The default `mntopt` value is `rw` (Read/Write).

lvm

You must specify the `lvm` option to create an LVM volume group and logical volume. Use this option with the `name` option.

name=<lvm_volume_name>

Replace *lvm_volume_name* with a name for the LVM volume. If you do not specify the `group` option, this name is also used as the LVM volume group name, which is used in the cluster resource name. For LVM logical volume naming conventions, see “[Linux LVM Volume Group and Logical Volume Names](#)” in the *OES 11 SP1: NLVM Reference*.

If you use the `NCP` option, the `NCP` volume name uses the same name as the LVM logical volume name, but all letters are uppercase. `NCP` volume names can be up to 14 alphanumeric characters, and underscores are allowed.

[group=<lvm_group_name>]

Replace *lvm_group_name* with a name for the LVM volume group. The group name is also used in the cluster resource name. If you do not specify a volume group name, the group is automatically named the same as the LVM volume.

shared

You must specify the `shared` option to create a clustered LVM volume group and logical volume.

ip=<IP_address_for_LVM_volgroup_cluster_resource>

Replace `IP_address_for_LVM_volgroup_cluster_resource` with a static unique IP address to use for the LVM cluster resource. Specify the address in IPv4 format.

ncp

Specify the `ncp` option to NCP enable the LVM logical volume for NCP file access.

- ♦ If you enable NCP, NSSMU automatically adds commands to the resource scripts that mount, dismount, and monitor an NCP volume. The NCP volume is named the same as the LVM logical volume name, and all letters in the name are uppercase. The tools automatically create an NCP Virtual Server object for the resource.
- ♦ If you do not enable NCP at this time, you can manually create a virtual server and add NCP file access later. For information, see [Section 13.5, “Creating a Virtual Server Object for an LVM Volume Group Cluster Resource,”](#) on page 343.

Example For example, at the command prompt, enter the following (all on the same line):

```
nlvm create linux volume
type=ext3
device=sdf
mp=/mnt/vol44
mntopt=rw
lvm
name=vol44
group=vg44
shared
ip=10.10.10.44
ncp
```

If the command is successful, the response is

```
Linux clustered volume vol44 created.
```

Typically, the volume creation takes less than 10 seconds. However, if you have a large tree or the server does not hold an eDirectory replica, the create time can take up to 3 minutes.

5 Verify that the cluster resource was created and brought online by entering

```
cluster status
```

The resource is named <lv_name>_resource. In the following example, `vol44_resource` is in the *Running* state.

```
avalon:~/Desktop # cluster status
Master_IP_Address_Resource      Running    avalon      1
lvmvol134_resource              Running    avalon      1
vol44_resource                   Running    avalon      1
```

6 Continue with [Section 13.3.4, “Configuring the LVM Cluster Resource Settings,”](#) on page 321.

13.3.4 Configuring the LVM Cluster Resource Settings

After you create a clustered LVM volume group and logical volume, use the procedure in this section to verify that the LVM volume group cluster resource was created and is online. You can customize the resource policies, monitoring, and preferred nodes settings.

- 1 Open Novell iManager in a Web browser, then log in as a cluster administrator user.
- 2 In *Roles and Tasks*, select *Clusters > My Clusters*.
- 3 Select the cluster where you created the volume group cluster resource.

If the cluster you want to manage is not in the list, you can add it. Click *Add*, browse to locate and select the Cluster object of the cluster, then click *OK*.

- 4 In the list of resources, locate the new resource, such as `vol144_resource`, and notice the state of the resource. It should be online and running.

My Clusters > clus1.ncs.novell

clus1.ncs.novell

View the status of both servers and resources in the cluster. You can also migrate cluster resources to different servers in the cluster, or change the state of resources to offline or online.

Cluster Manager BCC Manager Cluster Event Log Cluster Options

Run Report

Epoch: 0

avalon

Cluster State						
Online Offline Migrate Respond to Alert Refresh						
2 Item(s)						
<input type="checkbox"/>	Type	Name	State	Location	Lives	Up Since
<input type="checkbox"/>		Master_IP_Address_Resource	Running	avalon	1	Aug 8, 2012 10:01:33 AM
<input type="checkbox"/>		vol144_resource	Running	avalon	1	

Close

- 5 Click the resource's name link to open its Properties page.

You can alternatively go to the Cluster Options page, then click the resource's name link, or select the check box next to the resource and click *Details*.

- 6 On the Resource Policies page, view and modify the resource's Policy settings if needed. Click *Apply* if you make changes.

My Clusters > clus1.ncs.novell > vol44_resource

vol44_resource

Policies Monitoring Preferred Nodes Scripts Business Continuity

Set Start, Failover and Failback modes for the new resource. Changes other than business continuity changes made to a resource will not take affect until the resource is reloaded.

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 a 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 10.7, “Enabling Monitoring and Configuring the Monitor Script,” on page 192](#).

My Clusters > clust1.ncs.novell > vol44_resource

vol44_resource

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. Changes other than business continuity changes made to a resource will not take affect until the resource is reloaded.

☒ Enable Resource Monitoring

Polling Interval: 1 Minutes

Failure Rate

Maximum Local Failures: 3

Time Interval: 10 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 monitor script for this resource to run.

You can specify the value in minutes or seconds. For information, see [“Polling Interval” on page 192](#).

- 7d Specify the number of failures (*Maximum Local Failures*) for the specified amount of time (*Time Interval*).

For information, see [“Failure Rate” on page 193](#).

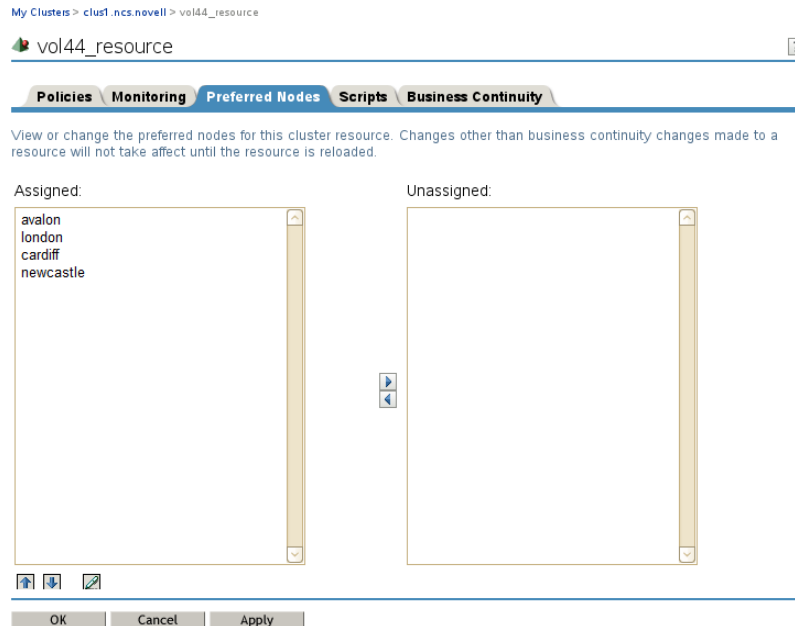
- 7e Specify the *Failover Action* by indicating whether you want the resource to be set to a comatose state, to migrate to another server, or to reboot the hosting node (without synchronizing or unmounting the disks) if a failure action initiates. The reboot option is normally used only for a mission-critical cluster resource that must remain available.

For information, see [“Failure Action” on page 193](#).

- 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 Continue with [Section 13.3.5, “Viewing or Modifying the LVM Resource Scripts,” on page 324](#).

13.3.5 Viewing or Modifying the LVM Resource Scripts

You can customize the scripts by adding lines for other products that use a shared LVM volume group resource. Compare the generic script with the templates for those products to identify what lines need to be added or modified.

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the cluster where you created the volume group cluster resource.
- 3 Select the cluster resource's name link to open the Properties page, then click the *Scripts* tab.
The *Scripts* tab automatically opens to the load script.
- 4 On the *Load Script* page, view or modify the load script. Click *Apply* if you make changes.
Ensure that all letters in the NCP volume name are uppercase letters A to Z. See the [“Sample LVM Resource Load Scripts Created by NSS Management Tools” on page 327](#).

My Clusters > clus1.ncs.novell > vol44_resource

vol44_resource

Policies Monitoring Preferred Nodes **Scripts** Business Continuity

Load Script | Unload Script | Monitor Script

View or edit the load script for this cluster resource. Changes other than business continuity changes made to a resource will not take affect until the resource is reloaded.

Script:

```

MOUNT_FS=ext3
# define the volume group name
VOL_GROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/vol44/vol44
# define the mount point
MOUNT_POINT=/usr/novell/vol44
# define NCP server name
NCP_SERVER=clus1-vol44-SERVER
# define NCP volume name
NCP_VOLUME=VOL44

# activate the volume group
exit_on_error vgchange -a ey $VOL_GROUP_NAME

```

Timeout: 6 Minutes

OK Cancel Apply

- Click the *Unload Script* link to view or modify the unload script. Click *Apply* if you make changes.

Ensure that all letters in the NCP volume name are uppercase letters A to Z. See the [“Sample LVM Resource Unload Scripts Created by NSS Management Tools”](#) on page 328.

My Clusters > clus1.ncs.novell > vol44_resource

vol44_resource

Policies Monitoring Preferred Nodes **Scripts** Business Continuity

Load Script | Unload Script | Monitor Script

View or edit the unload script for this cluster resource. Changes other than business continuity changes made to a resource will not take affect until the resource is reloaded.

Script:

```

RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOL_GROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/vol44/vol44
# define the mount point
MOUNT_POINT=/usr/novell/vol44
# define NCP server name
NCP_SERVER=clus1-vol44-SERVER
# define NCP volume name
NCP_VOLUME=VOL44

ignore_error ncpcon unbind --ncpservname=$NCP_SERVER --ipaddress=$RE

```

Timeout: 6 Minutes

OK Cancel Apply

- Click the *Monitor Script* link to view or modify the monitor script. Click *Apply* if you make changes.

Ensure that all letters in the NCP volume name are uppercase letters A to Z. See the [“Sample LVM Resource Monitor Scripts Created by NSS Management Tools”](#) on page 329.

My Clusters > clus1.ncs.novell > vol44_resource

vol44_resource

Policies Monitoring Preferred Nodes **Scripts** Business Continuity

Load Script | Unload Script | Monitor Script

View or edit the monitor script for this cluster resource. Changes other than business continuity changes made to a resource will not take affect until the resource is reloaded.

Script

```
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOL_GROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/vol44/vol44
# define the mount point
MOUNT_POINT=/usr/novell/vol44
# define NCP server name
NCP_SERVER=clus1-vol44-SERVER
# define NCP volume name
NCP_VOLUME=VOL44

# check the IP address
ignore_error status_secondary_ipaddress $RESOURCE_IP
```

Timeout: 6 Minutes

OK Cancel Apply

7 At the bottom of the page, click **OK** to close the Properties page and save your changes.

The changes do not take effect until the resource is taken offline and brought online again.

8 If you modified the settings on any of the pages, you must take the resource offline and bring it online in order for the changes to take effect.

8a In Roles and Tasks, select *Clusters > My Clusters*.

8b On the Cluster Manager page, select the check box next to the resource, then click *Offline*.
Wait for the status to report that it is offline, then continue.

8c Select the check box next to the resource, then click *Online*.

8d Verify that the resource comes online and reports a *Running* state.

If the resource goes into a Comatose state, it is probably because you made a mistake in the lines you added or modified in the scripts. Take the resource offline, then go back to correct the scripts, and try to bring it online again.

13.3.6 Sample LVM Resource Scripts

The sample LVM resource scripts in this section are automatically generated for an LVM volume group cluster resource when you use NLVM or NSSMU to create a clustered LVM2 volume, as described in [Section 13.3, “Configuring an LVM Volume Group Cluster Resource with NSS Management Tools,” on page 312](#). See [Section 13.3.1, “Sample Values,” on page 312](#) for information about the sample values used in these scripts. Ensure that you replace sample values with those for your own system.

- ♦ [“Sample LVM Resource Load Scripts Created by NSS Management Tools” on page 327](#)
- ♦ [“Sample LVM Resource Unload Scripts Created by NSS Management Tools” on page 328](#)
- ♦ [“Sample LVM Resource Monitor Scripts Created by NSS Management Tools” on page 329](#)

Sample LVM Resource Load Scripts Created by NSS Management Tools

Compare the load scripts in this section to identify the lines that are added when you enable the LVM logical volume for NCP file access:

- ♦ [“Without NCP File Access” on page 327](#)
- ♦ [“With NCP File Access” on page 327](#)

Without NCP File Access

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name (nssmu uses volume name for group name)
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# activate the volume group
exit_on_error vgchange -a ey $VOLGROUP_NAME

# create the mount point if it does not exist on the node
ignore_error mkdir -p $MOUNT_POINT

# mount the file system
exit_on_error mount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

exit 0
```

With NCP File Access

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name (nssmu uses volume name for group name)
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# define NCP server name
NCP_SERVER=clus1-vol44-SERVER
# define NCP volume name
NCP_VOLUME=VOL44

# activate the volume group
exit_on_error vgchange -a ey $VOLGROUP_NAME

# create the mount point if it does not exist on the node
ignore_error mkdir -p $MOUNT_POINT
```



```
# mount the file system
exit_on_error mount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# mount the NCP volume
exit_on_error ncpcon mount $NCP_VOLUME=251,PATH=$MOUNT_POINT

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

# bind the NCP virtual server to the resource IP address
exit_on_error ncpcon bind --ncpservname=$NCP_SERVER --ipaddress=$RESOURCE_IP

exit 0
```

Sample LVM Resource Unload Scripts Created by NSS Management Tools

Compare the unload scripts in this section to identify the lines that are added when you enable the LVM logical volume for NCP file access:

- ♦ [“Without NCP File Access” on page 328](#)
- ♦ [“With NCP File Access” on page 328](#)

Without NCP File Access

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name (nssmu uses volume name for group name)
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# unmount the volume
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error umount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# deactivate the volume group
exit_on_error vgchange -a n $VOLGROUP_NAME

exit 0
```

With NCP File Access

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name (nssmu uses volume name for group name)
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# define NCP server name
```



```

NCP_SERVER=clus1-vol44-SERVER
# define NCP volume name
NCP_VOLUME=VOL44

# unbind the NCP virtual server from the resource IP address
ignore_error ncpcon unbind --ncpservname=$NCP_SERVER --ipaddress=$RESOURCE_IP

# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# dismount the NCP volume
ignore_error ncpcon dismount $NCP_VOLUME

# unmount the volume
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error umount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# deactivate the volume group
exit_on_error vgchange -a n $VOLGROUP_NAME

exit 0

```

Sample LVM Resource Monitor Scripts Created by NSS Management Tools

Compare the monitor scripts in this section to identify the lines that are added when you enable the LVM logical volume for NCP file access:

- ♦ [“Without NCP File Access” on page 329](#)
- ♦ [“With NCP File Access” on page 330](#)

Without NCP File Access

```

#!/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

```


With NCP File Access

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# define NCP server name
NCP_SERVER=clus1-vol44-SERVER
# define NCP volume name
NCP_VOLUME=VOL44

# check the LVM logical volume
exit_on_error status_lv $MOUNT_DEV

# check the NCP volume
exit_on_error ncpcon volume $NCP_VOLUME

# check the file system
exit_on_error status_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# status the IP address
exit_on_error status_secondary_ipaddress $RESOURCE_IP

exit 0
```

13.4 Configuring an LVM Cluster Resource with LVM Commands and the Generic File System Template

This section describes how to use Linux Logical Volume Manager (LVM) commands to create a shared LVM volume group. You use the Generic File System template (`Generic_FS_Template`) to create a volume group cluster resource that cluster-enables the existing LVM volume group.

After you create the resource, you can add lines to its load script, unload script, and monitor script to customize the resource for other uses. Compare the `Generic_FS_Template` to the resource template for your product to determine which lines need to be added or modified.

- [Section 13.4.1, “Creating a Shared LVM Volume with LVM Commands,” on page 330](#)
- [Section 13.4.2, “Creating a Generic File System Cluster Resource for an LVM Volume Group,” on page 335](#)
- [Section 13.4.3, “Sample Generic LVM Resource Scripts,” on page 341](#)

13.4.1 Creating a Shared LVM Volume with LVM Commands

You can create a shared LVM volume group and logical volume by using native Linux LVM commands.

- [“Sample Values for the LVM Volume Group and Logical Volume” on page 331](#)
- [“Creating a Shared LVM Volume Group \(Quick Reference\)” on page 331](#)
- [“Creating a Shared LVM Volume Group \(Detailed Instructions\)” on page 332](#)

Sample Values for the LVM Volume Group and Logical Volume

The procedures in this section use 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>btrfs</code> , <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)

You can create the volume group and logical volume by issuing the following LVM commands as the `root` user on the cluster node.

Before you begin, initialize the SAN device as described in [Section 13.2, “Initializing a SAN Device,” on page 308](#).

When you are done, continue with [Section 13.4.2, “Creating a Generic File System Cluster Resource for an LVM Volume Group,” on page 335](#).

Action	Command
1. Create the LVM physical volume.	<code>pvccreate <device></code>
2. Create the clustered LVM volume group.	<code>vgcreate -c y <vg_name> <device></code>
3. Activate the volume group exclusively on the node.	<code>vgchange -a ey <vg_name></code>
4. Create the LVM logical volume.	<code>lvcreate -n <lv_name> -L size <vg_name></code>
5. Add a file system to the LVM logical volume.	<code>mkfs -t <fs_type> /dev/<vg_name>/<lv_name> [fs_options]</code>
6. If it does not exist, create the mount point path.	<code>mkdir -p <full_mount_point_path></code>
7. Mount the LVM logical volume.	<code>mount -t <fs_type> /dev/<vg_name>/<lv_name> <mount_point></code>

Action	Command
8. Dismount the LVM logical volume and deactivate the volume group.	<code>umount /dev/<vg_name>/<lv_name></code> <code>vgchange -a n <vg_name></code>
After you create the Generic LVM cluster resource, use the cluster resource to control when the volume is online or offline.	

Creating a Shared LVM Volume Group (Detailed Instructions)

For detailed instructions, use the following procedure to create the LVM volume group and logical volume:

- 1 Log in as the Linux root user to the first node of the cluster, then open a terminal console.
- 2 Initialize the SAN device as described in [Section 13.2, “Initializing a SAN Device,”](#) on page 308. LVM allows only one device per volume group.
- 3 Create an LVM physical volume on the device (such as `/dev/sdd`) by entering:

```
pvccreate <device>
```

For example:

```
pvccreate /dev/sdd
No physical volume label read from /dev/sdd
Physical volume "/dev/sdd" successfully created
```

- 4 Display information about the physical volume by entering:

```
pvddisplay [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
```

- 5 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.

- 6 Activate the volume group exclusively on the current server by entering:

```
vgchange -a ey <vg_name>
```


The `-a` option activates the volume. The `ey` parameter specifies the values exclusively and yes.
For example:

```
vgchange -a ey myclustervg01
logical volume(s) in volume group "myclustervg01" now active
```

7 View information about the volume group by using the `vgdisplay` command.

```
vgdisplay <vg_name>
```

Notice that 4 MB of the device is used for the volume group's Physical Extent (PE) table. You must consider this reduction in available space on the volume group when you specify the size of the LVM logical volume in the next step ([Step 8](#)).

For example:

```
vgdisplay 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-1oO7-jBH3-qHvySe
```

8 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`.

9 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-a1RR-i4Lw-Wwdt-v5io-2hDN-grWTLH
LV Write Access        read/write
LV Status              available
# open                 0
LV Size                500.00 MB
Current LE             125
Segments               1
Allocation             inherit
Read ahead sectors     auto
- currently set to    1024
Block device           253:1
```

- 10** Create a file system (such as Btrfs, Ext2, Ext3, ReiserFS, or XFS) on the LVM logical volume by entering:

```
mkfs -t <fs_type> /dev/<vg_name>/<lv_name> [fs_options]
```

For example:

```
mkfs -t ext3 /dev/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.
```

- 11** If the mount point does not exist, create the full directory path for the mount point.

```
mkdir -p <full_mount_point_path>
```

For example:

```
mkdir -p /mnt/myclustervol01
```

- 12** Mount the logical volume on the desired mount point by entering:

```
mount -t <fs_type> /dev/<vg_name>/<lv_name> <mount_point>
```

For example:

```
mount -t ext3 /dev/myclustervg01/myclustervol01 /mnt/myclustervol01
```

- 13** Dismount the volume and deactivate the LVM volume group by entering:

```
umount /dev/<vg_name>/<lv_name>
vgchange -a n "<vg_name>"
```


This allows you to use the load script and unload script to control when the volume group is activated or deactivated in the cluster.

For example, to unmount the volume and deactivate the myclustervg01 volume group, enter

```
umount /dev/myclustervg01/myclustervol01
vgchange -a n "myclustervg01"
0 logical volume(s) in volume group "myclustervg01" now active
```

- 14 Continue with [Section 13.4.2, “Creating a Generic File System Cluster Resource for an LVM Volume Group,”](#) on page 335.

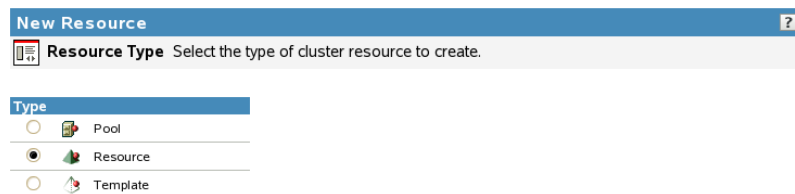
13.4.2 Creating a Generic File System Cluster Resource for an LVM Volume Group

You can use the Novell Cluster Services Generic File System template (Generic_FS_Template) to create a cluster resource for the LVM volume group.

- 1 Ensure that the SAN device is attached to all of the nodes in the cluster.
- 2 Ensure that the existing LVM volume you want to cluster-enable is deactive on all nodes in the cluster.

For information, see [Step 13](#) in [Section 13.4.1, “Creating a Shared LVM Volume with LVM Commands,”](#) on page 330.

- 3 In iManager, select *Clusters > My Clusters*.
- 4 Select the cluster you want to manage, then click *Cluster Options*.
- 5 Click the *New* link.
- 6 Click the *Resource* radio button, then click *Next*.

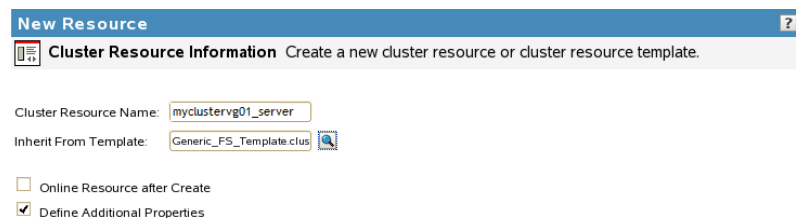


- 7 Specify information to define your new cluster resource:

- 7a Specify the name of the resource you want to create, such as 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.

- 7b In the *Inherit From Template* field, browse to select the *Generic_FS_Template*.



- 7c Deselect or select *Online Resource after Create*.

This option is deselected by default. Deselect this option to allow an opportunity to verify the scripts and resource policies before you bring the resource online for the first time.

Select this option to automatically online the resource on its most preferred node as soon as it is created and configured. The most preferred node is the first server in the assigned nodes list for the resource. You should also select the *Define Additional Properties* check box to configure the resource now.

7d Select the *Define Additional Properties* check box.

When this option is enabled, a wizard walks you through the load, unload, and monitor scripts, then allows you to set up the policies and preferred nodes for the cluster resource.

7e Click *Next*.

If you selected *Define Additional Properties*, continue with the setup. Otherwise, select the resource in the Cluster Objects list to open its properties page, and configure the resource.

8 Configure the load script for the resource by replacing the variables with your own values, specify the *Load Script Timeout* value, then click *Next*.

The following is the default Generic_FS template load script. The variable values must be specified before you can successfully bring the resource online. Be mindful that Linux path names are case sensitive. The mount path must already exist. You can add a `mkdir` command in the script for the mount point, or create the directory manually on each node before you allow the resource to fail over to other nodes. A sample script is available in [“Sample Generic LVM Resource Load Script” on page 341](#).

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

# define the IP address
RESOURCE_IP=a.b.c.d
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=name
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/volume_name
# define the mount point
MOUNT_POINT=/mnt/mount_point

# activate the volume group
exit_on_error vgchange -a ey $VOLGROUP_NAME

# create the mount point on the node if it does not exist
ignore_error mkdir -p $MOUNT_POINT

# mount the file system
exit_on_error mount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

exit 0
```


- 9 Configure the unload script for the resource by replacing the variables with your own values, specify the *Unload Script Timeout* value, then click *Next*.

The following is the default Generic_FS template unload script. The variable values must be specified before you bring the resource online. Be mindful that Linux path names are case sensitive. A sample script is available in [“Sample Generic LVM Resource Unload Script” on page 342](#).

```
#!/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
```

- 10 Configure the monitor script for the resource by replacing the variables with your own values, specify the *Monitor Script Timeout* value, then click *Next*.

The following is the default Generic_FS template monitor script. Be mindful that Linux path names are case sensitive. The variable values must be specified before you bring the resource online. A sample script is available in [“Sample Generic LVM Resource Monitor Script” on page 342](#).

```
#!/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
```


- 11** On the Resource Policies page, view and modify the resource's Policy settings:
- 11a** (Optional) Select the *Resource Follows Master* check box if you want to ensure that the resource runs only on the master node in the cluster.

If the master node in the cluster fails, the resource fails over to the node that becomes the master.
 - 11b** (Optional) Select the *Ignore Quorum* check box if you don't want the cluster-wide timeout period and node number limit enforced.

The quorum default values were set when you installed 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.
 - 11c** (Optional) By default, the Generic File System resource template sets the Start mode and Failover mode to *Auto* and the Failback Mode to *Disable*. You can change the default settings as needed.
 - ♦ **Start Mode:** If the Start mode is set to *Auto*, the resource automatically loads on a designated server when the cluster is first brought up. If the Start mode is set to *Manual*, you can manually start the resource on a specific server when you want, instead of having it automatically start when servers in the cluster are brought up.
 - ♦ **Failover Mode:** If the Failover mode is set to *Auto*, the resource automatically moves to the next server in the 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.
 - 11d** Click *Next*.

- 12 On the Preferred Nodes page, assign preferred nodes for the resource by moving them from the *Unassigned* list to the *Assigned* list.

When you bring a resource online, it is automatically loaded on the most preferred node in the list. If the node is not available, the other nodes are tried in the order that they appear in the list. You can modify the order of the nodes by clicking the Edit (pen) icon to open the list a text editor. In the editor, click *OK* to close the editor.

My Clusters > clus1.ncs.novell > vol44_resource

vol44_resource

Policies Monitoring Preferred Nodes Scripts Business Continuity

View or change the preferred nodes for this cluster resource. Changes other than business continuity changes made to a resource will not take affect until the resource is reloaded.

Assigned:

- avalon
- london
- cardiff
- newcastle

Unassigned:

OK Cancel Apply

- 13 Click *Finish*.

Typically, the resource 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.

- 14 Verify that the resource was created by viewing its entry in the *Cluster Objects* list on the Cluster Options page.

My Clusters > clus1.ncs.novell

clus1.ncs.novell

View cluster resource configuration information and administer cluster resources for the selected cluster.

Cluster Manager BCC Manager Cluster Event Log Cluster Options

Properties...

Cluster Objects

New | Delete | Details 14 Item(s)

Type	Name	IP Address	Distinguished Name	Pool Name
	Master_IP_Address_Resource	10.10.10.134	cn=clus1,ou=ncs,o=novell	
	avalon	10.10.10.37	cn=avalon,cn=clus1,ou=ncs,o=novell	
	myclustervq01_resource			
	AV_Template			
	DHCP_Template			
	DNS_Template			
	Generic_FS_Template			
	Generic_IP_Service			
	iFolder_Template			
	iPrint_Template			

Close

15 (Optional) Enable and configure resource monitoring:

Resource monitoring is disabled by default. Enable the monitoring function and modify the settings for the resource. For detailed information, see [Section 10.7, “Enabling Monitoring and Configuring the Monitor Script,” on page 192](#).

15a In iManager on the Cluster Manager page or Cluster Options page, select the resource name link to open its Properties dialog box, then select the *Monitoring* tab.

15b Select the *Enable Resource Monitoring* check box to enable resource monitoring for the selected resource.

15c Specify the *Polling Interval* to control how often you want the resource monitor script for this resource to run.

You can specify the value in minutes or seconds. For information, see [“Polling Interval” on page 192](#).

15d Specify the number of failures (*Maximum Local Failures*) for the specified amount of time (*Time Interval*).

For information, see [“Failure Rate” on page 193](#).

15e Specify the *Failover Action* by indicating whether you want the resource to be set to a comatose state, to migrate to another server, or to reboot the hosting node (without synchronizing or unmounting the disks) if a failure action initiates. The reboot option is normally used only for a mission-critical cluster resource that must remain available.

For information, see [“Failure Action” on page 193](#).

16 Bring the resource online:

16a Select *Clusters > My Clusters*.

16b On the Cluster Manager page, select the check box next to the new CSM cluster resource, then click *Online*.

Ensure that the resource successfully enters the *Running* state. If the resource goes comatose instead, it is probably because you made an error when typing values in the script definitions. Typical issues are that the target LVM volume is still active locally on a node, the path names were not entered as case-sensitive, the IP address is not unique in the network, or the mount path does not exist.

Take the resource offline, go to the resource’s *Properties > Scripts* page to review and modify its scripts as needed to correct errors, then try again to online the resource.

My Clusters > clus1.ncs.novell

clus1.ncs.novell

View the status of both servers and resources in the cluster. You can also migrate cluster resources to different servers in the cluster, or change the state of resources to offline or online.

Cluster Manager BCC Manager Cluster Event Log Cluster Options

Run Report

Epoch: 0

avalon

Cluster State						
Online Offline Migrate Respond to Alert Refresh						
2 Item(s)						
Type	Name	State	Location	Lives	Up Since	
<input type="checkbox"/>	Master_IP_Address_Resource	Running	avalon	1	Aug 8, 2012 10:01:33 AM	
<input type="checkbox"/>	vol44_resource	Running	avalon	1		

Close

13.4.3 Sample Generic LVM Resource Scripts

This section contains sample scripts for the Generic LVM resource.

The sample scripts in this section use the following sample parameters. Ensure that you replace the sample values with your values.

Parameter	Sample Value
RESOURCE_IP	10.10.10.136
MOUNT_FS	ext3
VOLGROUP_NAME	myclustervg01
MOUNT_DEV	/dev/\$VOLGROUP_NAME/myclustervol01
MOUNT_POINT	/mnt/myclustervol01

- ♦ [“Sample Generic LVM Resource Load Script” on page 341](#)
- ♦ [“Sample Generic LVM Resource Unload Script” on page 342](#)
- ♦ [“Sample Generic LVM Resource Monitor Script” on page 342](#)

Sample Generic LVM Resource Load Script

Use the following sample load script to complete the fields for your LVM volume group cluster resource:

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

# define the IP address
RESOURCE_IP=10.10.10.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=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 Generic LVM Resource Unload Script

Use the following sample unload script to complete the fields for your LVM volume group cluster resource:

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

# define the IP address
RESOURCE_IP=10.10.10.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=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 Generic 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=10.10.10.136
# 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
```


13.5 Creating a Virtual Server Object for an LVM Volume Group Cluster Resource

You can create an NCP virtual server for an existing clustered LVM volume group cluster resource by using the `/opt/novell/ncs/bin/ncs_ncpserv.py` script to create an `NCS:NCP Server` object. Having a virtual server object allows the resource to be viewed in the *Browse* panel in Novell iManager. You can also bind the virtual server name to the IP address to allow users to access the resource via an assigned name in addition to its IP address.

The `NCS:NCP Server` object does not give users NCP file access to the data on the Linux LVM logical volume. You can add commands to the load script to create an NCP volume on the Linux file system for NCP users. For information, see “NCP Enabling a clustered LVM volume”.

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 13.5.1, “Creating an NCP Virtual Server for the LVM Resource,” on page 343](#)
- [Section 13.5.2, “Adding NCP Virtual Server Commands to the Resource Scripts,” on page 344](#)
- [Section 13.5.3, “Sample LVM Resource Scripts with an NCP Virtual Server,” on page 346](#)

13.5.1 Creating an NCP Virtual Server for the LVM Resource

You use the `/opt/novell/ncs/bin/ncs_ncpserv.py` script to create a virtual server object (`NCS:NCP Server`) in eDirectory for the LVM volume group cluster resource. If the resource does not have NCP volumes on it, you do not use the `-v` option. For information about the `ncs_ncpserv.py` script, see [Section A.7, “Virtual NCS NCP Server Object \(ncs_ncpserv.py\) Script,” on page 433](#).

- 1 On the master cluster node, open a terminal console, then log in as the root user.
- 2 In the console, use the `cd` command to go to the `/opt/novell/ncs/bin` directory.
- 3 At the command prompt, enter

```
./ncs_ncpserv.py -c vg_name -i resource_ip_address
```

Replace the `vg_name` and `resource_ip_address` with the information for your particular solution.

Do not use periods in cluster resource names. 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 13.5.2, “Adding NCP Virtual Server Commands to the Resource Scripts,” on page 344](#).

13.5.2 Adding NCP Virtual Server Commands to the Resource Scripts

After you have created an NCS:NCP Server object, you must modify the resource's load, unload, and monitor scripts to bind and unbind the NCS:NCP Server object with the resource IP address.

- 1 In iManager, select *Clusters > My Clusters*, then select the cluster.
- 2 Click the name link of the LVM volume group cluster resource to open its Cluster Resource Properties page, then click the *Scripts* tab.
- 3 Modify the LVM resource's load script:
 - 3a On the *Scripts* tab, click *Load Script*.
 - 3b At the end of the definition area, add the following lines to define the virtual NCP server name:

```
# define NCP server name
NCP_SERVER=cluster1-myclustervg01-server
```

Replace the NCP server name with the name for your virtual NCP server.

- 3c Under the `add_secondary_ipaddress` command line, add a line to bind the NCP server name to the resource IP address:

```
# bind the NCP server name
exit_on_error ncpcon bind --ncpservname=$NCP_SERVER --ipaddress=$RESOURCE_IP
```

- 3d Verify the order of the commands against the ["Sample LVM Resource Load Script with an NCP Virtual Server" on page 346](#).

- 3e Click *Apply* to save your changes.

The script changes are not active until the next time the cluster resource is taken offline, and then brought online. Do not active the script changes at this time.

- 4 Modify the LVM resource's unload script:

- 4a On the *Scripts* tab, click *Unload Script*.

- 4b At the end of the definition area, add the following lines to define the virtual NCP server name:

```
# define NCP server name
NCP_SERVER=cluster1-myclustervg01-server
```

Specify the same NCP virtual server name that you used in the load script.

- 4c Under the definition, add a line to unbind the NCP server name from the resource IP address:

```
# unbind the NCP server name
ignore_error ncpcon unbind --ncpservname=$NCP_SERVER --ipaddress=$RESOURCE_IP
```

- 4d Verify the order of the commands against the ["Sample LVM Resource Unload Script with an NCP Virtual Server" on page 347](#).

- 4e Click *Apply* to save your changes.

The script changes are not active until the next time the cluster resource is taken offline, and then brought online.

5 Modify the LVM resource's monitor script:

5a On the *Scripts* tab, click *Monitor Script*.

5b At the end of the definition area, add the following lines to define the virtual NCP server name:

```
# define NCP server name
NCP_SERVER=cluster1-myclustervg01-server
```

Specify the same NCP virtual server name that you used in the load script.

5c Verify the order of the commands against the [“Sample LVM Resource Monitor Script with an NCP Virtual Server” on page 348](#).

5d Click *Apply* to save your changes.

The script changes are not active until the next time the cluster resource is taken offline, and then brought online.

6 At the bottom of the Properties dialog box, click *OK* to save your changes.

7 Activate the modified scripts by taking the resource offline and bringing it online.

7a Click the *Cluster Manager* tab.

7b Select the check box next to the LVM volume group cluster resource, then click *Offline*.

Wait until the resource is reports an Offline status before continuing.

7c Select the check box next to the LVM volume group cluster resource, then click *Online*.

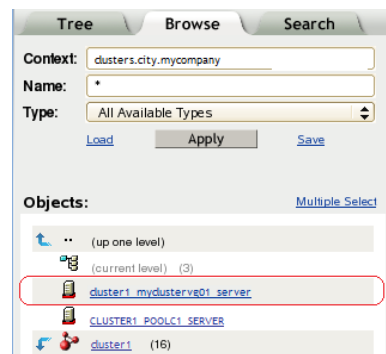
Wait until the resource is reports an Online status before continuing.

8 Verify that an `NCS:NCP Server` object appears in the *Browse* panel in iManager.

8a In the iManager toolbar, click the *View Objects* icon.

8b In the left panel, click *Browse*.

8c Browse to the Cluster container to see the virtual server object for the cluster resource, such as `cluster1-myclustervg01-server`.



9 (Optional) To add NCP file access for the Linux volume, continue with [Section 13.6, “Enabling NCP File Access for a Clustered LVM Volume,” on page 348](#)

13.5.3 Sample LVM Resource Scripts with an NCP Virtual Server

This section contains sample scripts for the Generic LVM resource where commands have been added to bind and unbind an NCP Virtual Server object with the resource IP address.

The sample scripts in this section use the following sample parameters. Ensure that you replace the sample values with your values.

Parameter	Sample Value
RESOURCE_IP	10.10.10.136
MOUNT_FS	ext3
VOLGROUP_NAME	myclustervg01
MOUNT_DEV	/dev/\$VOLGROUP_NAME/myclustervol01
MOUNT_POINT	/mnt/myclustervol01
NCP_SERVER	clus1-myclustervg01-server

- ♦ [“Sample LVM Resource Load Script with an NCP Virtual Server” on page 346](#)
- ♦ [“Sample LVM Resource Unload Script with an NCP Virtual Server” on page 347](#)
- ♦ [“Sample LVM Resource Monitor Script with an NCP Virtual Server” on page 348](#)

Sample LVM Resource Load Script with an NCP Virtual Server

Use the following sample load script to complete the fields for your LVM volume group cluster resource:

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

# define the IP address
RESOURCE_IP=10.10.10.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=myclustervg01
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/myclustervol01
# define the mount point
MOUNT_POINT=/mnt/myclustervol01

# define NCP server name
NCP_SERVER=cluster1-myclustervg01-server

# activate the volume group
exit_on_error vgchange -a ey $VOLGROUP_NAME
```



```

# create the mount point on the node if it does not exist
ignore_error mkdir -p $MOUNT_POINT

# mount the file system
exit_on_error mount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

# bind the NCP server name
exit_on_error ncpcon bind --ncpservname=$NCP_SERVER --ipaddress=$RESOURCE_IP

exit 0

```

Sample LVM Resource Unload Script with an NCP Virtual Server

Use the following sample unload script to complete the fields for your LVM volume group cluster resource:

```

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

# define the IP address
RESOURCE_IP=10.10.10.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=myclustervg01
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/myclustervol01
# define the mount point
MOUNT_POINT=/mnt/myclustervol01

# define NCP server name
NCP_SERVER=cluster1-myclustervg01-server

# unbind the NCP server name
ignore_error ncpcon unbind --ncpservname=$NCP_SERVER --ipaddress=$RESOURCE_IP

# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# unmount the volume
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error umount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

#deactivate the volume group
exit_on_error vgchange -a n $VOLGROUP_NAME

exit 0

```


Sample LVM Resource Monitor Script with an NCP Virtual Server

Use the following sample monitor script to complete the fields for your LVM volume group cluster resource:

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

# define the IP address
RESOURCE_IP=10.10.10.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=myclustervg01
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/myclustervol01
# define the mount point
MOUNT_POINT=/mnt/myclustervol01

# define NCP server name
NCP_SERVER=cluster1-myclustervg01-server

# check the logical volume
exit_on_error status_lv $MOUNT_DEV

# test the file system
exit_on_error status_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# status the IP address
exit_on_error status_secondary_ipaddress $RESOURCE_IP

# check the status of NCP services
rcnstd status
if test $? != 0; then
    exit_on_error rcnstd restart
fi
sleep 5

exit 0
```

13.6 Enabling NCP File Access for a Clustered LVM Volume

After you create an NCP virtual server object for the clustered LVM volume, you can enable NCP file access to a clustered LVM volume by adding NCP commands to the load, unload, and monitor script.

- [Section 13.6.1, “Adding NCP Volume Commands to the Resource Scripts,” on page 348](#)
- [Section 13.6.2, “Creating a Shared NCP Volume Object,” on page 351](#)
- [Section 13.6.3, “Sample LVM Resource Scripts Enabled for NCP File Access,” on page 353](#)

13.6.1 Adding NCP Volume Commands to the Resource Scripts

After you have created an NCS:NCP Server object and the Volume object for the NCP volume, you can add NCP file access for users by adding NCP volume commands to the resource's load, unload, and monitor scripts.

- 1 In iManager, select *Clusters > My Clusters*, then select the cluster.
- 2 Click the name link of the LVM volume group cluster resource to open its Cluster Resource Properties page, then click the *Scripts* tab.

3 Modify the LVM resource's load script:

3a On the *Scripts* tab, click *Load Script*.

3b At the end of the definition area, add the following lines to define the NCP volume name:

```
# define NCP volume name
NCP_VOLUME=MYCLUSTERVOL01
```

Replace the NCP server name with the name for your virtual NCP server. NCP volume names can be up to 14 alphanumeric characters. Supported characters are A to Z, 0 to 9, and underscores (_).

3c After the `add_secondary_ipaddress` command line, add a line to mount the NCP volume:

```
# mount the NCP volume
exit_on_error ncpcon mount $NCP_VOLUME=251,PATH=$MOUNT_POINT
```

3d Verify the order of the commands against the [“Sample LVM Resource Load Script with an NCP Volume” on page 353](#).

3e Click *Apply* to save your changes.

The script changes are not active until the next time the cluster resource is taken offline, and then brought online. Do not active the script changes at this time.

4 Modify the LVM resource's unload script:

4a On the *Scripts* tab, click *Unload Script*.

4b At the end of the definition area, add the following lines to define the NCP volume name:

```
# define NCP volume name
NCP_VOLUME=MYCLUSTERVOL01
```

Specify the same volume name that you used in the load script.

4c Under the `ncpcon unbind` command, add a line to dismount the NCP volume:

```
# dismount the NCP volume
ignore_error ncpcon dismount $NCP_VOLUME
```

4d Verify the order of the commands against the [“Sample LVM Resource Unload Script with an NCP Volume” on page 354](#).

4e Click *Apply* to save your changes.

The script changes are not active until the next time the cluster resource is taken offline, and then brought online.

5 Modify the LVM resource's monitor script:

5a On the *Scripts* tab, click *Monitor Script*.

5b At the end of the definition area, add the following lines to define the NCP volume name:

```
# define NCP volume name
NCP_VOLUME=MYCLUSTERVOL01
```

Specify the same volume name that you used in the load script.


5c Verify the order of the commands against the [“Sample LVM Resource Monitor Script with an NCP Volume” on page 355](#).

5d Click *Apply* to save your changes.

The script changes are not active until the next time the cluster resource is taken offline, and then brought online.

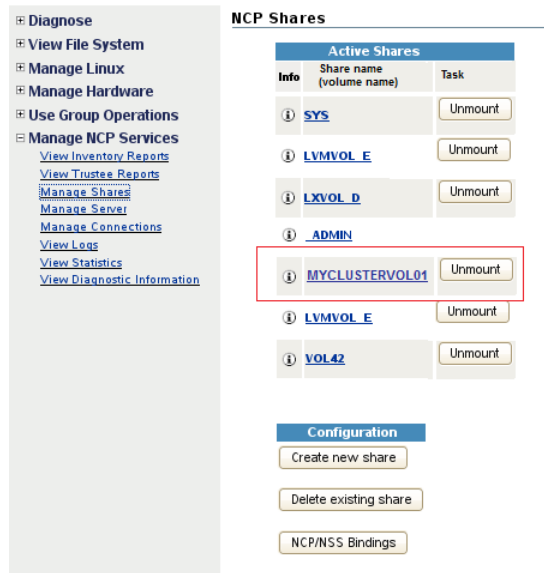
6 At the bottom of the Properties dialog box, click *OK* to save your changes.

- 7 Activate the modified scripts by taking the resource offline and bringing it online.
 - 7a Click the *Cluster Manager* tab.
 - 7b Select the check box next to the LVM volume group cluster resource, then click *Offline*.
Wait until the resource is reports an Offline status before continuing.
 - 7c Select the check box next to the LVM volume group cluster resource, then click *Online*.
Wait until the resource is reports an Online status before continuing.
- 8 Verify that a clustered NCP Volume object appears in the *Browse* panel in iManager.
 - 8a In the iManager toolbar, click the *View Objects* icon.
 - 8b In the left panel, click *Browse*.
 - 8c Browse to the cluster container to see the NCP volume, such as `clus1_MYCLUSTERVG01`.



- 9 Verify the NCP volume by using the Manage NCP Services plug-in for Novell Remote Manager:
 - 9a Log in as the root user to Novell Remote Manager.
 - 9b Select *Manage NCP Services*, then select *Manage Shares*.
 - 9c Locate the NCP volume in the *Active Shares* list.

The clustered NCP volume is defined in the LVM volume group cluster resource load script. The NCP volume is mounted when you bring the resource online, and dismounted when you take the resource offline. You do not use the *Mount* and *Unmount* buttons in this interface.



- 9d Click the *Information* (i) icon to the left of the NCP volume to view details about the NCP volume and mount point.

MYCLUSTERVOL01 Share Information

Information	
Description	Value
File system path	/mnt/myclustervol01
File system shadow path	n/a
Loaded name spaces	DOS LONG
File system type	EXT3
NCP volume ID	253
Status	mounted cluster resource
Sector Size	512
Sectors per Cluster	8
Capacity	665.19 MB

Advanced Information [View](#)

[Open File Information](#)

Volume tasks

Available Actions

[Add Shadow volume](#)

[Perform Inventory](#)

[Share Management Home](#)

9e Exit Novell Remote Manager.

13.6.2 Creating a Shared NCP Volume Object

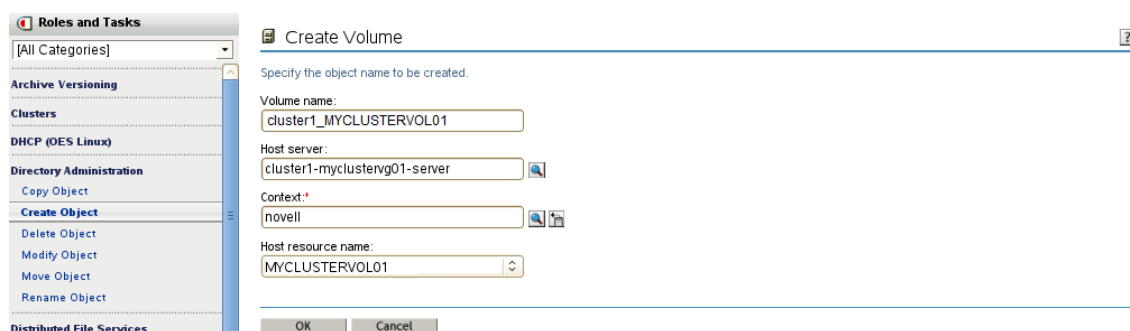
Use the following procedure to create an NCP volume object on the NCP virtual server you created in [Section 13.5, “Creating a Virtual Server Object for an LVM Volume Group Cluster Resource,”](#) on page 343

- 1 Log in to iManager as a cluster administrator.
- 2 Select *Directory Administration > Create Object*.
- 3 On the Create Object page, select Volume, then click OK.

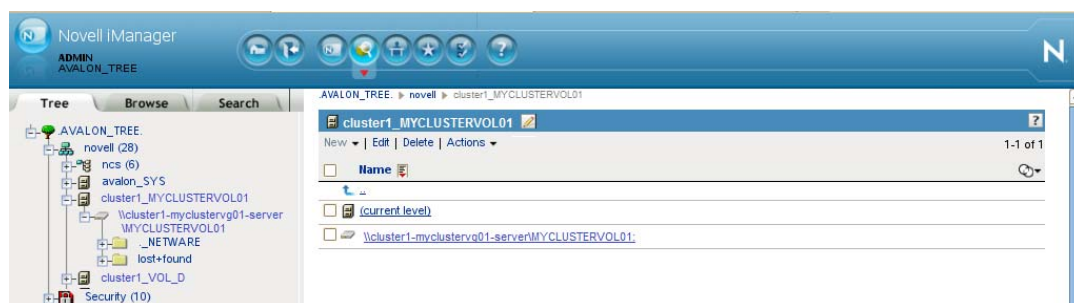
- 4 On the Create Volume page, complete the following information, then click OK.

Parameter	Action
Volume name	Specify the name that you defined in the load script in the format <code><cluster_name>_<ncp_volumen_name></code> , such as <code>cluster1_MYCLUSTERVOL01</code> .

Parameter	Action
Host server	Browse to select the NCP virtual server object for the LVM cluster resource, such as <code>cluster1-myclustervg01-server</code> .
Context	Browse to select the context where you want to create the Volume object for the NCP volume, such as <code>novell</code> .
Host resource name	From the drop-down list of available NCP volumes on the server, select the NCP volume that you defined in the load script, such as <code>MYCLUSTERVOL01</code> .



- 5 After the Volume object is successfully created, click *OK* to dismiss the confirmation message.
- 6 Verify that the object is accessible.
 - 6a In the iManager toolbar, select the *View Objects* icon.
 - 6b Browse in the Tree view to locate and view information about the NCP volume's Volume object.



13.6.3 Sample LVM Resource Scripts Enabled for NCP File Access

This section contains sample scripts for the Generic LVM resource where commands have been added to bind and unbind an NCP Virtual Server object with the resource IP address and to mount and dismount an NCP volume for the Linux file system.

The sample scripts in this section use the following sample parameters. Ensure that you replace the sample values with your values.

Parameter	Sample Value
RESOURCE_IP	10.10.10.136
MOUNT_FS	ext3
VOLGROUP_NAME	myclustervg01
MOUNT_DEV	/dev/\$VOLGROUP_NAME/myclustervol01
MOUNT_POINT	/mnt/myclustervol01
NCP_SERVER	clus1-myclustervg01-server
NCP_VOLUME	MYCLUSTERVOL01
Volume ID	251
	Ensure that the volume ID that you assign to the NCP volume is unique across all nodes in the cluster. Clustered NCP volume IDs are assigned from 254 to 0, in descending order.

- ♦ [“Sample LVM Resource Load Script with an NCP Volume” on page 353](#)
- ♦ [“Sample LVM Resource Unload Script with an NCP Volume” on page 354](#)
- ♦ [“Sample LVM Resource Monitor Script with an NCP Volume” on page 355](#)

Sample LVM Resource Load Script with an NCP Volume

Use the following sample load script to complete the fields for your LVM volume group cluster resource that has an NCP virtual server and an NCP volume:

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

# define the IP address
RESOURCE_IP=10.10.10.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=myclustervg01
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/myclustervol01
# define the mount point
MOUNT_POINT=/mnt/myclustervol01

# define NCP server name
NCP_SERVER=cluster1-myclustervg01-server
# define NCP volume name
NCP_VOLUME=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

# mount the NCP volume
exit_on_error ncpcon mount $NCP_VOLUME=251,PATH=$MOUNT_POINT

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

# bind the NCP server name
exit_on_error ncpcon bind --ncpservname=$NCP_SERVER --ipaddress=$RESOURCE_IP

exit 0

```

Sample LVM Resource Unload Script with an NCP Volume

Use the following sample unload script to complete the fields for your LVM volume group cluster resource that has an NCP virtual server and an NCP volume:

```

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

# define the IP address
RESOURCE_IP=10.10.10.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=myclustervg01
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/myclustervol01
# define the mount point
MOUNT_POINT=/mnt/myclustervol01

# define NCP server name
NCP_SERVER=cluster1-myclustervg01-server
# define NCP volume name
NCP_VOLUME=MYCLUSTERVOL01

# unbind the NCP server name
ignore_error ncpcon unbind --ncpservname=$NCP_SERVER --ipaddress=$RESOURCE_IP

# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# dismount the NCP volume
ignore_error ncpcon dismount $NCP_VOLUME

# unmount the Linux volume
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error umount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

#deactivate the volume group
exit_on_error vgchange -a n $VOLGROUP_NAME

exit 0

```


Sample LVM Resource Monitor Script with an NCP Volume

Use the following sample monitor script to complete the fields for your LVM volume group cluster resource that has an NCP virtual server and an NCP volume:

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

# define the IP address
RESOURCE_IP=10.10.10.136
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=myclustervg01
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/myclustervol01
# define the mount point
MOUNT_POINT=/mnt/myclustervol01

# define NCP server name
NCP_SERVER=cluster1-myclustervg01-server
# define NCP volume name
NCP_VOLUME=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

# check the status of NCP services
rcnstd status
  if test $? != 0; then
    exit_on_error rcnstd restart
  fi
sleep 5

exit 0
```

13.7 Modifying the LVM Resource Scripts for Other Services

After you create an LVM volume group cluster resource, you can add lines to its load script, unload script, and monitor script to customize the resource for other uses. Compare the Generic File System template to the resource template for your product to determine which lines need to be added or modified.

The Novell Cluster Services resource templates are available in the Clusters plug-in for Novell iManager. From the Cluster Options page, select the resource template to view its properties, then click the *Scripts* tab to access its default load, unload, and monitor scripts. You can also select the check box next to the resource template, then click *Details* to access its properties. For information, see [Section 10.3, “Using Cluster Resource Templates,” on page 184](#).

13.8 Renaming the Mount Point Path for a Clustered LVM Volume

For a clustered LVM volume, you can rename the mount point path by modifying the mount point variable in the cluster load, unload, and monitor scripts.

- 1 Open Novell iManager in a Web browser, then log in as an administrator user.
- 2 In *Roles and Tasks*, select *Clusters > My Clusters*.
- 3 Select the cluster where you created the volume group cluster resource.
- 4 On the Cluster Manager page, select the check box next to the resource, then click *Offline*.
Wait for the status to report that it is offline, then continue.
- 5 Modify the mount point path value in the load, unload, and monitor scripts for the LVM volume group cluster resource:

- 5a On the Cluster Manager page, select the resource's name link to open its Cluster Properties page, then click the *Scripts* tab.

The Scripts tab automatically displays the load script.

- 5b Modify the load script:

- 5b1 In the load script, type the new value for the mount point in the MOUNT_POINT variable:

```
MOUNT_POINT=/usr/novell/ext3/vol44
```

- 5b2 Ensure that a `mkdir` command is added above the Linux `mount` command line in the load script to create the path on nodes if it does not exist.

```
# create the mount point if it does not exist on the node
ignore_error mkdir -p $MOUNT_POINT
```

You can alternatively make the new path by using the `mkdir` command in a terminal console on each node. If the master node is not the most preferred node, ensure that you make the path before you bring the resource online.

- 5b3 Click *Apply*.

- 5c Click the *Unload Script* link, type the new value for the mount point in the MOUNT_POINT variable, then click *Apply*.

```
MOUNT_POINT=/usr/novell/ext3/vol44
```

- 5d Click the *Monitor Script* link, type the new value for the mount point in the MOUNT_POINT variable, then click *Apply*.

```
MOUNT_POINT=/usr/novell/ext3/vol44
```

- 5e At the bottom of the page, click *OK* to close the Properties page and save your changes.
The changes do not take effect until the resource is brought online.

- 6 Bring the resource online to allow the script changes to take effect.

- 6a In *Roles and Tasks*, select *Clusters > My Clusters*, then select the cluster.

- 6b On the Cluster Manager page, select the check box next to the resource, then click *Online*.

- 6c Verify that the resource comes online and reports a *Running* state.

If the resource goes into a Comatose state, it is probably because you made a mistake in the lines you added or modified in the scripts. Take the resource offline, then go back to correct the scripts, and try to bring it online again.

- 7 In NSSMU, verify that the new mount point is used when the clustered LVM volume resource is brought online:

7a Log in as the root user to the node that is hosting the resource, then start NSSMU by entering:

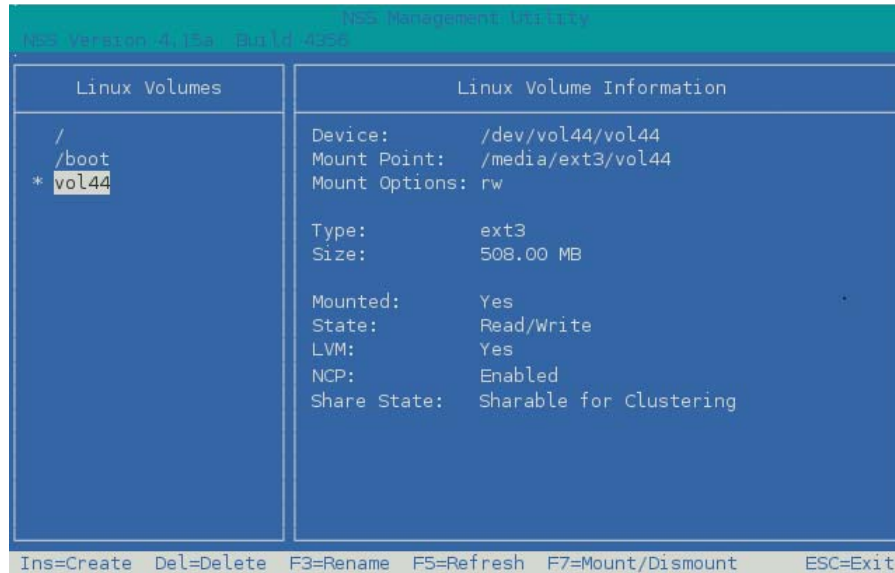
```
nssmu
```

7b From the NSSMU main menu, select *Linux Volumes*, then press Enter.

7c In the *Linux Volumes* list, select the clustered LVM volume.

7d View the volume details to verify that the mount point has changed.

For example, the mount point is `/usr/novell/ext3/vol44`.



7e Click Escape twice to exit NSSMU.

13.9 Disabling Clustering for an LVM Volume Group and Logical Volume

Before you attempt to disable clustering for an LVM volume group and logical volume, you must take the resource offline.

WARNING: If you attempt to delete a cluster resource without first offlining it, deletion errors occur, and the data associated with the clustered volume group might not be recoverable.

Deleting a volume group cluster resource disables only the Novell Cluster Services cluster settings. There are additional tasks to perform to prepare the LVM volume group and logical volume for non-clustered access on one of the cluster nodes. To delete a resource and create a new one with the same name, you must wait to create the new one until after Novell eDirectory synchronizes all of the objects in the tree related to the deleted resource.

If the volume was NCP enabled, the definition of the NCP volume in the load script is automatically deleted, but you must manually remove its NCP Volume object in Novell eDirectory. To re-create the NCP volume on the non-clustered Linux LVM logical volume, can use the Manage NCP Services plug-in for Novell Remote Manager or use the `ncpcon create volume` command.

All resource configuration should happen from the master node. On the Cluster Options page for Novell iManager, you are automatically connected to the Cluster object, which is associated with the master node.

We strongly recommend that when you need to delete a cluster resource, that you do so only from the master node in the cluster. You might want to delete a cluster resource on a non-master node if the master node is not a preferred node for the resource, or if there is a resource mutual exclusion rule that prohibits the resource from being moved there. If the resource cannot be migrated to the master node, follow the additional steps that are annotated as “Non-Master”.

Use the procedure in this section to disable clustering for a Linux LVM volume group and logical volume. Afterwards, you can mount and dismount the volume on only one node in the cluster.

- 1 Ensure that the resource is online on the node where you want mount the volume after clustering is disabled.

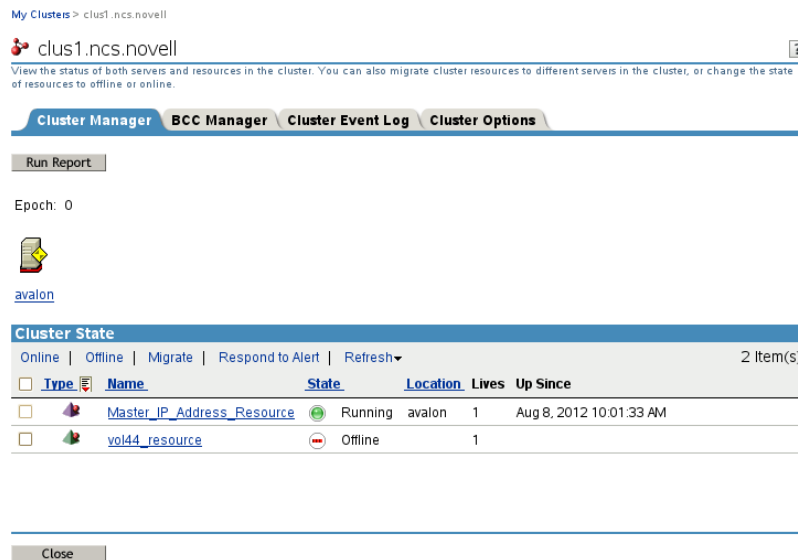
If the resource is not on the desired node, open a terminal console as the `root` user, then enter

```
cluster migrate <resource_name> <from_node> <to_node>
```

For example, if node2 is the to node, enter

```
cluster migrate vg44_resource node1 node2
```

- 2 Log in to iManager as a cluster administrator.
- 3 In *Roles and Tasks*, select *Clusters > My Clusters*, then select the cluster.
- 4 Take the volume group cluster resource offline, by using one of the following methods:
 - ♦ In iManager on the Cluster Manager page, select the check box next to the LVM volume group cluster resource (such as `vg44_resource`), then click *Offline*. Wait until the resource is offline to continue.



- ♦ Open a terminal console as the `root` user, then enter

```
cluster offline <resource_name>
```

The unload script unbinds the NCP virtual server from the resource IP address, dismounts the NCP volume, dismounts the Linux LVM logical volume, and deactivates the Linux LVM volume group.

5 Delete the volume group cluster resource:

5a (Non-Master) If the volume group was taken offline on a non-master node, verify that the revision and resource information is the same on the non-master node as on the master node before you delete the resource.

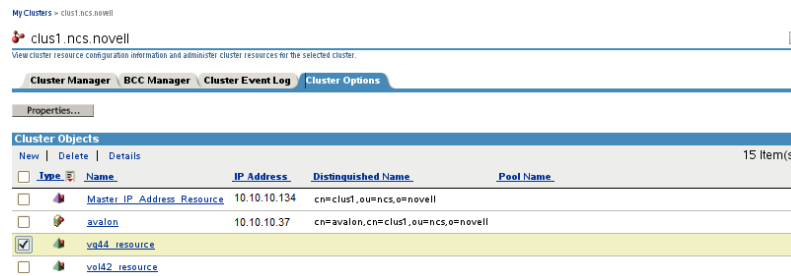
5a1 On the non-master node, log in as the `root` user, then open a terminal console. At a command prompt on the non-master node, enter

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


5a2 Look at the file `/var/opt/novell/ncs/resource-priority.conf` to verify that it has the same information (REVISION and NUMRESOURCES) as the file on the master node.

5b Click the *Cluster Options* tab.

5c Select the check box next to the volume group cluster resource, then click *Delete*.



5d When you are prompted to confirm the deletion, click *OK* to continue, or click *Cancel* to abort the deletion.



5e (Non-Master) If the volume group was taken offline on a non-master node, verify that the revision and resource information is the same on the master node as on the non-master node after you delete the resource.

5e1 On the master node, log in as the `root` user, then open a terminal console. At a command prompt on the master node, enter

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

5e2 Look at the file `/var/opt/novell/ncs/resource-priority.conf` to verify that it has the same information (REVISION and NUMRESOURCES) as that of the non-master node where you deleted the cluster resource.

5e3 In iManager, select *Clusters > Cluster Options*, then browse to select the Cluster object.

5e4 Click *Properties*, select the *Priorities* tab, then click *Apply* on the Priorities page.

5e5 At a command prompt, enter

```
cluster view
```

The cluster view should be consistent.

- 5e6** Look at the file `/var/opt/novell/ncs/resource-priority.conf` on the master node to verify that the revision number increased.

If the revision number increased, you are done. Do not continue with [Step 5e7](#).

If the deleted resource is the only one in the cluster, the priority won't force the update. A phantom resource might appear in the interface. You need to restart Cluster Services to force the update, which also removes the phantom resource.

- 5e7** If the revision number did not automatically update in the previous steps, restart 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.

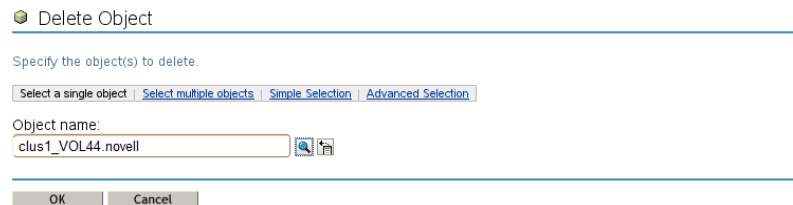
For example:

```
cluster restart 120
```

- 6** If the Linux volume was NCP enabled, delete the NCP volume object.

6a In iManager, select *Directory Administration > Delete Object*.

6b On the Delete Object page, browse to locate and select the resource's NCP Volume object, such as `clus1_VOL44.novell`, then click **OK**.

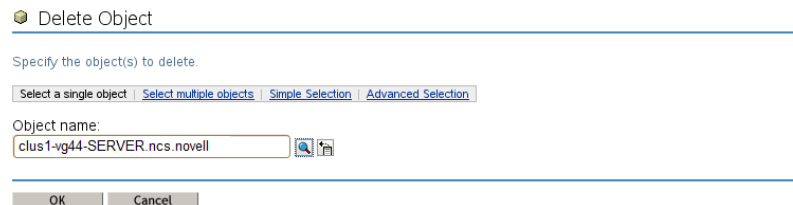


6c Click **OK** to delete the NCP volume object, then click **OK** again to close the success message.

- 7** If the Linux volume was NCP enabled, or if you manually created an NCP virtual server for the resource, delete the NCP virtual server object:

7a In iManager, select *Directory Administration > Delete Object*.

7b On the Delete Object page, browse to locate and select the resource's NCS:NCP Server object, such as `clus1-vg44-SERVER.ncs.novell`, then click **OK**.



7c Click **OK** to delete the NCP virtual server, then click **OK** again to close the success message.

7d Visually confirm that the resource does not appear in the list of Cluster objects, then exit iManager.

- 8** Log in as the `root` user on the server where the resource was online, then open a terminal console.

- 9** Deactive the volume group from Clustered LVM:

```
vgchange -c n <vg_name>
```

For example:


```
vgchange -c n vg44
```

- 10** Activate the volume group locally on the current node:

```
vgchange -a ey <vg_name>
```

For example:

```
vgchange -a ey vg44
```

- 11** Mount the LVM logical volume locally on the current node:

```
mount -t <fstype> <device_name> <full_mount_point_path>
```

Replace *fstype* with the file system type of the volume.

Replace *device_name* with the full device path of the logical volume, such as `/dev/<vg_name>/<lv_name>`.

Replace *full_mount_point_path* with the volume's mount point.

For example:

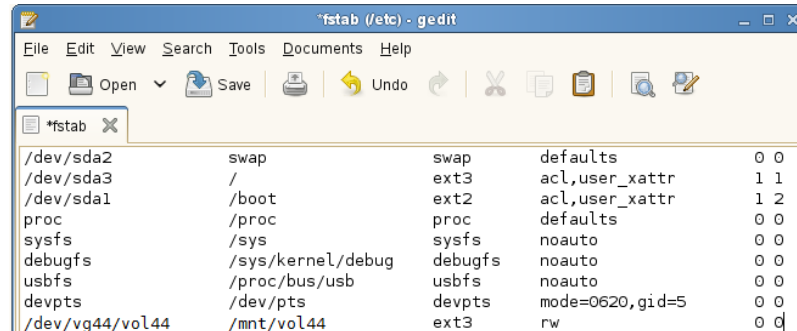
```
mount -t ext3 /dev/vg44/vol144 /mnt/vol144
```

- 12** Add an entry for the LVM volume in the `/etc/fstab` file to allow the volume to be mounted automatically on reboot. It also provides the automatic mount information for NSSMU to use to mount and dismount the volume on the node.

In a text editor, modify the `/etc/fstab` file to specify the mount point information and file system type.

For example, complete the line for the volume's device path, such as:

```
/dev/vg44/vol144    /mnt/vol144        ext3                rw                  0 0
```



- 13** View details about the non-clustered LVM logical volume:

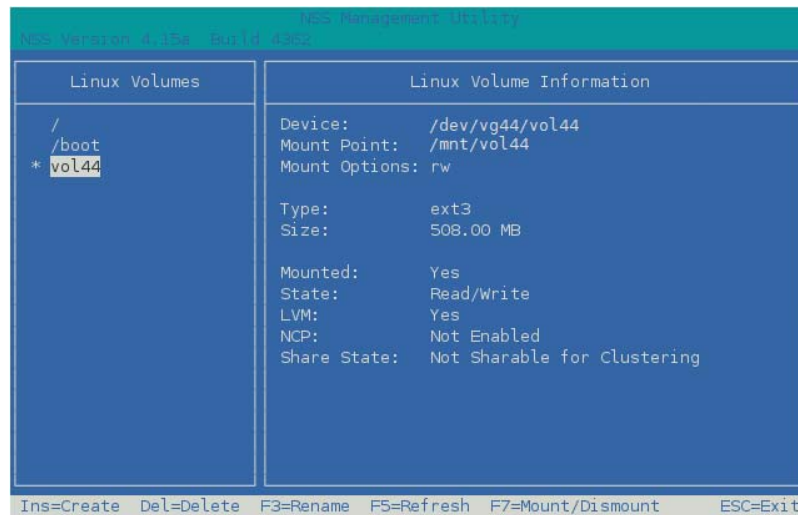
- 13a** Open an terminal console as the `root` user, launch NSSMU:

```
nssmu
```

- 13b** In the NSSMU main menu, select *Linux Volumes*, then press Enter.

13c On the Linux Volumes page, then select the volume to view its details.

The volume is no longer cluster enabled. You can mount and dismount the volume only on the current node.



The screenshot shows the NSS Management Utility interface. The title bar reads "NSS Management Utility" and "NSS Version 4.154 Build 4352". The interface is divided into two main panels. The left panel, titled "Linux Volumes", contains a list of volumes: "/", "/boot", and "* vol44". The "vol44" entry is highlighted. The right panel, titled "Linux Volume Information", displays the following details for the selected volume: Device: /dev/vg44/vol44, Mount Point: /mnt/vol44, Mount Options: rw, Type: ext3, Size: 508.00 MB, Mounted: Yes, State: Read/Write, LVM: Yes, NCP: Not Enabled, and Share State: Not Sharable for Clustering. At the bottom of the window, a status bar shows keyboard shortcuts: Ins=Create, Del=Delete, F3=Rename, F5=Refresh, F7=Mount/Dismount, and ESC=Exit.

Linux Volumes	Linux Volume Information
/	
/boot	
* vol44	Device: /dev/vg44/vol44 Mount Point: /mnt/vol44 Mount Options: rw Type: ext3 Size: 508.00 MB Mounted: Yes State: Read/Write LVM: Yes NCP: Not Enabled Share State: Not Sharable for Clustering

14 (Optional) Create an NCP volume for the Linux volume by using Novell Remote Manager. For information, see [“Creating NCP Volumes on Linux File Systems”](#) in the *OES 11 SP1: NCP Server for Linux Administration Guide*.

To use the same volume name that was used in the cluster resource load script, wait to re-create the NCP volume until after Novell eDirectory synchronizes all of the objects in the tree related to the deleted resource.

13.10 Deleting a Clustered LVM Volume Group and Logical Volume


Before you delete a clustered LVM volume group, you must take the volume group cluster resource offline and delete the cluster resource object and related objects in Novell eDirectory.

All resource configuration must happen from the master node. On the Cluster Options page for Novell iManager, you are automatically connected to the Cluster object, which is associated with the master node.

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

3 Take the volume group cluster resource offline:


- 3a** On the Cluster Manager page, select the check box next to the volume group cluster resource, then click *Offline*. Wait until the resource is offline to continue.




The unload script unbinds the NCP virtual server from the resource IP address, dismounts the NCP volume, dismounts the Linux LVM logical volume, and deactivates the Linux LVM volume group.

4 Delete the volume group cluster resource:

- 4a** Click the *Cluster Options* tab.
- 4b** Select the check box next to the volume group cluster resource, then click *Delete*.



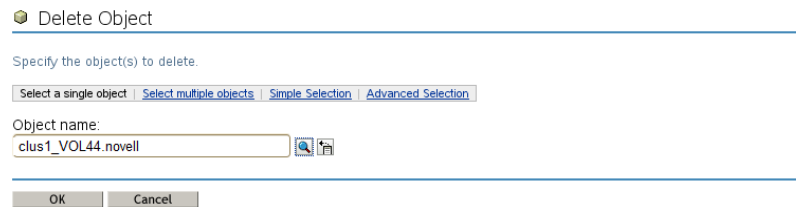
- 4c** When you are prompted to confirm the deletion, click *OK* to continue, or click *Cancel* to abort the deletion.



5 If the Linux volume was enabled for NCP file access, delete the NCP volume object.

5a In iManager, select *Directory Administration > Delete Object*.

5b On the Delete Object page, browse to locate and select the resource's NCP Volume object, such as `clus1_VOL44.novell`, then click *OK*.

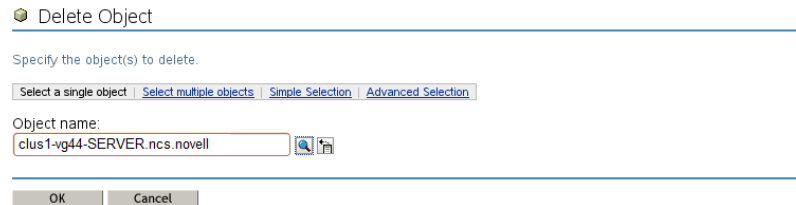


5c Click *OK* to delete the NCP volume object, then click *OK* again to close the success message.

6 If the Linux volume was NCP enabled, or if you manually created an NCP virtual server for the resource, delete the NCP virtual server object.

6a In iManager, select *Directory Administration > Delete Object*.

6b On the Delete Object page, browse to locate and select the resource's NCS:NCP Server object, such as `clus1-vg44-SERVER.ncs.novell`, then click *OK*.



6c Click *OK* to delete the NCP virtual server, then click *OK* again to close the success message.

7 (Optional) Delete the LVM volume group and logical volume:


7a Log in as the `root` user to the cluster node where the resource was online, then open a terminal console.

7b At a command prompt, launch NSSMU:

```
nssmu
```

7c In the NSSMU main menu, select *Linux Volumes*, then press *Enter*.

7d In the *Linux Volumes* list, select the clustered linux volume, then press *Delete*.



7e When you are prompted to confirm the delete action, press *y* (Yes) to continue, or press *n* (No) to cancel the delete action.


Deleting the volume and volume group puts the device in an uninitialized state.

7f Press *Esc* to return to the main menu

8 Re-initialize the device:

8a In the NSSMU main menu, select *Devices* and press *Enter*.

8b In the *Devices* list, select the device.



8c Press *F3* to initialize the device.

8d When you are prompted to confirm the initialization action, read the warning message, then press *y* (Yes) to continue, or press *n* (No) to cancel the action.

8e Select *DOS* or *GPT* as the partitioning type, then press *Enter*.

8f Visually verify that the device was initialized and the specified partition type is displayed.

8g Press *Esc* twice to exit NSSMU.

- 9 Verify that the LVM volume group is deleted by entering the following the native LVM command:

```
vgdisplay
```

13.11 Deleting a Clustered LVM Volume (Created in NSSMU or NLVM)

We strongly recommend that you delete a cluster-enabled LVM volume only from the master node in the cluster. This allows the cluster information to be automatically updated.

The procedures in this section assume that you created the clustered LVM volume in NSSMU or NLVM. There are default naming conventions applied by these tools that might not apply to LVM volume groups that you created and cluster-enabled by using native LVM tools and the Generic File System template.

WARNING: Deleting an LVM volume destroys all data on it.

NSSMU and the `nlvm delete linux volume <volume_name>` command delete the cluster-enabled LVM volume as well as the related objects in eDirectory:

- ♦ Linux LVM volume group and logical volume from the file system
- ♦ Cluster Resource object for the LVM resource
- ♦ If the LVM volume is NCP-enabled:
 - ♦ Volume object for the LVM volume
 - ♦ Virtual server for the LVM resource (NCS:NCP Server object)

When the LVM volume resides on the master node, the cluster information is automatically updated.

When the LVM volume resides on a non-master node, additional steps are required to update the cluster information. A cluster restart might be needed to force the information to be updated.

Use the following procedures to delete a cluster-enabled LVM volume:

- ♦ [Section 13.11.1, “Deleting a Cluster-Enabled LVM Volume on the Master Node,” on page 366](#)
- ♦ [Section 13.11.2, “Deleting a Cluster-Enabled LVM Volume on a Non-Master Node,” on page 367](#)

13.11.1 Deleting a Cluster-Enabled LVM Volume on the Master Node

- 1 If the LVM resource is on a non-master node in the cluster, migrate it to the master node. As the root user, open a terminal console, then enter

```
cluster migrate <resource_name> <master_node_name>
```

To migrate the resource, the master node must be in the resource’s preferred nodes list.

- 2 Use the `cluster status` command to check the resource status. If the resource is online or comatose, take it offline.

As the root user, enter

```
cluster offline <resource_name>
```

Use the `cluster status` command to verify that the resource has a status of *Offline* before you continue.

- 3 Delete the LVM volume on the master node by using NSSMU.
You can alternatively use the `nlvm delete linux volume <lx_volume_name>` command.
 - 3a In NSSMU, select *Linux Volumes*, then press Enter.
 - 3b Select the unmounted LVM volume, then press Delete.
 - 3c Select OK to confirm, then press Enter.
- 4 In the *Tree View* in iManager, browse the objects to verify that the following objects were deleted:
 - ♦ LVM resource object (from the Cluster container)
 - ♦ If the LVM volume was NCP-enabled:
 - ♦ Volume object for the LVM volume
 - ♦ Virtual server for the LVM resource (NCS:NCP Server object)
- 5 Re-initialize the device that contained the LVM volume.
When NLVM or NSSMU removes the Linux LVM volume group, it leaves the device in an uninitialized state.
 - 5a In NSSMU, select *Devices*, then press Enter.
 - 5b Select the device, then press F3 (Initialize).
 - 5c Press y (Yes) to confirm.
 - 5d Select the partitioning scheme as DOS or GPT, then press Enter.
- 6 (Optional) Use a third-party SAN management tool to assign the device to only one desired server.

13.11.2 Deleting a Cluster-Enabled LVM Volume on a Non-Master Node

- 1 Log in as the root user to the non-master node where the cluster resource currently resides, then open a terminal console.
- 2 Use the `cluster status` command to check the resource status. If the resource is online or comatose, take it offline by using one of the following methods:


```
cluster offline <resource_name>
```

Use the `cluster status` command to verify that the resource has a status of *Offline* before you continue.
- 3 At the command prompt on the non-master node, enter


```
/opt/novell/ncs/bin/ncs-configd.py -init
```
- 4 Look at the file `/var/opt/novell/ncs/resource-priority.conf` to verify that it has the same information (REVISION and NUMRESOURCES) as the file on the master node.
- 5 Delete the LVM volume on the master node by using NSSMU.
You can alternatively use the `nlvm delete linux volume <lx_volume_name>` command.
 - 5a In NSSMU, select *Linux Volumes*, then press Enter.
 - 5b Select the unmounted LVM volume, then press Delete.
 - 5c Select OK to confirm, then press Enter.
- 6 In the *Tree View* in iManager, browse the objects to verify that the following objects were deleted:
 - ♦ LVM resource object (from the Cluster container)

- ♦ If the LVM volume was NCP-enabled:
 - ♦ Volume object for the LVM volume
 - ♦ Virtual server for the LVM resource (NCS:NCP Server object)
- 7 Re-initialize the device that contained the LVM volume.
 When NLVM or NSSMU removes the Linux LVM volume group, it leaves the device in an uninitialized state.
 - 7a In NSSMU, select *Devices*, then press Enter.
 - 7b Select the device, then press F3 (Initialize).
 - 7c Press y (Yes) to confirm.
 - 7d Select the partitioning scheme as DOS or GPT, then press Enter.
- 8 On the master node, log in as the root user, open a terminal console, then enter


```
/opt/novell/ncs/bin/ncs-configd.py -init
```
- 9 Look at the file `/var/opt/novell/ncs/resource-priority.conf` to verify that it has the same information (REVISION and NUMRESOURCES) as that of the non-master node where you deleted the cluster resource.
- 10 In iManager, select *Clusters > My Clusters*, select the cluster, then select the *Cluster Options* tab.
- 11 Click *Properties*, select the *Priorities* tab, then click *Apply* on the Priorities page.
- 12 At a command prompt, enter


```
cluster view
```

The cluster view should be consistent.
- 13 Look at the file `/var/opt/novell/ncs/resource-priority.conf` on the master node to verify that the revision number increased.
 If the revision number increased, skip [Step 14](#).
 If the deleted resource is the only one in the cluster, the priority won't force the update. A phantom resource might appear in the interface. You need to restart Cluster Services to force the update, which also removes the phantom resource.
- 14 If the revision number did not automatically update in the previous steps, restart 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
```
- 15 (Optional) Use a third-party SAN management tool to assign the devices to only the desired server.

13.12 Linux LVM Management Tools

LVM tools are available in the YaST Expert Partitioner and in native Linux command line commands.

- ♦ [Section 13.12.1, "Using the YaST Expert Partitioner," on page 369](#)
- ♦ [Section 13.12.2, "Using LVM Commands," on page 369](#)

13.12.1 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 SP2: Storage Administration Guide* (http://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html).

13.12.2 Using LVM Commands

For information about using LVM commands, see the man pages for the commands described in [Table 13-1](#). Perform the commands as the root user.

Table 13-1 LVM Commands

Command	Description
<code>pvcreate <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.
<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>	Activates (<code>-a ey</code>) or deactivates (<code>-a n</code>) a volume group and its logical volumes for input/output. IMPORTANT: Ensure that you use the <code>ey</code> option to exclusively activate a volume group on a cluster node. This option is used by default in the load script.
<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>	Displays information about a specified volume group. To find the total physical extent of a volume group, enter <code>vgdisplay vg_name grep "Total PE"</code>
<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>	Removes a logical volume, such as <code>/dev/vg_name/lv_name</code> . Before using this command, close the logical volume by dismounting it with the <code>umount</code> command.
<code>vgextend <vg_name><device></code>	Adds a specified physical volume to an existing volume group

Command	Description
<code>vgreduce <vg_name> <device></code>	Removes a specified physical volume from an existing volume group. IMPORTANT: Ensure that the physical volume is not currently being used by a logical volume. If it is, you must move the data to another physical volume by using the <code>pvmove</code> command.
<code>lvextend -L size</dev/vg_name/lv_name></code>	Extends the size of a specified logical volume. Afterwards, you must also expand the file system to take advantage of the newly available space.
<code>lvreduce -L size </dev/vg_name/lv_name></code>	Reduces the size of a specified logical volume. IMPORTANT: Ensure that you reduce the size of the file system first before shrinking the volume, otherwise you risk losing data.
<code>lvrename </dev/vg_name/old_lv_name> </dev/vg_name/new_lv_name></code>	Renames an existing LVM logical volume in a volume group from the old volume name to the new volume name. It does not change the volume group name.

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

Novell Cluster Services for Novell Open Enterprise Server (OES) 11 and later 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 SP3 or earlier OES servers. This section describes how to modify the scripts and cluster settings for the existing resources so they can run on OES 11 or later clusters.

IMPORTANT: Beginning in OES 11, Novell Cluster Services creates LVM volume group resources for Linux POSIX file systems. To create new resources, see [Chapter 13, “Configuring and Managing Cluster Resources for Shared LVM Volume Groups,”](#) on page 305.

- ♦ [Section 14.1, “Requirements for Using CSM Cluster Resources,”](#) on page 371
- ♦ [Section 14.2, “Identifying CSM-Based Resources,”](#) on page 372
- ♦ [Section 14.3, “Deporting the CSM Containers with Segment Managers,”](#) on page 373
- ♦ [Section 14.4, “Modifying the Scripts for CSM Resources without a Segment Manager,”](#) on page 374
- ♦ [Section 14.5, “Modifying the Scripts for CSM Resources with a Segment Manager,”](#) on page 378
- ♦ [Section 14.6, “Configuring and Adding OES 11 SP1 Nodes to the OES 2 SP3 Cluster,”](#) on page 385
- ♦ [Section 14.7, “Configuring the Preferred Nodes and Cluster Settings for a CSM Cluster Resource,”](#) on page 386
- ♦ [Section 14.8, “Verifying the Load and Unload Scripts,”](#) on page 387
- ♦ [Section 14.9, “Creating an NCP Virtual Server Object for a CSM Cluster Resource,”](#) on page 387
- ♦ [Section 14.10, “Deleting a Cluster Resource,”](#) on page 390

14.1 Requirements for Using CSM Cluster Resources

On OES 1 and 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 and later.

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 SP3 nodes to OES 11 SP1 nodes in a mixed-mode cluster, and use them after upgrading the cluster to OES 11 SP1:

- ♦ Recognizes and supports the EVMS-based file system structure.

- Provides a way for managing a resource in the load, unload, and monitor scripts on OES 11 or later clusters that works with its existing 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. The commands activate, deactivate, or check the status of the CSM container. You must add `csmport` commands in resource's the load script, unload script, and monitor script before you bring the resource online on OES 11 SP1 servers during a cluster upgrade. 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 427](#), or see the `csmport(8)` man page.

A mixed-mode cluster is supported as a temporary configuration scenario for converting an OES 2 SP3 cluster to an OES 11 SP1 cluster. After the CSM-based resource is configured to run on OES 11 SP1, it should fail over only to the OES 11 SP1 nodes in the mixed-mode cluster. For information, see [Section 7.1.2, "Linux POSIX File Systems on an EVMS CSM \(Compatibility Only\)," on page 106](#).

You cannot create new CSM-based cluster resources on OES 11 or later clusters. For new Linux POSIX volume resources, you must use the Linux Logical Volume Manager (LVM) volume groups as described in [Chapter 13, "Configuring and Managing Cluster Resources for Shared LVM Volume Groups," on page 305](#).

14.2 Identifying CSM-Based Resources

[Table 14-1](#) lists the typical cluster resources for OES 2 SP3 and identifies whether the resource might manage a Linux POSIX file system on a CSM container. After you identify the resources on your OES 2 SP3 cluster, continue with [Section 14.3, "Deporting the CSM Containers with Segment Managers," on page 373](#). Do not add OES 11 SP1 nodes to the OES 2 SP3 cluster at this time.

Table 14-1 OES 2 SP3 Cluster Resources that Manage Linux POSIX File Systems on CSM Containers

OES 2 SP3 Cluster Resources	Uses a Linux POSIX file system on a CSM container in OES 2 SP3
Archive and Version Services	Yes The cluster resource manages a clustered Linux POSIX volume for the PostgreSQL database and a clustered NSS pool and volume for the Archive volume.
DFS VLDB (Distributed File Services volume location database)	No
DHCP Server	Possible The DHCP resource template supports using either an NSS volume or Linux POSIX volume. Inspect the script to determine which type of file system is used.
DNS Server	No
File, AFP (Apple Filing Protocol)	No
File, CIFS (Windows File Services)	No

OES 2 SP3 Cluster Resources	Uses a Linux POSIX file system on a CSM container in OES 2 SP3
File, FTP	No
File, NetStorage	No
File, Samba	Yes
iFolder 3.9	Yes
iPrint	Yes
MySQL	Yes
QuickFinder (Server Synchronization Feature)	No
Storage, DST shadow volumes built with NSS volumes	No
Storage, Linux POSIX file system on a CSM containers	Yes, based on a Generic_FS template
Storage, NCP volumes	Yes, based on a Generic_FS template
Storage, NSS pools and volumes	No
Xen virtual machines	No

14.3 Deporting the CSM Containers with Segment Managers

If a resource's CSM container does not have a segment manager, it contains a single volume, and the CSMPORT utility automatically converts the volume to a compatibility volume and deports the CSM container. If a resource's CSM container has a segment manager on it, you must manually convert the Linux POSIX volumes in the container to compatibility volumes, and then deport the container by using the EVMS GUI. Special handling is also required to convert the scripts for CSM resources with a segment manager on them.

IMPORTANT: A CSM container must have its volumes converted and be deported in order to work in compatibility mode on OES 11 SP1 nodes.

Having a segment manager on the CSM container is not a typical configuration. In an OES 2 cluster, you might have added a segment manager on a 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*.

Check each of the CSM containers to determine if it contains a segment manager, and if it does, how many volumes it contains. If a resource's CSM container contains a segment manager, convert the container's volumes to compatibility volumes, then deport the container. If a resource's CSM container does not contain a segment manager, continue with [Section 14.4, “Modifying the Scripts for CSM Resources without a Segment Manager,”](#) on page 374.

Before you add OES 11 SP1 nodes to an OES 2 SP3 cluster, perform the following for each CSM cluster resource:

- 1 Log in as the `root` user on the OES 2 SP3 server in the cluster where the Linux POSIX volume cluster resource is online, then open a terminal console.
If the resource is offline, log in to an OES 2 SP3 node in the resource's *Preferred Nodes* list, then bring the resource online by entering


```
cluster online <resource_name>
```

2 Launch the EVMSGUI utility by entering

```
evmsgui
```

3 Check to see if the CSM container has a segment manager on it, then do one of the following:

- ♦ If the CSM container has a segment manager, determine which volumes are on it, then continue with the next step.
- ♦ If a resource's CSM container does not have a segment manager, it has only a single volume and it is not necessary to convert the volume or deport the container. Continue with [Section 14.4, "Modifying the Scripts for CSM Resources without a Segment Manager," on page 374.](#)

4 Convert each volume in the CSM container to a compatibility volume:

- 4a** Select the *Volumes* tab.
- 4b** Unmount the volume.
- 4c** Right-click the volume and select *Convert to Compatibility Volume*.
- 4d** Click *Convert*.
- 4e** Click *Save* to save all of the changes.
- 4f** Repeat [Step 4a](#) through [Step 4e](#) for each volume in the CSM container.

5 Modify the CSM container's properties to set its type to *Deported*:

- 5a** Select the *Containers* tab.
- 5b** Right-click the container and select *Modify Properties*.
- 5c** Select the container, then click *Next*.
- 5d** Select *Type* and change it to *Deported*.
- 5e** Click *Modify*, then click *OK* to confirm the change.
- 5f** Click *Save* to save the container type change.

6 Repeat the previous steps for each CSM resource in the OES 2 SP3 cluster that manages a clustered Linux POSIX volume.

7 Continue with [Section 14.5, "Modifying the Scripts for CSM Resources with a Segment Manager," on page 378.](#)

14.4 Modifying the Scripts for CSM Resources without a Segment Manager

Before you add OES 11 SP1 nodes to an OES 2 SP3 cluster, you must take all of the CSM-based resources offline on the OES 2 SP3 nodes in the cluster. Modify their scripts to run on OES 11 SP1 nodes by adding `csmpport` commands to activate, deactivate, or check the status of the CSM container. After you modify the scripts, the resources cannot be used on OES 2 SP3 nodes. They can be mounted successfully only on OES 11 SP1 nodes in the cluster.

IMPORTANT: If the CSM container for your cluster resource has a segment manager on it, follow the instructions in [Section 14.3, "Deporting the CSM Containers with Segment Managers," on page 373,](#) then go to [Section 14.5, "Modifying the Scripts for CSM Resources with a Segment Manager," on page 378.](#)

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 to prepare the CSM resources for an OES 11 SP1 node in a mixed-mode cluster. Do not add the OES 11 SP1 nodes at this time.

- ♦ [Section 14.4.1, “Offlining the CSM Cluster Resources without Segment Managers,” on page 375](#)
- ♦ [Section 14.4.2, “Configuring Scripts for a CSM Cluster Resource without a Segment Manager,” on page 375](#)
- ♦ [Section 14.4.3, “Sample Load Script for a CSM Resource without a Segment Manager,” on page 377](#)
- ♦ [Section 14.4.4, “Sample Unload Script for CSM Resource without a Segment Manager,” on page 377](#)
- ♦ [Section 14.4.5, “Sample Monitor Script for a CSM Resource without a Segment Manager,” on page 378](#)

14.4.1 Offlining the CSM Cluster Resources without Segment Managers

Offline every OES 2 SP3 cluster resource that manages a Linux POSIX file system on CSM containers without segment managers:

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the cluster you want to manage.
- 3 On the Cluster Manager page, select the check box next to each of the CSM-based cluster resources, then click *Offline*.
- 4 Continue with [Section 14.4.2, “Configuring Scripts for a CSM Cluster Resource without a Segment Manager,” on page 375](#).

14.4.2 Configuring Scripts for a CSM Cluster Resource without a Segment Manager

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the cluster you want to manage.
- 3 Select *Cluster Options*.
- 4 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.
- 5 Click the *Scripts* tab to view the Load Script page.
- 6 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
```

If the path to the mount point does not exist on OES 11 SP1 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
```

For an example, see [Section 14.4.3, “Sample Load Script for a CSM Resource without a Segment Manager,” on page 377](#).

- 7 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 14.4.4, “Sample Unload Script for CSM Resource without a Segment Manager,” on page 377](#).

- 8 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 14.4.5, “Sample Monitor Script for a CSM Resource without a Segment Manager,” on page 378](#).

- 9 Repeat [Step 3](#) to [Step 8](#) for each of the resources that you took offline in [Section 14.4.1, “Offlining the CSM Cluster Resources without Segment Managers,” on page 375](#).

Do not bring the CSM cluster resources online again until OES 11 SP1 nodes have joined the cluster, and each resource’s *Preferred Nodes* list has been modified to use only OES 11 SP1 nodes.

- 10 Continue with [Section 14.6, “Configuring and Adding OES 11 SP1 Nodes to the OES 2 SP3 Cluster,” on page 385](#).

14.4.3 Sample Load Script for a CSM Resource without a Segment Manager

Use the following sample load script to complete the fields for your CSM cluster resource on OES 11 SP1:

```
#!/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

# create the mount point path when loading on a new 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
```

14.4.4 Sample Unload Script for CSM Resource without a Segment Manager

Use the following sample unload script to complete the fields for your CSM cluster resource on OES 11 SP1:

```
#!/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
```


14.4.5 Sample Monitor Script for a CSM Resource without a Segment Manager

Use the following sample monitor script to complete the fields for your CSM cluster resource on OES 11 SP1:

```
#!/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
```

14.5 Modifying the Scripts for CSM Resources with a Segment Manager

Before you add OES 11 SP1 nodes to an OES 2 SP3 cluster, you must take all of the CSM-based resources offline on the OES 2 SP3 nodes in the cluster. Modify their scripts to run on OES 11 SP1 nodes by adding `csmport` commands to activate, deactivate, or check the status of the CSM container. After you modify the scripts, the resources cannot be used on OES 2 SP3 nodes. They can be mounted successfully only on OES 11 SP1 nodes in the cluster.

If the CSM container has a segment manager on it, you must first convert the volumes in the container to compatibility volumes and deport them before you offline them and modify their scripts for OES 11 and later. For information, see [Section 14.3, “Deporting the CSM Containers with Segment Managers,” on page 373](#).

IMPORTANT: If the CSM container does not have a segment manager on it, follow the instructions in [Section 14.4, “Modifying the Scripts for CSM Resources without a Segment Manager,” on page 374](#).

If the CSM container has a segment manager on it, the container can have one or more volumes in it. 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.

In addition to adding the csmport commands to the scripts, you must 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 to prepare the CSM resources for an OES 11 SP1 node in a mixed-mode cluster. Do not add the OES 11 SP1 nodes at this time.

- ♦ [Section 14.5.1, “Offlining the CSM Cluster Resources with a Segment Manager,” on page 380](#)
- ♦ [Section 14.5.2, “Configuring the Scripts for a CSM Cluster Resource with a Segment Manager,” on page 380](#)
- ♦ [Section 14.5.3, “Sample Load Script for a CSM Resource with a Segment Manager,” on page 383](#)
- ♦ [Section 14.5.4, “Sample Unload Script for a CSM Resource with a Segment Manager,” on page 384](#)
- ♦ [Section 14.5.5, “Sample Monitor Script for a CSM Resource with a Segment Manager on It,” on page 385](#)

14.5.1 Offlining the CSM Cluster Resources with a Segment Manager

Offline every OES 2 SP3 cluster resource that manages a Linux POSIX file system on CSM containers with segment managers:

- 1 In iManager, select *Clusters > My Clusters*, select the cluster, then click *Cluster Manager*.
- 2 Select the check box next to each of the CSM-based cluster resources that you want to manage, then click *Offline*.

Wait until the resources report an *Offline* status.

- 3 Continue with [Section 14.5.2, “Configuring the Scripts for a CSM Cluster Resource with a Segment Manager,”](#) on page 380.

14.5.2 Configuring the Scripts for a CSM Cluster Resource with a Segment Manager

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the cluster you want to manage.
- 3 Select *Cluster Options*.
- 4 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.

- 5 Click the *Scripts* tab to view the Load Script page.
- 6 On the Load Script page, modify the script to handle the CSM container and multiple volumes on it.

For an example, see [Section 14.5.3, “Sample Load Script for a CSM Resource with a Segment Manager,”](#) on page 383.

- 6a** 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=xfs
```

- 6b** Add a device definition for each volume:

```
#define the volume devices
VOLUME_DEV1=csm44p1
VOLUME_DEV2=csm44p2
```

- 6c** Add a mount device definition for each volume:

```
#define the devices
MOUNT_DEV1=/dev/mapper/$VOLUME_DEV1
MOUNT_DEV2=/dev/mapper/$VOLUME_DEV2
```

- 6d** Add a mount point definition for each volume:

```
#define the mount points
MOUNT_POINT1=/mnt/lxvolext3
MOUNT_POINT2=/mnt/lxvolxfs
```

- 6e** Add a `csmport` command to activate the CSM container:

```
#activate the container
exit_on_error csmport -i $CONTAINER_NAME
```


- 6f** Add a `kpartx` command to activate the partitions. The command should follow the `csmport` command in the load script.

```
#activate the partitions
exit_on_error /sbin/kpartx -a /dev/mapper/$CONTAINER_NAME
```

- 6g** 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
```

- 6h** 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
```

- 6i** Click *Apply* to save the load script changes.

- 7** 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 14.5.4, “Sample Unload Script for a CSM Resource with a Segment Manager,” on page 384](#).

- 7a** Modify the definitions as described in [Step 6a](#) through [Step 6d](#) 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/lxvolxext3
MOUNT_POINT2=/mnt/lxvolxfs
```

- 7b** Add an `umount` 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
```

- 7c** 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
```

- 7d** Add a `csmport` command to deactivate the CSM container:

```
#deactivate the container
exit_on_error csmport -e $CONTAINER_NAME
```

- 7e** Click *Apply* to save the unload script changes.

- 8** 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 14.5.5, “Sample Monitor Script for a CSM Resource with a Segment Manager on It,”](#) on page 385.

- 8a** Modify the definitions as described in [Step 6a](#) through [Step 6d](#) 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/lxvolxext3
MOUNT_POINT2=/mnt/lxvolxfs
```

- 8b** 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
```

- 8c** Add a `csmport` command to check the status of the CSM container:

```
#check the container
exit_on_error csmport -c $CONTAINER_NAME
```

- 8d** Click *Apply* to save the monitor script changes.

- 9** Repeat [Step 5](#) to [Step 8](#) for each of the resources that you took offline in [Section 14.4.1, “Offlining the CSM Cluster Resources without Segment Managers,”](#) on page 375.

Do not bring the CSM cluster resources online again until OES 11 SP1 nodes have joined the cluster, and each resource’s *Preferred Nodes* list has been modified to use only OES 11 SP1 nodes.

- 10** Continue with [Section 14.6, “Configuring and Adding OES 11 SP1 Nodes to the OES 2 SP3 Cluster,”](#) on page 385.

14.5.3 Sample Load Script for a CSM Resource with a Segment Manager

Use the following sample load script to complete the fields for your CSM cluster resource on OES 11 SP1:

```
#!/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=xfst

#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/lxvolxfst

#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
```


14.5.4 Sample Unload Script for a CSM Resource with a Segment Manager

Use the following sample unload script to complete the fields for your CSM cluster resource on OES 11 SP1:

```
#!/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=xfst

#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/lxvolxfst

#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
```


14.5.5 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 SP1:

```
#!/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=xfss

#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/lxvolxext3
MOUNT_POINT2=/mnt/lxvolxfss

#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
```

14.6 Configuring and Adding OES 11 SP1 Nodes to the OES 2 SP3 Cluster

Configure Novell Cluster Services on an OES 11 SP1 server, join the node to the cluster, then add a line for each of the CSM containers in its `/etc/opt/novell/nss/nlvm.conf` file. Do not bring the CSM resources online at this time.

- 1 Prepare an OES 11 SP1 server for the cluster, but do not allow it to join the cluster. On page 3 of the Cluster Services Configuration in YaST2, ensure that you deselect (disable) the check box for *Start Cluster Services now*.

Use `chkconfig novell-ncs off` before rebooting to the SLES11 SP2 Installation CD. This prevents Cluster Services from starting automatically after reboots.

IMPORTANT: Do not start cluster services on this node or join the OES 2 SP3 cluster, until after all of the Linux POSIX cluster resources have been taken offline.

For information about configuring the OES 11 SP1 node for the existing OES 2 SP3 cluster, see [Step 2 in Section 7.4, “Adding OES 11x Nodes to an OES 2 SP3 Cluster \(Rolling Cluster Upgrade\),” on page 108.](#)

2 Log in as the root user on the OES 11 SP1 server, then open a terminal console.

3 Add the OES 11 SP1 server to the OES 2 SP3 cluster by entering

```
cluster join
```

4 In a text editor, add the following line for each CSM container to the `/etc/opt/novell/nss/nlvm.conf` file on the OES 11 SP1 node, then save the file.

```
Include devices /dev/mapper/<container_name>
```

Replace `container_name` with the name of the CSM container for the Linux POSIX volume.

5 Continue with [Section 14.7, “Configuring the Preferred Nodes and Cluster Settings for a CSM Cluster Resource,” on page 386.](#)

14.7 Configuring the Preferred Nodes and Cluster Settings for a CSM Cluster Resource

After you have modified the CSM cluster resources’ scripts and added at least one OES 11 SP1 node to the cluster, the CSM resources are ready to run only on OES 11 SP1 nodes in the cluster. You must modify the preferred nodes for each cluster resource so that it fails over only to other nodes running OES 11 SP1. You should also disable the *Resource Follows Master* option until the upgrade is complete and all of the OES 2 SP3 nodes have left the cluster. 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 SP3 nodes and converted the load, unload, and monitor scripts as described in the following:
 - ♦ [Section 14.4, “Modifying the Scripts for CSM Resources without a Segment Manager,” on page 374](#)
 - ♦ [Section 14.5, “Modifying the Scripts for CSM Resources with a Segment Manager,” on page 378](#)
- 2 Ensure that you have added one or more OES 11 SP1 nodes to the OES 2 SP3 cluster as described in [Section 14.6, “Configuring and Adding OES 11 SP1 Nodes to the OES 2 SP3 Cluster,” on page 385](#), but do not bring the CSM resources online at this time.
- 3 In iManager, select *Clusters > My Clusters*, 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 SP1 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 10.9, "Configuring Preferred Nodes and Failover Order for a Resource,"](#) on page 198, then click *Finished*.

IMPORTANT: Ensure that the preferred nodes list contains only OES 11 SP1 nodes.

- 10 Configure resource priorities for load order as described in [Section 10.10, "Configuring Resource Priorities for Load Order,"](#) on page 200.
- 11 (Optional) Assign the resource to a Resource Mutual Exclusion Group as described in [Section 10.11, "Configuring Resource Mutual Exclusion Groups,"](#) on page 201.
- 12 Continue with [Section 14.8, "Verifying the Load and Unload Scripts,"](#) on page 387.

14.8 Verifying the Load and Unload Scripts

After you have modified the resource scripts, ensure that the load and unload scripts are working.

- 1 In iManager, select *Clusters > My Clusters*, then select the cluster.
- 2 On the Cluster Manager page, select the check box next to the new CSM cluster resource, then click *Online*.

Ensure that the resource successfully enters the *Running* state. If the resource goes comatose instead, it is probably because you made an error when typing values in the script definitions. 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 14.9, "Creating an NCP Virtual Server Object for a CSM Cluster Resource,"](#) on page 387.

14.9 Creating an NCP Virtual Server Object for a CSM Cluster Resource

Having an NCP 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. When you cluster-enable a Linux POSIX volume on OES 2 SP3, an NCP virtual server object is not created automatically as it is when you cluster-enable an NSS pool.

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 SP1: NCP Server for Linux Administration Guide*.

You can use an existing Virtual Server Object that you created before you moved the CSM resource to OES 11 (or later). After a conversion is finalized, 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.

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 14.9.1, “Creating the Virtual Server Object,” on page 388](#)
- ♦ [Section 14.9.2, “Modifying the Load Script,” on page 389](#)
- ♦ [Section 14.9.3, “Modifying the Unload Script,” on page 389](#)
- ♦ [Section 14.9.4, “Activating the Script Changes,” on page 390](#)
- ♦ [Section 14.9.5, “Verifying the Virtual Server,” on page 390](#)

14.9.1 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.7, “Virtual NCS NCP Server Object \(`ncs_ncpserv.py`\) Script,” on page 433](#).

- 1 On the master cluster node, open a terminal console, then log in as the root user.
- 2 In the console, use the `cd` command to go to the `/opt/novell/ncs/bin` directory.
- 3 At the command prompt, enter

```
./ncs_ncpserv.py -c 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 14.9.2, “Modifying the Load Script,” on page 389](#)

14.9.2 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 > My Clusters*, then select the cluster.
- 2 On the Cluster Manager page, 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 14.9.3, “Modifying the Unload Script,”](#) on page 389.

14.9.3 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 > My Clusters*, then select the cluster.
- 2 On the Cluster Manager page, 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 14.9.4, “Activating the Script Changes,”](#) on page 390.

14.9.4 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 > My Clusters*, then select the cluster.
- 2 On the Cluster Manager page, 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 14.9.5, “Verifying the Virtual Server,” on page 390](#).

14.9.5 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`.

14.10 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.

15 Configuring Novell Cluster Services in a 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 SP1 server. It can be used in the host and guest virtual environments.

For a virtual machine, you can use any virtualization hypervisor (such as Xen, KVM, VMware, and Hyper-V) that supports the 64-bit SUSE Linux Enterprise Server 11 SP2 or later operating system, which is the platform used for OES 11 SP1 services. For information about using OES 11 SP1 with Xen and KVM, see the following resources:

- ♦ For information on setting up virtualized OES 11 SP1, see “[Installing, Upgrading, or Updating OES on a VM](#)” in the *OES 11 SP1: Installation Guide*.
- ♦ To get started with Xen virtualization, see the [Virtualization with Xen documentation \(http://www.suse.com/documentation/sles11/book_xen/data/book_xen.html\)](http://www.suse.com/documentation/sles11/book_xen/data/book_xen.html).
- ♦ To get started with KVM virtualization, see the [Virtualization with KVM documentation \(http://www.suse.com/documentation/sles11/book_kvm/data/book_kvm.html\)](http://www.suse.com/documentation/sles11/book_kvm/data/book_kvm.html).

You can install Novell Cluster Services on an OES 11 SP1 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 SP1 Xen host server environment. You can modify these templates to work with other types of hypervisors.

All templates except Xen and XenLive are valid on guest servers (DomU). When Novell Cluster Services is installed on an OES 11 SP1 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.

Use the information in the following sections to help you deploy and run Novell Cluster Services in a virtual environment: Although many different cluster virtualization scenarios are possible, only those outlined in the following sections have been tested.

- ♦ [Section 15.1, “Prerequisites for Xen Host Server Environments,” on page 392](#)
- ♦ [Section 15.2, “Virtual Machines as Cluster Resources,” on page 392](#)
- ♦ [Section 15.3, “Virtual Machines as Cluster Nodes,” on page 400](#)
- ♦ [Section 15.4, “Virtual Cluster Nodes in Separate Clusters,” on page 401](#)
- ♦ [Section 15.5, “Mixed Physical and Virtual Node Clusters,” on page 401](#)
- ♦ [Section 15.6, “Additional Information,” on page 403](#)

15.1 Prerequisites for Xen Host Server Environments

When using Novell Cluster Services to create cluster resources in a Xen host server environment, one network adapter on each cluster node must be available without Xen bridging to allow Novell Cluster Services to monitor for network state changes in the cluster. You need at least one other network adapter on each cluster node to use as a bridge for the virtual servers that will be hosted on the Xen host.

For each cluster node, determine which network adapter you want to use for Novell Cluster Services communications and network state monitoring, then do the following:

- ♦ Disable bridging on the network adapter.
- ♦ Assign the IP address of the host server to the network adapter.

For information about why this configuration is necessary, see “[TID 7004595: Novell Cluster Services is unable to detect a network link down when installed on a Xen domain 0 \(dom0\)](http://www.novell.com/support/viewContent.do?externalId=7004595&sliceId=1)” (<http://www.novell.com/support/viewContent.do?externalId=7004595&sliceId=1>) on Novell Support (<http://www.novell.com/support>).

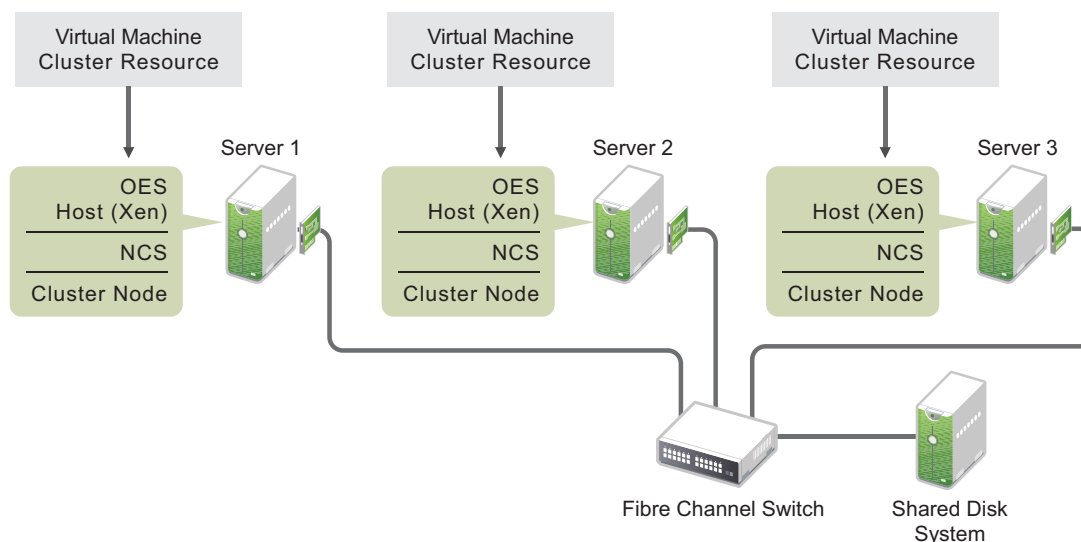
15.2 Virtual Machines as Cluster Resources

In this scenario, you have an OES 11 SP1 cluster configured on physical machines. OES 11 SP1 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 SP1 cluster without virtualization.

You can create OES virtual machines on each cluster node and configure those virtual machines to be cluster resources. You can then fail over or migrate virtual machine cluster resources (entire virtual machines) to different physical nodes in your cluster.

Figure 15-1 depicts how this setup might look. Novell Cluster Services (NCS) is installed and running on the virtual machine (VM) host server.

Figure 15-1 Virtual Machines as Cluster Resources



The following sections describe how to create a cluster resource and its cluster scripts for each virtual machine:

- [Section 15.2.1, “Creating a Xen Virtual Machine Cluster Resource,” on page 393](#)
- [Section 15.2.2, “Configuring Virtual Machine Load, Unload, and Monitor Scripts,” on page 394](#)
- [Section 15.2.3, “Setting Up Live Migration,” on page 399](#)

15.2.1 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 > My Clusters*, select the cluster, then click *Cluster Options*.
- 4 On the Cluster Options page, click *New*.
- 5 Click the *Resource* radio button to specify *Resource* as the resource type you want to create, then click *Next*.
- 6 Specify a name for the virtual machine resource.
- 7 In the *Inherit From Template* field, browse to the Cluster object container, then select the desired Xen template name from the list of templates in the container.
The Xen templates are named `Xen_Template` and `XenLive_Template`.
- 8 Select the *Define Additional Properties* check box, click *Next*, then continue with [“Configuring Virtual Machine Load, Unload, and Monitor Scripts” on page 394](#).

15.2.2 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 10.8, “Configuring the Start, Failover, and Failback Modes for Cluster Resources,” on page 197](#) and [Section 10.9, “Configuring Preferred Nodes and Failover Order for a Resource,” on page 198](#).

If you are modifying the scripts, continue with the following sections:

- ♦ [“Configuring the Load Script” on page 394](#)
- ♦ [“Configuring the Unload Script” on page 395](#)
- ♦ [“Configuring the Monitor Script” on page 397](#)

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 394](#)
- ♦ [“Sample XenLive_Template Load Script” on page 395](#)

- 2 Click *Next* and continue with [“Configuring the Unload Script” on page 395](#).

Sample Xen_Template Load Script

The Xen_Template load 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

#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 396](#)
- ♦ [“Sample XenLive_Template Unload Script” on page 396](#)

- 2 Click *Next*, then continue the setup by configuring the resource policies and the resource server assignments.

For information, see [Section 10.8, “Configuring the Start, Failover, and Failback Modes for Cluster Resources,”](#) on page 197 and [Section 10.9, “Configuring Preferred Nodes and Failover Order for a Resource,”](#) on page 198.

- 3 If you want to enable monitoring for the resource, continue with [“Configuring the Monitor Script”](#) on page 397.

Sample Xen_Template Unload Script

The Xen_Template unload script appears similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
export OCF_ROOT=/usr/lib/ocf/

#define domU name
OCF_RESOURCE_INSTANCE=xen_vm_name

#define shared volume name
VOLUME_NAME=${OCF_RESOURCE_INSTANCE}

#define the volume group name
CONTAINER_NAME=name
# define the device
MOUNT_DEV=/dev/${CONTAINER_NAME}/${VOLUME_NAME}

# filesystem settings
export OCF_RESKEY_device=$MOUNT_DEV
export OCF_RESKEY_directory=/mnt/${OCF_RESOURCE_INSTANCE}
export OCF_RESKEY_fstype=ext3
#export OCF_RESKEY_options=

# service settings
export OCF_RESKEY_xmfile=${OCF_RESKEY_directory}/${OCF_RESOURCE_INSTANCE}

# stop the service
exit_on_error ocf_stop Xen

# umount the file system
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error ocf_stop Filesystem

#deactivate the volume group
exit_on_error vgchange -a n $VOLGROUP_NAME

# return status
exit 0
```

Sample XenLive_Template Unload Script

The XenLive_Template unload script appears similar to the following example:


```
#!/bin/bash
. /opt/novell/ncs/lib/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 monitor script that you can customize. You use the script to monitor the health of a virtual machine cluster resource.

Resource monitoring is disabled by default. If you want to enable resource monitoring for a virtual machine cluster resource, you must enable it prior to customizing the resource monitor script.

- ♦ [“Enabling Resource Monitoring” on page 397](#)
- ♦ [“Viewing or Modifying the Monitor Script” on page 398](#)
- ♦ [“Sample Xen_Template Monitor Script” on page 398](#)
- ♦ [“Sample XenLive_Template Monitor Script” on page 399](#)

Enabling Resource Monitoring

To enable resource monitoring for a virtual machine cluster resource:

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the cluster.
- 3 On the Cluster Manager page, select the check box next to the virtual machine resource, then click the *Details* link.
- 4 Click the *Monitoring* tab, then select the *Enable Resource Monitoring* check box to enable resource monitoring for the resource.
Resource monitoring is disabled by default.
- 5 For the polling interval, specify how often you want the resource monitor script for this resource to run.
You can choose to specify the number in minutes or seconds.
- 6 Specify the number of failures (*Maximum Local Failures*) for the specified amount of time (*Time Interval*).

If the resource monitor detects that the resource fails the number of times specified in the amount of time specified, a failure action initiates.

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

If the failure action initiates and you chose the option to migrate the resource to another server, the resource migrates to the next server in its *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 > My Clusters*.
- 2 Select the cluster.
- 3 On the Cluster Manager page, select the check box next to the virtual machine resource that you created, then click the *Details* link.
- 4 Click the *Scripts* tab, then click the *Monitor Script* link.
- 5 View or edit the commands in the script that monitor the resource on the server.

You can use the same commands that would be used at the Linux terminal console.

See the following examples of the default *Xen_Template* and *XenLive_Template* monitor scripts:

- ♦ ["Sample Xen_Template Monitor Script" on page 398](#)
- ♦ ["Sample XenLive_Template Monitor Script" on page 399](#)

- 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 397](#). 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
```

15.2.3 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>).

15.3 Virtual Machines as Cluster Nodes


In this scenario, you have OES 11 SP1 (Xen) installed and configured on each node (physical machine). You create the OES 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.

You can then create and configure cluster resources that run on the virtual cluster nodes. The process for creating and configuring cluster resources on a virtual cluster node is the same as on a physical cluster node.

Cluster resources can be failed over or migrated between virtual cluster nodes that are on the same physical node or on separate physical nodes.

Figure 15-2 depicts using virtual machines as cluster nodes.

Figure 15-2 Virtual Machines as Cluster Nodes



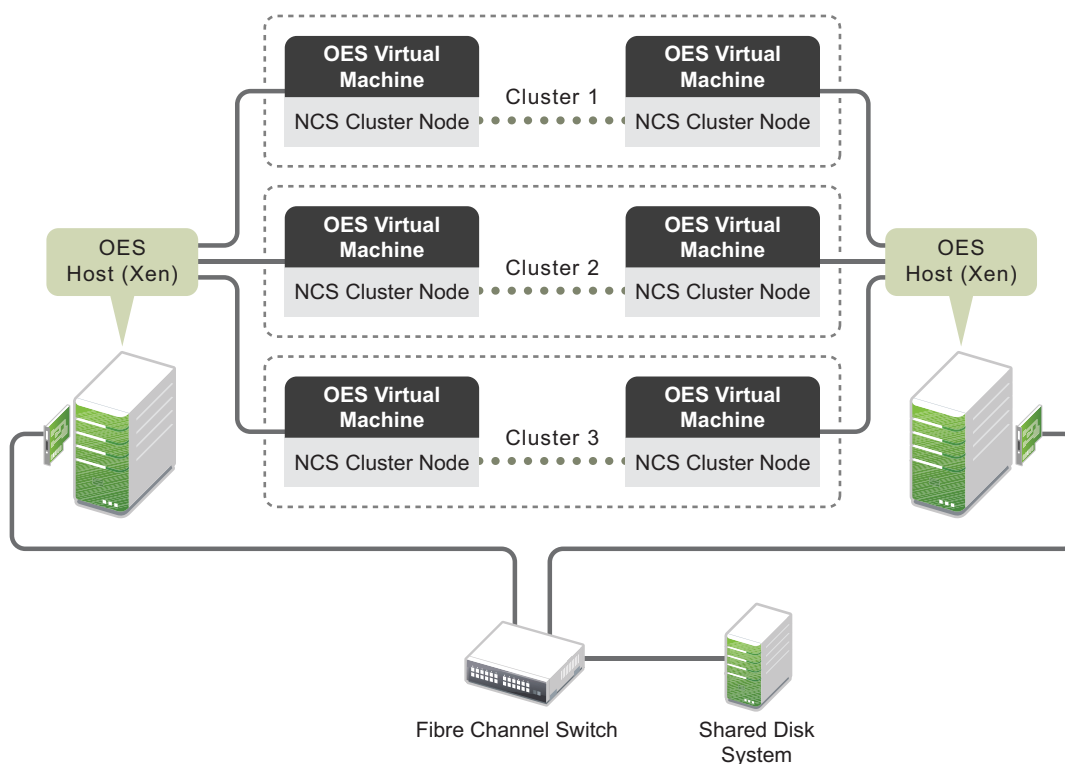
15.4 Virtual Cluster Nodes in Separate Clusters

In this scenario, you have OES 11 SP1 (Xen) installed and configured on each node (physical machine). You create multiple OES 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.

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 15-3 depicts using virtual cluster nodes in separate clusters.

Figure 15-3 Virtual Cluster Nodes in Separate Clusters



15.5 Mixed Physical and Virtual Node Clusters

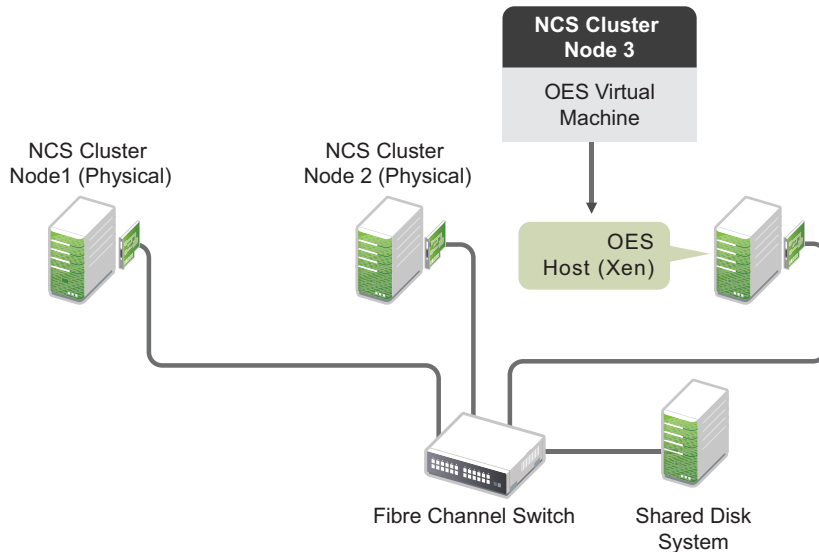
This is a temporary scenario that is used for upgrading cluster nodes or converting clusters from physical to virtual cluster nodes.

This can be done through several different methods. One method is to add a virtual OES cluster node to an existing physical OES cluster. To do this, you install an OES 11 SP1 (Xen) server (physical machine). You then create a OES virtual machine on the physical machine and install and configure Novell Cluster Services on the OES virtual machine. During the Novell Cluster Services installation, you add the OES virtual cluster node to your existing OES cluster (physical nodes).

You can then migrate the desired resources from the physical cluster nodes to the virtual cluster node. This lets you offload resources from physical nodes so you can upgrade hardware and software and then replace the physical cluster nodes with virtual cluster nodes.

Figure 15-4 depicts how this setup might look.

Figure 15-4 *Mixed Physical and Virtual Node Cluster*

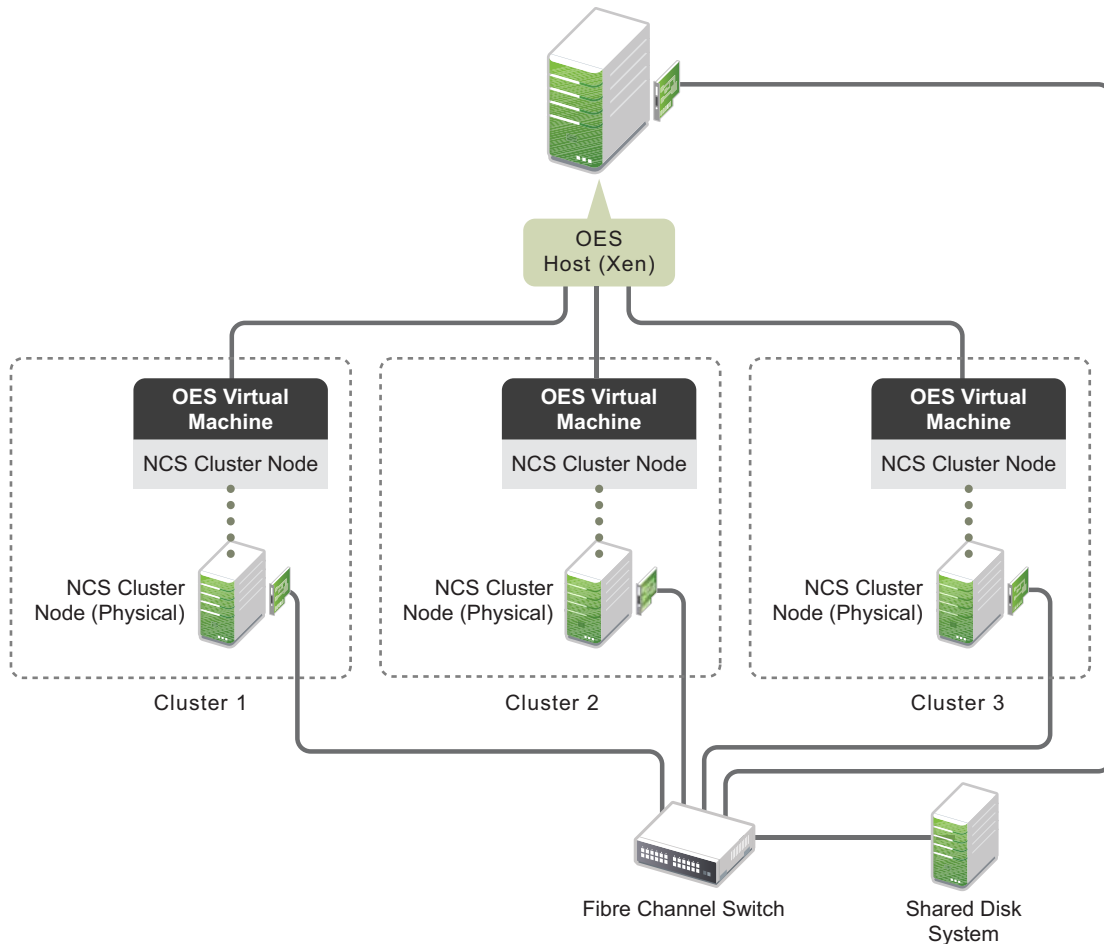


Another method is to install Novell Cluster Services on physical nodes and create a separate cluster for each node. You then install an OES 11 SP1 (Xen) server (physical machine) and create OES virtual machines and install Novell Cluster Services on each virtual machine. You can then add one virtual OES cluster node to each cluster to create multiple two-node clusters, each containing one physical and one virtual cluster node.

This allows you to migrate the desired resources from each physical cluster node to the virtual cluster node in the same cluster. Using this setup, you offload resources from physical nodes so you can upgrade hardware and software and then replace the physical cluster nodes in each cluster with virtual cluster nodes.

Figure 15-5 depicts how this setup might look.

Figure 15-5 *Separate Mixed Physical and Virtual Node Clusters*



15.6 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 SP1: Installation Guide* guide.

For information on setting up virtualized OES 11 SP1, see “[Installing, Upgrading, or Updating OES on a VM](#)” in the *OES 11 SP1: Installation Guide* guide.

16 Troubleshooting Novell Cluster Services

This section describes known issues and frequently asked questions for managing Novell Cluster Services.

- ♦ [Section 16.1, "File Location Information," on page 406](#)
- ♦ [Section 16.2, "Version Issues," on page 406](#)
- ♦ [Section 16.3, "Diagnosing a Poison Pill Condition," on page 406](#)
- ♦ [Section 16.4, "Diagnosing Cluster Problems," on page 407](#)
- ♦ [Section 16.5, "Why didn't my resource go to the node I specified?," on page 407](#)
- ♦ [Section 16.6, "iManager: Error 500 Occurs while Connecting to a Cluster," on page 408](#)
- ♦ [Section 16.7, "iManager: Error 500 Occurs after Installing or Updating the Clusters Plug-In for iManager 2.7.5," on page 408](#)
- ♦ [Section 16.8, "iManager: Problems with Clusters Plug-In After Update for Novell iManager 2.7.5," on page 408](#)
- ♦ [Section 16.9, "A Device Name Is Required to Create a Cluster Partition," on page 409](#)
- ♦ [Section 16.10, "Cluster Resource Goes Comatose Immediately After Migration or Failover," on page 409](#)
- ♦ [Section 16.11, "Cannot Authenticate to Remote Servers during Cluster Configuration," on page 409](#)
- ♦ [Section 16.12, "Cannot Connect to an iSCSI Target," on page 409](#)
- ♦ [Section 16.13, "Errors when Deleting a Cluster Resource or Clustered Pool," on page 409](#)
- ♦ [Section 16.14, "Is there a way to uninstall Novell Cluster Services from a server?," on page 410](#)
- ♦ [Section 16.15, "Running supportconfig Concurrently on Multiple Nodes Causes Cluster Problems," on page 410](#)
- ♦ [Section 16.16, "Error Bad XML: Cluster Search Times Out," on page 410](#)
- ♦ [Section 16.17, "Error NFS Stale Handle Reported on Host when Cluster Failover Occurs," on page 410](#)
- ♦ [Section 16.18, "Error 20897 - This node is not a cluster member," on page 411](#)
- ♦ [Section 16.19, "Error 499 - Could not delete this resource Data_Server," on page 411](#)
- ♦ [Section 16.20, "Error 603 - smdr.novell is not registered with SLP for a new cluster resource," on page 411](#)
- ♦ [Section 16.21, "Migration Authentication Error Occurs when Connecting to a New Cluster Resource on the Target Server," on page 412](#)
- ♦ [Section 16.22, "Communications Errors," on page 412](#)

16.1 File Location Information

Information about Novell Cluster Services scripts and configuration can be found in the following locations. Ensure that you use the Clusters plug-in to Novell iManager to configure cluster and resource settings and to edit the load, unload, and monitor scripts.

Information	File Location
Load scripts	/var/opt/novell/ncs
Unload scripts	
Monitor 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 ldnscs and unldnscs	/opt/novell/ncs/bin
Resource templates	/opt/novell/ncs/templates

16.2 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 437.

16.3 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

16.4 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 *Action > Run Report* option on the *Clusters > My Clusters* page of the Clusters plug-in to iManager.
- ♦ SAN configuration:
 - ♦ Host bus adapter, hub or switch type
 - ♦ Device driver name and revision
 - ♦ Storage controller type
 - ♦ Storage controller firmware revision
- ♦ SBD configuration - Single or mirrored partition?
- ♦ LAN configuration:
 - ♦ Network interface card, hub, or switch type
 - ♦ Device driver name and revision
 - ♦ Dedicated heartbeat or shared public LAN?
- ♦ Server configuration:
 - ♦ Type, memory, and number of CPUs
 - ♦ Software revisions and patches
 - ♦ List of RPMs
 - ♦ `/opt/novell/ncs/bin/ldncs` file
- ♦ LAN packet trace
- ♦ Console log files
- ♦ Server failure log files
- ♦ Server core dump 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?

16.5 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.

16.6 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/).

16.7 iManager: Error 500 Occurs after Installing or Updating the Clusters Plug-In for iManager 2.7.5

After installing or updating the Clusters plug-in to version 3.3.1 for Novell iManager 2.7.5, you get a Status Code 500 error whenever you attempt to manage the cluster from iManager.

This problem can occur if iManager loads a version of the class loader file earlier than `/var/opt/novell/iManager/nps/WEB-INF/lib/commons-lang-2.6.jar`. Novell iManager uses one class loader for all of its plug-ins. If multiple versions of the class loader file exist in the `/var/opt/novell/iManager/nps/WEB-INF/lib` directory, iManager loads whichever one comes first alphabetically.

To resolve this problem, you must remove the earlier versions (such as `commons-lang.jar` or `commons-lang-2.4.jar`) to ensure that iManager loads the 2.6 or later version of the class loader file, then restart Tomcat.

For more information, see [Section 5.8, "Installing or Updating the Clusters Plug-in for iManager," on page 96](#).

This problem has been resolved in the Cluster plug-in for iManager 2.7.6 and later.

16.8 iManager: Problems with Clusters Plug-In After Update for Novell iManager 2.7.5

The Clusters plug-in for Novell iManager 2.7.5 in OES 11 SP1 was reorganized to display two tasks in the left panel: *My Clusters* and *My Resources*. If you have problems displaying or using the plug-in after updating to this version, ensure that you performed the following tasks after the update:

- ♦ [Step 5 in Section 5.8.3, "Uninstalling and Reinstalling the Clusters Plug-In," on page 97](#)
- ♦ [Section 5.8.4, "Updating Role-Based Services for the Clusters Plug-In for OES 11 SP1," on page 98](#)

16.9 A Device Name Is Required to Create a Cluster Partition

If you are planning to work with shared-disk NSS pools and volumes, you must create a Split Brain Detector (SBD) partition when you configure the first node on the cluster by selecting the name of a device (such as `sdC`) that is already initialized and marked shareable for clustering. If you don't enter a device name, you must manually create an SBD before you configure clustering on a second node and before you attempt to create storage objects or cluster-enable NSS pools.

16.10 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 58](#).

16.11 Cannot Authenticate 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.

16.12 Cannot Connect to an iSCSI Target

If you are connecting to an iSCSI target that already has NSS partitions and pools created on it, you might not be able to access those NSS partitions and pools until you reboot the Linux initiator server. This is required for each Linux initiator server that will access the iSCSI target.

For instructions on configuring an OES 11 or later server as an iSCSI initiator and connecting to an iSCSI target, go to [“Mass Storage over IP Networks--iSCSI”](http://www.suse.com/documentation/sles11/stor_admin/data/cha_inst_system_iscsi.html) (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)

16.13 Errors when Deleting a Cluster Resource or Clustered Pool

If you attempt to delete a cluster resource or clustered pool without first taking the cluster resource offline, deletion errors occur, and the data associated with the clustered pool is not recoverable.

To avoid this problem, you must take the cluster resource offline before you attempt to delete either the cluster resource or the clustered pool. For instructions, see [Section 10.15, “Deleting Cluster Resources, or Disabling Clustering for a Pool, LVM Volume Group, or Service,” on page 209](#).

16.14 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 9.15, “Deleting a Cluster Node from a Cluster,” on page 157](#).

16.15 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.

16.16 Error Bad XML: Cluster Search Times Out

If you are using in Novell eDirectory 8.7.3x, timeouts 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 SP7 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.

16.17 Error NFS Stale Handle Reported on Host when Cluster Failover Occurs

The host I/O reports an `NFS Stale Handle` error immediately after failover is performed on the cluster.

This problem can occur if a user attempts to access the IP address of the resource before the volume is mounted or after the volume is dismounted. You can modify the order of commands in a load script to: `NSS, NFS, IP, NCP, [AFP], and [CIFS | Samba]`. Reverse the order in an unload script. This allows

storage to be mounted before the IP address is bound and to be dismounted after the IP address is unbound. Adjust your scripts to this order, then take the resource offline and bring it online to apply the changes. For sample scripts, see [Section 12.10, “Adding NFS Export for a Clustered Pool Resource,” on page 247](#).

16.18 Error 20897 - This node is not a cluster member

If Novell Cluster Services is installed on a node, but an SBD does not exist, NSS tools (including Novell Linux Volume Manager (NLVM) commands, NSSMU, NSS console commands, and the Storage plug-in to Novell iManager) return the following error:

```
Error 20897 - This node is not a cluster member.
```

In a Novell Cluster Services cluster, NSS tools use the cluster's SBD to detect if a node is a cluster member and to lock against concurrent changes to physically shared storage. Without an SBD, NLVM cannot detect whether a node is a member of the cluster and cannot acquire the locks it needs to execute tasks. In this state, you can use the `-s` option with NLVM commands to prepare a device and create an SBD partition. To minimize the risk of corruption, you must ensure that nobody else is changing any storage on any nodes at the same time.

For information about creating an SBD partition, see [Section 9.16, “Creating or Deleting Cluster SBD Partitions,” on page 157](#).

16.19 Error 499 - Could not delete this resource Data_Server

When you try to delete a cluster resource, you get a message “Could Not Delete This Resource Data_Server (Error 499). The error means that the object does not exist in eDirectory. This error can occur if you try to delete an orphaned cluster resource. An orphaned cluster resource does not have a pool or volume associated with it, does not have an IP address, and the scripts are empty.

To delete an orphaned cluster resource:

- 1 On the master node of the cluster, open a terminal as the `root` user. At the command prompt, enter the following commands:

```
opt/novell/ncs/bin/ncs-configd.py -init  
echo -n "exec rm -f /var/opt/novell/ncs/data_server.*" > /proc/ncs/cluster
```

16.20 Error 603 - smdr.novell is not registered with SLP for a new cluster resource

You might get an error after you create a cluster resource, indicating that `smdr.novell` is not registered with SLP for cluster resources, but the `smdr.novell` service for the node is registered.

```
Error: "cluster--<212>: Read ResVol error -603"
```

The first time a cluster resource is created, `smdrd` cannot figure it out. The `smdrd` service is designed to work with resources that contain a data volume.

Ensure that the resource contains a volume and is configured properly, then restart `smdrd`. Thereafter, `smdrd` is aware of the cluster resource, and advertises it correctly.

- 1 Log in to the server as the `root` user, open a terminal console, then enter

```
rcnovell-smdrd restart
```

16.21 Migration Authentication Error Occurs when Connecting to a New Cluster Resource on the Target Server

You might get a Migration Authentication error when you attempt to connect to a newly created cluster resource on the target server.

```
2013-05-03 12:34:20,095 INFO - Migration Framework:Authentication:nbackup:
Unable to retrieve the Target Service Name list from <resource_IP_address>, trying
with DNS name <resource_DNS_name>
2013-05-03 12:34:20,096 INFO - Migration Framework:Authentication:
2013-05-03 12:34:20,096 INFO - Migration Framework:Authentication:nbackup:
Unable to retrieve the Target Service Name list from <resource_DNS_name>
```

There are two possible causes:

- ♦ The Storage Management Service (SMS) `smdr.novell` service running on the target server is not aware of the newly created resource.

The Migration tool uses SMS to transfer data between nodes. When a cluster resource is first created, the `smdrd` service is not automatically aware of the new resource. For information, see [Section 16.20, “Error 603 - `smdr.novell` is not registered with SLP for a new cluster resource,” on page 411](#).

Restart `smdrd`. Thereafter, `smdrd` is aware of the cluster resource, and advertises it correctly.

1. Log in to the target server as the `root` user, open a terminal console, then enter

```
rcnovell-smdrd restart
```

- ♦ The cluster resource does not contain a volume.

A properly created pool cluster resource or LVM cluster resource contains at least one volume. Even if you have restarted the `smdrd` service, the service by design does not recognize a resource unless it contains a data volume.

Ensure that the resource contains a volume and is configured properly, then restart `smdrd`. Thereafter, `smdrd` is aware of the cluster resource, and advertises it correctly.

1. Create a volume on the pool or LVM volume group, update the resource scripts as needed, then take the resource offline and bring it online to apply the changes.
2. Log in to the target server as the `root` user, open a terminal console, then enter

```
rcnovell-smdrd restart
```

16.22 Communications Errors

When you manage a cluster or cluster resources with the Clusters plug-in to Novell iManager, CIMOM errors can occur for the following reasons.

- ♦ The CIMOM broker daemon must be running and listening on port 5989. For information, see [the Section 4.6.8, “SFCB and CIMOM,” on page 61](#).

- ♦ You must be able to connect to the cluster master node. Verify the ability to connect to the cluster master node by pinging the cluster IP address from a terminal console on a different computer.
Modify the master IP address only by using the Clusters plug-in for iManager as described in [Section 8.10, “Viewing or Modifying the Cluster IP Address or Port,” on page 133](#). A cluster restart is required before Novell iManager can continue to manage the cluster.
- ♦ The cluster IP address must be the same in the *NCS:Network Address* attribute and *Network Address* attribute for the Cluster object. Addresses might not match if you modified the master IP address for the cluster and did not restart the cluster.

17 Security Considerations

This section describes security issues and recommendations for Novell Cluster Services for Novell Open Enterprise Server (OES) 11 SP1. 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 17.1, “Cluster Administration Rights,” on page 415](#)
- ♦ [Section 17.2, “Ports,” on page 415](#)
- ♦ [Section 17.3, “Email Alerts,” on page 416](#)
- ♦ [Section 17.4, “Log Files,” on page 416](#)
- ♦ [Section 17.5, “Configuration Files,” on page 416](#)

17.1 Cluster Administration Rights

For information about the rights needed to install and manage Novell Cluster Services, see the following sections:

- ♦ [“Novell eDirectory 8.8.x” on page 54](#)
- ♦ [“Novell Domain Services for Windows” on page 66](#)
- ♦ [“SFCB and CIMOM” on page 61](#)
- ♦ [“SLP” on page 58](#)
- ♦ [“eDirectory Tree” on page 55](#)
- ♦ [“Cluster Installation Administrator” on page 49](#)
- ♦ [“Cluster Administrator or Administrator-Equivalent User” on page 52](#)

17.2 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 8.10, “Viewing or Modifying the Cluster IP Address or Port,” on page 133](#).

If you enable Email Notification, Cluster Services uses Postfix to send messages. The default port number is 25. For information about how to address port conflicts, see [Section 8.5.1, “Configuring an Email Port,” on page 124](#).

You must specify the Novell 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 Novell iManager. Port 5989 is the default setting for secure HTTP (HTTPS) communications.

17.3 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 8.5, “Configuring Cluster Event Email Notification,” on page 124](#).

17.4 Log Files

NCS normally writes log messages to `/var/log/messages`.

NCS installation writes log messages to `/var/opt/novell/install/ncslog`.

NCS configuration writes log messages to `/var/log/YaST2/y2log`.

You can view cluster events logged in the `/admin/Novell/Cluster/EventLog.xml` file by using the Clusters plug-in for Novell iManager. Go to *Clusters > My Clusters*, select the cluster, then click *Event Log*. For information, see [Section 9.5, “Viewing the Cluster Event Log in iManager,” on page 146](#).

You can also copy the `/admin/Novell/Cluster/EventLog.xml` file to a working location, and then use the `cat` command to display the file. For information, see [Section 9.6, “Viewing the Cluster Event Log at the Command Line,” on page 149](#).

NCS writes cluster resource runtime messages to files under `/var/opt/novell/log/ncs/`. Output log files are named with the resource name, the type of script (load, unload, or monitor), and the `.out` extension (`<resource_name>.<load|unload|monitor>.out`). For example: `/var/opt/novell/log/ncs/myresource.unload.out` reports the log for the myresource unload script messages.

17.5 Configuration Files

You should never modify configuration files directly, unless you are directed to do so by a documented procedure or Novell Support.

The cluster configuration file is `/etc/opt/novell/ncs/clstrlib.conf`.

The cluster configuration for a specific node is `/etc/opt/novell/ncs/<nodename>`.

The cluster node numbers are found in the `/var/opt/novell/ncs/gipc.conf` file.

The Novell Cluster Services reboot behavior after shutting down a node conforms to the kernel panic setting for the Linux operating system. By default the kernel panic setting is set for no reboot after a node shutdown. You can use directive `kernel.panic` in file `/etc/sysctl.conf` file to allow an automatic reboot and to specify the number of seconds to delay the reboot. For information, see [Section 8.7, “Configuring the Cluster Node Reboot Behavior,” on page 129](#).

If you use Device Mapper Multipath for devices, ensure that you modify the settings in the `/etc/modprobe.conf.local` file and the `/etc/multipath.conf` file so that multipath I/O works correctly for Novell Cluster Services cluster resource fail over and fail back. For information, see [Section 4.10, “Multipath I/O Configuration Requirements,” on page 70](#).

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 417](#)
- ♦ [Section A.2, “Business Continuity Clustering Commands,” on page 423](#)
- ♦ [Section A.3, “extend_schema Command,” on page 423](#)
- ♦ [Section A.4, “SBD Utility,” on page 424](#)
- ♦ [Section A.5, “Cluster Segment Manager Import/Export \(csmport\) Utility,” on page 427](#)
- ♦ [Section A.6, “Configuration Update \(ncs-configd.py\) Script,” on page 433](#)
- ♦ [Section A.7, “Virtual NCS NCP Server Object \(ncs_ncpserv.py\) Script,” on page 433](#)
- ♦ [Section A.8, “NCPCON Mount Command for NCP and NSS Volumes,” on page 435](#)

A.1 Cluster Management Commands

To execute a cluster console command, enter `cluster` followed by the command. For example, if you want to display cluster statistics, enter `cluster stats display` at the terminal console. You can also enter `cluster help` at the command prompt to get information on the commands and their functions.

The functions of many of the commands can also be performed using iManager. See the other sections of this document for additional information.

[Table A-1](#) lists the cluster-related terminal console commands and gives a brief description of each command.

Table A-1 Cluster Console Commands

Cluster Console Command	Description
ALERT {resource}{YES NO}	<p>The resource start, failover, or failback mode is set to manual and the resource is waiting for administrator intervention to start on a node, or to fail over or fail back to another node. Alert types are Start Alert, Failback Alert, and Failover Alert.</p> <p>For the Start Alert, you must address the alert, then offline and online the resource.</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>
CONVERT preview [resource] 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>

Cluster Console Command	Description
JOIN	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.
LEAVE	Removes the node where the command is executed from the cluster. The node will not be visible to other servers in the cluster.
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>

Cluster Console Command	Description
<pre>MIGRATE {resource} { <node_name> -list -most -next }</pre>	<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>Options</p> <p><node_name></p> <p>Migrates the resource to the specified node if possible. No action is taken if the specified node is not in the resource's preferred nodes list, there is a Resource Mutual Exclusion conflict on the specified node, or the specified node is not currently active in the cluster.</p> <p>-l, -list</p> <p>Shows the preferred nodes list for the specified resource.</p> <p>-m, -most</p> <p>Migrates the specified resource to the most preferred node currently in the cluster. If the resource is running on such a node, no action is taken.</p> <p>-n, -next</p> <p>Migrates the resource to the node in the preferred node list that is next in order to the node where the resource is currently running. If there is no next node or all such nodes are not in the cluster currently, it searches from the beginning of the resource's preferred nodes list. No action is taken if a new destination could not be determined after the search is exhausted.</p> <p>Examples</p> <p>Move <code>res1</code> from the current node to <code>node2</code> in the resource's preferred nodes list. The resource is migrated only if there are no Resource Mutual Exclusion conflicts on the target node.</p> <pre>cluster migrate res1 node2</pre> <p>List the preferred nodes for resource <code>POOL_14_SERVER</code>.</p> <pre>cluster migrate POOL_14_SERVER -l</pre> <pre>Preferred nodes for resource 'POOL_14_SERVER' CG-06 CG-05 CG-04 CG-03 CG-02 CG-01 CG-08 CG-07</pre> <pre>Status for Resource: POOL_14_SERVER Running on CG-05 Lives: 65 Revision: 5</pre>

Cluster Console Command	Description
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>
RESOURCES	<p>Lists all resources that currently exist in the cluster. The resources do not need to be online or running.</p>
RESTART [seconds]	<p>Restarts Novell Cluster Services software on all servers in the cluster. 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>

Cluster Console Command	Description
SCAN FOR NEW DEVICES	<p>IMPORTANT: This command is deprecated 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> <p>WARNING: In EMC PowerPath environments, do not use the <code>rescan-scsi-bus.sh</code> utility provided with the operating system or the HBA vendor scripts for scanning the SCSI buses. To avoid potential file system corruption, EMC requires that you follow the procedure provided in the vendor documentation for EMC PowerPath for Linux.</p>
SET {Parameter} {Value}	<p>Sets cluster parameters individually for the cluster. See Chapter 8, “Configuring Cluster Policies, Protocols, and Properties,” on page 113 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>

Cluster Console Command	Description
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}	Reports the status of the specified resource. This includes the number of times the resource has been migrated or failed over to another server, the resource state, and the node where the resource is currently running.
VIEW	Displays the node name, cluster epoch number, master node name, and a list of nodes that are currently members of the cluster.

A.2 Business Continuity Clustering Commands

If Novell Business Continuity Clustering (BCC) 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 423](#)
- ♦ [Section A.3.2, “Options,” on page 424](#)
- ♦ [Section A.3.3, “Example,” on page 424](#)

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/ncpserver.preldif
```

A.3.2 Options

Replace the parameters with the credentials to access and location of the eDirectory schema files.

Parameter	Description	Example
port_num	The port number you assigned for eDirectory communications in the tree where you plan to install clusters. The default port is 636.	636
admin_username	The typeful fully distinguished user name of the administrator who has the eDirectory rights needed to extend the schema.	cn=admin,o=example
admin_password	The password of the administrator user.	pas5W0rd
server_ip_address	The IP address of the eDirectory server that contains the schema files.	10.10.10.1

A.3.3 Example

For example, enter the following commands in the order shown, using the values for your particular solution:

```
/opt/novell/oes-install/util/extend_schema --port 636 cn=admin,o=example pas5W0rd  
10.1.1.1 /opt/novell/ncs/schema/ncs.ldif
```

```
/opt/novell/oes-install/util/extend_schema --port 636 cn=admin,o=example pas5W0rd  
10.1.1.1 /opt/novell/ncs/schema/ncpserver.preldif
```

A.4 SBD Utility

The SBD utility (sbdutil) allows you to create, find, or view a 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 425](#)
- ♦ [Section A.4.2, “Options,” on page 425](#)

- [Section A.4.3, “Return Value,” on page 426](#)
- [Section A.4.4, “Examples,” on page 426](#)

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 426](#) 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 and later servers.

The utility can be used to scan for CSM containers and check if they can be imported for use in the cluster. This is the default mode when you run it without any options, or with only the `-v` option. By default, the utility searches block devices (`/dev/sd*`) for containers. Information about each device is displayed. Details about the CSM container are present if the container exists and can be imported for use in the cluster.

The utility also allows you to do the following:

- Search for a specified CSM container and check if it can be imported for use in the cluster.
- Import a specified CSM container for exclusive use in the cluster, and generate a corresponding Device Mapper object for it as `/dev/mapper/csm_name`.
- Export a specified CSM container by removing its Device Mapper object from the `/dev/mapper` directory. This allows the container to be imported by other nodes.

You can specify devices to search for any of the options. Wildcards such as the question mark (?) and asterisk (*) can be used when specifying devices. A question mark replaces a character in that position. An asterisk replaces any number of characters in that position.

- [Section A.5.1, “Syntax,” on page 427](#)
- [Section A.5.2, “Options,” on page 427](#)
- [Section A.5.3, “Examples,” on page 429](#)

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_EM005G` for CSM containers.


```
/opt/novell/ncs/bin/csmport -v /dev/mapper/MD_EMC_005G
```

The following command uses three question mark wildcards to search multipath devices with any three characters in those positions in the device name:

```
/opt/novell/ncs/bin/csmport -v /dev/mapper/MD_EMC_???G
```

The following command uses an asterisk wildcard to search multipath devices with any number of characters in that position in the device name:

```
/opt/novell/ncs/bin/csmport -v /dev/mapper/MD_EMC_*
```

-h

Prints the help page for this command.

-c <container_name> [device...]

Searches the `/dev/sd*` devices or specified devices for the named CSM container.

```
/opt/novell/ncs/bin/csmport -c container_name
```

-i <container_name> [device...]

Use this option to import a CSM container on a node for exclusive use in the cluster. The utility achieves this by searching for the container and if found, extracting the volume inside the container and exposing it exclusively at `/dev/mapper/container_name`.

The utility uses the container name instead of the Linux POSIX volume name because some CSM containers might contain compatibility volumes, which do not have names. The `/dev/mapper/container_name` object can be mounted at the old mount point that was used for the shared Linux POSIX volume in that container. The mount point can have the same name or different name than the Device Mapper object.

```
/opt/novell/ncs/bin/csmport -i container_name
```

Example

Suppose that you created a Linux POSIX cluster resource with the following settings on an OES 2 server or an OES 1 server:

```
RESOURCE_IP:          10.10.10.44
CONTAINER_NAME:       csm44
EVMS_VOLUME_NAME:    lxvol44
MOUNT_DEV:            /dev/evms/$CONTAINER_NAME/$EVMS_VOLUME_NAME
MOUNT_POINT:          /mnt/users
```

On an OES 11 or later server, you can use the following command to generate the Device Mapper object at `/dev/mapper/csm44`:

```
/opt/novell/ncs/bin/csmport -i csm44
```

The Device Mapper object can be mounted at the old mount point.

The resulting settings are:

```
RESOURCE_IP:          10.10.10.44
CONTAINER_NAME:       csm44
MOUNT_DEV:            /dev/mapper/$CONTAINER_NAME
MOUNT_POINT:          /mnt/users
```

-e <devmapper_object> [device...]

Use this option to export the Device Mapper object from the node. It removes the object from the `/dev/mapper` directory. Basically, the option undoes what the `-i` option did, so that the container can be imported by other nodes.

```
/opt/novell/ncs/bin/csmport -e devmapper_object
```


Example

The following command exports the Device Mapper object `/dev/mapper/csm44`.

```
/opt/novell/ncs/bin/csmport -e csm44
```

-v

Verbose mode can be used with any of the options.

-s

Silent mode can be used with any of the options.

A.5.3 Examples

This section provides examples of typical tasks you perform when using the `csmport` command.

- ♦ [“Scanning for CSM Containers” on page 429](#)
- ♦ [“Importing a CSM Container for Exclusive Use on a Node in a Cluster” on page 432](#)
- ♦ [“Exporting a CSM Container” on page 432](#)
- ♦ [“Searching for a Container on a Specified Device” on page 433](#)
- ♦ [“Searching for Containers on Multiple Devices with Wildcards” on page 433](#)
- ♦ [“Importing a CSM Container on a Specified Device” on page 433](#)

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 `CSM5GNoEVMSVo1` and can be imported to `/dev/mapper/CSM5GNoEVMSVo1` for its exclusive use on the cluster node.

Scan for CSM containers

```
Search device '/dev/sdn' for CSM containers.
Retrieve device information.
Device '/dev/sdn': major=8, minor=208.
Refresh device '/dev/sdn'.
Open device '/dev/sdn' for reading.
Check CSM metadata on device '/dev/sdn'.
/dev/sdn is not a CSM container.
```

```
Search device '/dev/sdm' for CSM containers.
Retrieve device information.
Device '/dev/sdm': major=8, minor=192.
Refresh device '/dev/sdm'.
Open device '/dev/sdm' for reading.
Check CSM metadata on device '/dev/sdm'.
/dev/sdm is not a CSM container.
```

```
Search device '/dev/sdh' for CSM containers.
Retrieve device information.
Device '/dev/sdh': major=8, minor=112.
Refresh device '/dev/sdh'.
Open device '/dev/sdh' for reading.
Check CSM metadata on device '/dev/sdh'.
/dev/sdh has a CSM container.
Cluster ID doesn't match. On disk = 37037513, in memory = 135657986.
Container name = 'AJTestContainer01' on disk.
Container is accessible (1).
```



```

Search device '/dev/sdc' for CSM containers.
Retrieve device information.
Device '/dev/sdc': major=8, minor=32.
Refresh device '/dev/sdc'.
Open device '/dev/sdc' for reading.
Check CSM metadata on device '/dev/sdc'.
/dev/sdc is not a CSM container.

Search device '/dev/sdl' for CSM containers.
Retrieve device information.
Device '/dev/sdl': major=8, minor=176.
Refresh device '/dev/sdl'.
Open device '/dev/sdl' for reading.
Check CSM metadata on device '/dev/sdl'.
/dev/sdl is not a CSM container.

Search device '/dev/sdi' for CSM containers.
Retrieve device information.
Device '/dev/sdi': major=8, minor=128.
Refresh device '/dev/sdi'.
Open device '/dev/sdi' for reading.
Check CSM metadata on device '/dev/sdi'.
/dev/sdi has a CSM container.
Container name = 'CSM5GNoEVMSVol' on disk.
Container is accessible (4).
No EVMS volume.
New name: /dev/mapper/CSM5GNoEVMSVol
Size: 10485757 sectors.
Device '/dev/sdi' has a CSM container 'CSM5GNoEVMSVol' and can be converted to
/dev/mapper/CSM5GNoEVMSVol.

Search device '/dev/sdk' for CSM containers.
Retrieve device information.
Device '/dev/sdk': major=8, minor=160.
Refresh device '/dev/sdk'.
Open device '/dev/sdk' for reading.
Check CSM metadata on device '/dev/sdk'.
/dev/sdk is not a CSM container.

Search device '/dev/sdg' for CSM containers.
Retrieve device information.
Device '/dev/sdg': major=8, minor=96.
Refresh device '/dev/sdg'.
Open device '/dev/sdg' for reading.
Check CSM metadata on device '/dev/sdg'.
/dev/sdg is not a CSM container.

Search device '/dev/sdb' for CSM containers.
Retrieve device information.
Device '/dev/sdb': major=8, minor=16.
Refresh device '/dev/sdb'.
Open device '/dev/sdb' for reading.
Check CSM metadata on device '/dev/sdb'.
/dev/sdb is not a CSM container.

Search device '/dev/sdt' for CSM containers.
Retrieve device information.
Device '/dev/sdt': major=65, minor=48.
Refresh device '/dev/sdt'.
Open device '/dev/sdt' for reading.
Failed to read the 1st sector of device '/dev/sdt': Input/output error (5).

Search device '/dev/sdx' for CSM containers.
Retrieve device information.
Device '/dev/sdx': major=65, minor=112.
Refresh device '/dev/sdx'.
Open device '/dev/sdx' for reading.
Failed to read the 1st sector of device '/dev/sdx': Input/output error (5).

```



```

Search device '/dev/sdj' for CSM containers.
  Retrieve device information.
  Device '/dev/sdj': major=8, minor=144.
  Refresh device '/dev/sdj'.
  Open device '/dev/sdj' for reading.
  Check CSM metadata on device '/dev/sdj'.
    /dev/sdj has a CSM container.
    Cluster ID doesn't match. On disk = 37037513, in memory = 135657986.
    Container name = 'CMS4GCon' on disk.
    Container is accessible (4).

Search device '/dev/sdy' for CSM containers.
  Retrieve device information.
  Device '/dev/sdy': major=65, minor=128.
  Refresh device '/dev/sdy'.
  Open device '/dev/sdy' for reading.
  Failed to read the 1st sector of device '/dev/sdy': Input/output error (5).

Search device '/dev/sdu' for CSM containers.
  Retrieve device information.
  Device '/dev/sdu': major=65, minor=64.
  Refresh device '/dev/sdu'.
  Open device '/dev/sdu' for reading.
  Failed to read the 1st sector of device '/dev/sdu': Input/output error (5).

Search device '/dev/sdd' for CSM containers.
  Retrieve device information.
  Device '/dev/sdd': major=8, minor=48.
  Refresh device '/dev/sdd'.
  Open device '/dev/sdd' for reading.
  Failed to read the 1st sector of device '/dev/sdd': Input/output error (5).

Search device '/dev/sdq' for CSM containers.
  Retrieve device information.
  Device '/dev/sdq': major=65, minor=0.
  Refresh device '/dev/sdq'.
  Open device '/dev/sdq' for reading.
  Check CSM metadata on device '/dev/sdq'.
    /dev/sdq is not a CSM container.

Search device '/dev/sds' for CSM containers.
  Retrieve device information.
  Device '/dev/sds': major=65, minor=32.
  Refresh device '/dev/sds'.
  Open device '/dev/sds' for reading.
  Failed to read the 1st sector of device '/dev/sds': Input/output error (5).

Search device '/dev/sdv' for CSM containers.
  Retrieve device information.
  Device '/dev/sdv': major=65, minor=80.
  Refresh device '/dev/sdv'.
  Open device '/dev/sdv' for reading.
  Failed to read the 1st sector of device '/dev/sdv': Input/output error (5).

Search device '/dev/sdaa' for CSM containers.
  Retrieve device information.
  Device '/dev/sdaa': major=65, minor=160.
  Refresh device '/dev/sdaa'.
  Open device '/dev/sdaa' for reading.
  Failed to read the 1st sector of device '/dev/sdaa': Input/output error (5).

Search device '/dev/sde' for CSM containers.
  Retrieve device information.
  Device '/dev/sde': major=8, minor=64.
  Refresh device '/dev/sde'.
  Open device '/dev/sde' for reading.
  Check CSM metadata on device '/dev/sde'.
    /dev/sde is not a CSM container.

Search device '/dev/sdz' for CSM containers.

```



```

Retrieve device information.
Device '/dev/sdz': major=65, minor=144.
Refresh device '/dev/sdz'.
Open device '/dev/sdz' for reading.
Failed to read the 1st sector of device '/dev/sdz': Input/output error (5).

Search device '/dev/sdw' for CSM containers.
Retrieve device information.
Device '/dev/sdw': major=65, minor=96.
Refresh device '/dev/sdw'.
Open device '/dev/sdw' for reading.
Failed to read the 1st sector of device '/dev/sdw': Input/output error (5).

Search device '/dev/sdr' for CSM containers.
Retrieve device information.
Device '/dev/sdr': major=65, minor=16.
Refresh device '/dev/sdr'.
Open device '/dev/sdr' for reading.
Failed to read the 1st sector of device '/dev/sdr': Input/output error (5).

Search device '/dev/sdp' for CSM containers.
Retrieve device information.
Device '/dev/sdp': major=8, minor=240.
Refresh device '/dev/sdp'.
Open device '/dev/sdp' for reading.
Failed to read the 1st sector of device '/dev/sdp': Input/output error (5).

Search device '/dev/sdo' for CSM containers.
Retrieve device information.
Device '/dev/sdo': major=8, minor=224.
Refresh device '/dev/sdo'.
Open device '/dev/sdo' for reading.
Failed to read the 1st sector of device '/dev/sdo': Input/output error (5).

Search device '/dev/sdf' for CSM containers.
Retrieve device information.
Device '/dev/sdf': major=8, minor=80.
Refresh device '/dev/sdf'.
Open device '/dev/sdf' for reading.
Check CSM metadata on device '/dev/sdf'.
/dev/sdf is not a CSM container.

Search device '/dev/sda' for CSM containers.
Retrieve device information.
Device '/dev/sda': major=8, minor=0.
Refresh device '/dev/sda'.
Open device '/dev/sda' for reading.
Check CSM metadata on device '/dev/sda'.
/dev/sda is not a CSM container.

```

Importing a CSM Container for Exclusive Use on a Node in a Cluster

Use the `-i` option to import the CSM container named `CSM5GNoEVMSVol` for exclusive use on a node in a cluster.

```
/opt/novell/ncs/bin/csmpport -i CSM5GNoEVMSVol
```

This command generates the Device Mapper object `/dev/mapper/CSM5GNoEVMSVol`. You can map this object to the desired mount point.

Exporting a CSM Container

Use the `-e` option to export the Device Mapper object `/dev/mapper/CSM5GNoEVMSVol`.

```
/opt/novell/ncs/bin/csmpport -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 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.

A.7 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.7.1, “Syntax,” on page 434](#)
- ♦ [Section A.7.2, “Options,” on page 434](#)
- ♦ [Section A.7.3, “Examples,” on page 434](#)

A.7.1 Syntax

At the command prompt on the master node of the cluster, enter the following command as the root user:

```
/opt/novell/ncs/bin/ncs_ncpserv.py -c <lx_volume_name> -i <resource_ip_address>
[-v <ncp_volume_name| "ncp_volume_name1:ncp_volume_name2:..."> ]
```

A.7.2 Options

-c <lx_volume_name>

Create a virtual NCP Server object for the specified volume.

-i <resource_ip_address>

Specify a unique static IP address to use for the cluster resource.

-v <ncp_volume_name | "ncp_volume_name1:ncp_volume_name2:...">

Use this option to specify one or multiple volumes on the Linux cluster resource that need a virtual NCP Server object to be created and bound to the specified IP address. If the -v option is not specified, all of the NCP volumes that currently exist on the Linux POSIX cluster resource are bound to the IP address.

If you enter multiple volume names, use colons to delimit the names and put quotation marks around the list of names. Volume names can be listed by the volume name (MY_NNCP_VOL06) or by the volume distinguished name (cn=CLUS_02_MY_NNCP_VOL06,o=novell), or any combination of the two methods.

A.7.3 Examples

In the following examples, the resource IP address is 10.10.10.44, the cluster name is cluster1 and the cluster context is ou=clusters,o=mycompany.

Example 1

To specify a single NCP volume named USERS on the lxvol44 cluster resource, enter

```
./ncs_ncpserv.py -c lxvol44 -i 10.10.10.44 -v USERS
```

The following confirmation message is displayed:

```
NCP Server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' created.
```

```
Object 'cn=servername_USERS,ou=clusters,o=mycompany' renamed to
'cn=cluster1_USERS,ou=clusters,o=mycompany'.
```

The volume name you need to use in the scripts is: USERS

```
NCP server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' and volume
'cn=cluster1_USERS,ou=clusters,o=mycompany' are linked with each other.
```

Example 2

To specify multiple NCP volumes on the lxvol44 cluster resource, enter

```
./ncs_ncpserv.py -c lxvol44 -i 10.10.10.44 -v
"USERS:MY_NCP_VOL06:cn=servername_MY_NCP_VOL07,ou=clusters,o=novell"
```


The following confirmation message is displayed:

```
NCP Server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' created.
```

```
Object 'cn=servername_USERS,ou=clusters,o=mycompany' renamed to  
'cn=cluster1_USERS,ou=clusters,o=mycompany'.
```

The volume name you need to use in the scripts is: USERS

```
NCP server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' and volume  
'cn=cluster1_USERS,ou=clusters,o=mycompany' are linked with each other.
```

```
Object 'cn=servername_MY_NCP_VOL06,ou=clusters,o=mycompany' renamed to  
'cn=cluster1_MY_NCP_VOL06,ou=clusters,o=mycompany'.
```

The volume name you need to use in the scripts is: MY_NCP_VOL06

```
NCP server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' and volume  
'cn=cluster_MY_NCP_VOL06,ou=clusters,o=mycompany' are linked with each other.
```

```
Object 'cn=servername_MY_NCP_VOL07,ou=clusters,o=mycompany' renamed to  
'cn=cluster1_MY_NCP_VOL07,ou=clusters,o=mycompany'.
```

The volume name you need to use in the scripts is: MY_NCP_VOL07

```
NCP server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' and volume  
'cn=cluster_MY_NCP_VOL07,ou=clusters,o=mycompany' are linked with each other.
```

A.8 NCPCON Mount Command for NCP and NSS Volumes

Use the `ncpcon mount` command to mount an NCP volume or NSS volume on your Linux server. This command makes the volume accessible to NCP clients.

- [Section A.8.1, “Syntax,” on page 435](#)
- [Section A.8.2, “NCP Volume Examples,” on page 435](#)
- [Section A.8.3, “NSS Volume Examples,” on page 436](#)
- [Section A.8.4, “Other Examples,” on page 436](#)

A.8.1 Syntax

```
mount < all | volume_name | volume_name=volume_id,path=/volume_mntpoint >
```

Replace *volume_name* with the name of the NSS volume or NCP volume, such as VOL1. To mount all local NCP and NSS volumes, replace *volume_name* with `all`. A clustered NCP volume or NSS volume is mounted when its cluster resource is brought online.

Replace *volume_id* with a value from 0 to 254 as the server volume ID. It is not necessary to manually specify a volume ID for a locally mounted volume. NCP automatically assigns unique volume IDs to locally mounted NCP volumes and NSS volumes in increasing order, from 0 to 254. IDs 0 and 1 are reserved for the `sys` and `_admin` volumes. When the command is used in a cluster resource script, the volume ID must be specified to ensure that the volume ID is unique across all cluster nodes where the volume will be mounted. By convention in clusters, the volume IDs are assigned in decreasing order, from 254 to 0.

A.8.2 NCP Volume Examples

For an NCP volume created on a Linux POSIX file system, replace `/volume_mntpoint` with the Linux path of the mount point for the NCP share. Typically, this path is mount point of the Linux volume; that is, the NCP share is created at the root of the Linux volume. The path must exist on all nodes in the cluster.

If an NCP volume is mounted locally, the mount path is stored in the `/etc/fstab` file, so it is not necessary to specify a mount path. For example:


```
ncpcon mount VOL1
```

The mount path for the NCP volume is required when you use the command in a cluster resource load script. For example, if you use NSSMU to create and NCP enable a clustered Linux LVM volume myvol with a mount point of /usr/novell/myvol, the NCP volume name (share name) is MYVOL, and the resource load script includes the following definitions for variables used in the command:

```
# define the mount point path of the Linux volume
MOUNT_PATH=/usr/novell/myvol

# define the NCP volume name
NCP_VOLUME=MYVOL

exit_on_error ncpcon mount $NCP_VOL=253,path=$MOUNT_PATH
```

A.8.3 NSS Volume Examples

For an NSS volume, the default mount path is /media/nss/<nss_volume_name>. For example, if you create an NSS volume named USERS, the default mount path is /media/nss/USERS. If you use the default path for NSS volumes, it is not necessary to include the path option for local mounts:

```
ncpcon mount USERS
```

The default path is also not required in a cluster script, but a volume ID that is unique across all cluster nodes must be specified:

```
exit_on_error ncpcon mount USERS=252
```

In a cluster, you cannot rename the mount point path. The default mount point path of /media/nss/<volume_name> applies.

```
exit_on_error ncpcon mount <volumename>=<volume_ID>
```

For example:

```
exit_on_error ncpcon mount USERS=252
```

The volume is mounted at /media/nss/USERS.

You can specify mount options for the NSS volume by using the /opt switch in the command. For example, you can specify the name space to use by using the /opt=ns=<long|unix|dos|mac> command.

```
# specify the name space to use when mounting the volume
exit_on_error ncpcon mount USERS=252 /opt=ns=unix
```

A.8.4 Other Examples

```
ncpcon mount sys
ncpcon mount all
```


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
/etc/init.d/novell-ncs	LSB Compliant Service
/etc/opt/novell/ncs/<node_name>	Configuration file for a specific cluster node
/etc/opt/novell/ncs/clstrlib.conf	Cluster configuration file
/lib/modules/kernel_dir/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.

NCS File Name and Path	Purpose
<code>/lib/modules/kernel_dir/ncs/vipx.ko</code>	Kernel module. Replace <i>kernel_dir</i> with the current kernel directory. Use <code>uname -r</code> to see the current kernel directory.
<code>/lib/modules/kernel_dir/ncs/vll.ko</code>	Kernel module. Replace <i>kernel_dir</i> with the current kernel directory. Use <code>uname -r</code> to see the current kernel directory.
<code>/opt/novell/ncs/bin/adminfs</code>	Cluster management (iManager and CLI)
<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/csmport</code>	CSM Import/Export utility
<code>/opt/novell/ncs/bin/ldncs</code>	Loads NCS; used by the <code>Cluster Start</code> command.
<code>/opt/novell/ncs/bin/ncs_ncpserv.py</code>	Virtual NCS NCP Server Object script
<code>/opt/novell/ncs/bin/NCS_STONITH_SCRIPT file</code>	An executable file that you create to authenticate to your power controller and turn off the power, cycle the power, or reset the power for the node. The presence of the file enables the STONITH feature for the node.
<code>/opt/novell/ncs/bin/ncs-configd.py</code>	Cluster configuration daemon
<code>/opt/novell/ncs/bin/ncs-emaild</code>	Cluster email daemon
<code>/opt/novell/ncs/bin/ncs-resourced.py</code>	Daemon used to run load and unload scripts.
<code>/opt/novell/ncs/bin/ncstempl.py</code>	Used to install or update cluster resource templates in the Cluster container
<code>/opt/novell/ncs/bin/sbdutil</code>	SBD partition utility
<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

NCS File Name and Path	Purpose
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Maintenance.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Migrate.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Monitor.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Offline.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Online.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Pools.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Rename.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Resources.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Restart.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Set.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Stats.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Status.pm	Cluster CLI command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_View.pm	Cluster CLI command
/opt/novell/ncs/bin/uldncs	Unloads NCS; used by the Cluster Stop command.
/opt/novell/ncs/lib/ncs-1.0.0.so	NCS snap-in
/opt/novell/ncs/lib/ncsfuns	Shared library commands for load/unload scripts
/opt/novell/ncs/schema/ncpserver.ldif	NCS schema file
/opt/novell/ncs/schema/ncpserver.preldif	NCS schema file
/opt/novell/ncs/schema/ncs.ldif	NCS schema file
/opt/novell/ncs/schema/ncs.sch	NCS schema file
/opt/novell/oes-install/util/	Path to the extend_schema command
/usr/include/ncssdk.h	NCS SDK
/usr/lib/libncssdk.so	NCS SDK
/usr/lib/libncssdk.so.1.0.0	NCS SDK

NCS File Name and Path	Purpose
<code>/usr/sbin/rcnovell-ncs</code>	Link to <code>/etc/init.d/novell-ncs</code>
/usr/share/man/man7/sbdutil.7.gz	SBDUTIL man page
/usr/share/man/man8/cluster.8.gz	CLUSTER man page
/usr/share/man/man8/csmport.8.gz	CSMPORT man page

C Electing a Master Node

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

If a slave node fails to receive heartbeat packages from a master node within a predefined tolerance, the slave assumes that the master has left the cluster or is experiencing problems. The slave node, together with other slave nodes, tries to elect a new master node. The goal is to identify a node with a higher IP address that can be seen by the most member nodes. Each slave node follows the same process to identify a master candidate. If multiple masters are elected, the SBD (split-brain detector) guarantees that only one master survives and fences the other masters and their members.

Novell Cluster Services for Novell Open Enterprise Server (OES) 11 Support Pack 1 (SP1) introduces some intelligence in the master election process when the master leaves a cluster (voluntarily or involuntarily). The same master is elected under both the old and new algorithms, but the conclusion is reached sooner, especially for larger clusters. The new algorithm substantially reduces the time needed for master election in some cases. The algorithm is available for OES 11 SP1, and will be available as patches in September 2012 for OES 11 and OES 2 SP3.

This section briefly describes the old and new master-election algorithms.

- ♦ [Section C.1, “Master-Election Algorithm for OES 11 and Earlier,” on page 441](#)
- ♦ [Section C.2, “New Master-Election Algorithm for OES 11 SP1,” on page 442](#)

C.1 Master-Election Algorithm for OES 11 and Earlier

In Novell Cluster Services for OES 11 and earlier, after losing contact with the master node, a slave node attempts to find a node with a higher IP address that will act as master. The logic flow for the master-election algorithm is represented in [Figure C-1, “Master-Election Algorithm for OES 11 Initial Release and Earlier Versions,” on page 442](#).

Before polling other nodes, a node checks its own network adapter status. If the adapter is down, the node immediately promotes itself as the master. Rather than polling and waiting for replies that cannot come back, the node can jump directly to the ultimate conclusion for that node that it is the only node remaining in the cluster. This shortcut can substantially shorten the election process if the network outage is caused by a network adapter failure. However, if the cause for a network outage occurs further upstream by a cable or switch failure, the node follows the election process, unaware that it is awaiting a response that cannot arrive.

If a node's adapter is working, it begins a process to find a live node with a higher IP address than its own:

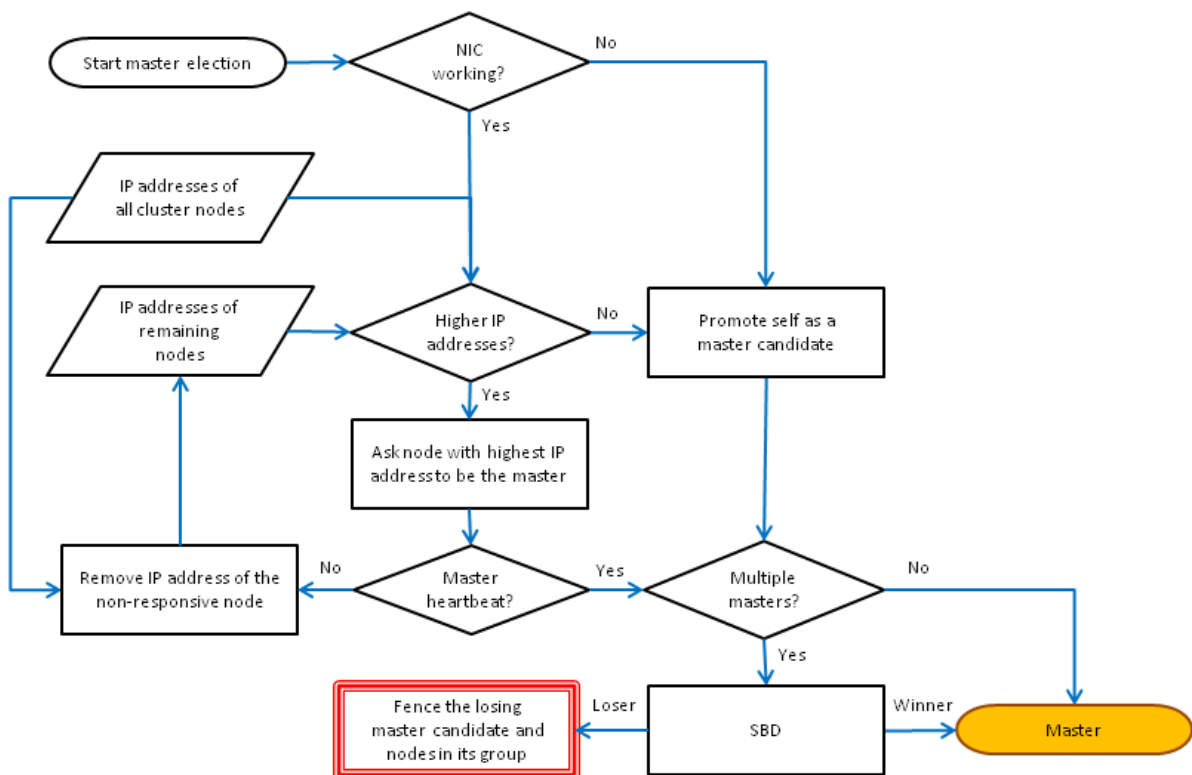
1. A node determines if there are other nodes (excluding the old master) with an IP address higher than its own. If it has the highest IP address, it promotes itself as the master.

2. If there are nodes with higher IP addresses, a node asks a member node with the highest IP address (excluding the old master) to be the new master. It waits for a predetermined tolerance to see if that master candidate node will start acting like a master and send it heartbeat packages. If a heartbeat package arrives, the node becomes a slave of its elected master.
3. If no heartbeat package arrives, the node then picks the member node with the second highest IP address and repeats the same procedure.
4. If there are no responses from nodes with a higher IP address than its own, a node promotes itself as the master.

If more than one node is elected as master at the end of the process, the SBD guarantees that only one master survives and fences the other masters and their members. Describing how the SBD addresses all possible situations to determine the master is beyond the scope of this document. Let's consider a simple scenario with multiple master candidates. The SBD kills any master candidate node whose adapter is down. If a master candidate was the old master, it becomes the new master. Otherwise, the master candidate with the highest IP address wins.

The master election process can take time because a node sequentially asks each potential master node to become the master and waits for a response before trying the next higher IP address. If a large cluster (8+ nodes) loses LAN connections among all the nodes, it can take up to 3 minutes to elect all of the new masters. The node with lowest IP address tries almost all of the other nodes.

Figure C-1 Master-Election Algorithm for OES 11 Initial Release and Earlier Versions



C.2 New Master-Election Algorithm for OES 11 SP1

In Novell Cluster Services for OES 11 SP1 (and fully patched OES 11 and OES 2 SP3), after losing contact with the master node, a slave node attempts to find a node with a higher IP address that will act as master, but it also uses intelligence to determine which other nodes are capable of responding. The new algorithm elects the same master as the old algorithm, but the time to make the decision is

deterministic and the process is more efficient. The logical flow for the master-election algorithm is represented in [Figure C-2, “Master-Election Algorithm for OES 11 SP1 \(and Fully Patched OES 11 and OES 2 SP3\),” on page 444.](#)

In the new master-election algorithm, after a slave determines that its adapter is working, it broadcasts a ping to all of the cluster nodes. The purpose of pinging is to identify whether the node can communicate with other nodes and which ones. This helps address a situation where LAN communications are down between some nodes, such as a cable or switch failure.

If there are nodes with higher IP addresses than its own, a node asks a member node with the highest IP address (excluding the old master) to be the new master. It waits for a predetermined tolerance to see if that master candidate node will start acting like a master and send it heartbeat packages. If a heartbeat package arrives, the node becomes a slave of its elected master.

Sending the first master request in the old manner serves two purposes:

- ♦ It maintains some compatibility with older versions of Novell Cluster Services.
- ♦ It provides some action while waiting for the ping replies to come in.

If the member with the highest IP address starts acting as the new master (in time), the election is over for this node.

If no heartbeat packages arrive, the node then focuses on member nodes with higher IP addresses that have either pinged or replied to its ping. This helps the node by-pass waiting periods for requests to nodes that are unable to communicate. The node picks the member node with the highest IP address, and asks it to be the new master. It waits for a predetermined tolerance to see if that master candidate node will start acting like a master and send it heartbeat packages. The request to this node is almost certain to succeed since it is known to be able to communicate.

If no heartbeat packages arrive, the node picks the member node with the next highest IP address that has either pinged or replied to the ping, and the election goes on.

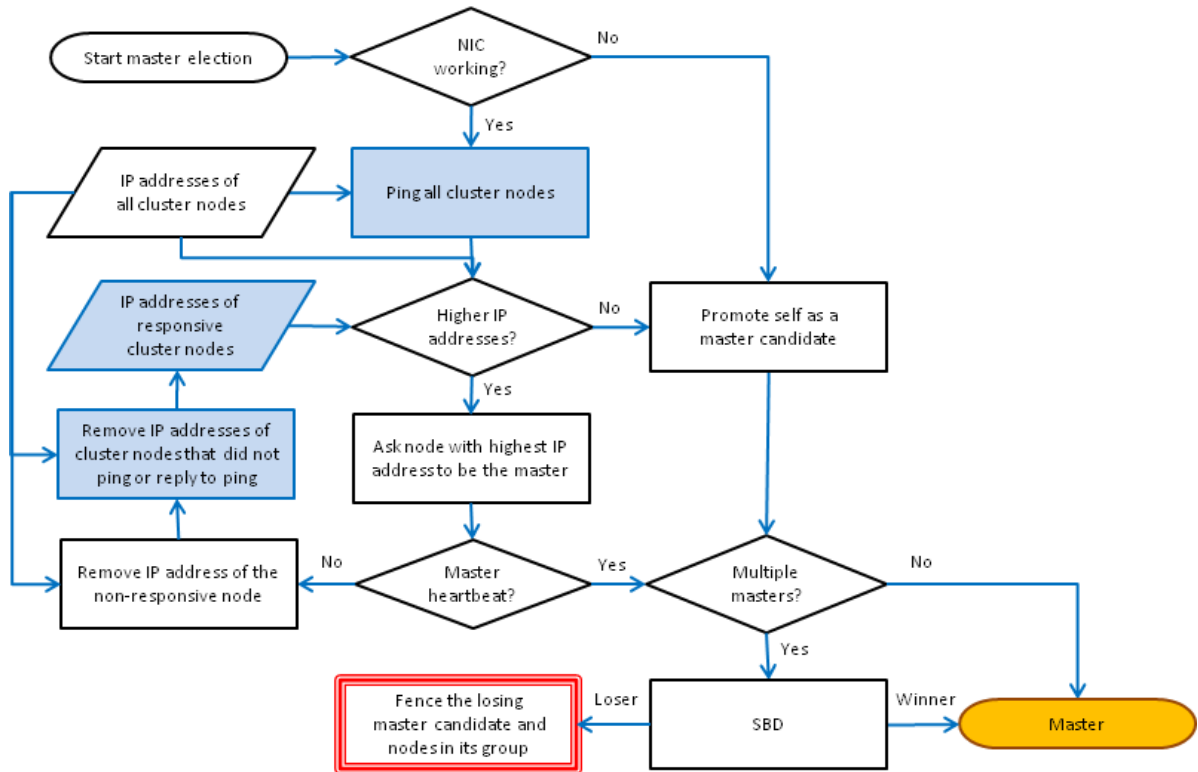
If there are no responses from nodes with a higher IP address than its own, a node promotes itself as the master.

If more than one node is elected as master at the end of the process, the SBD guarantees that only one master survives and fences the other masters and their members.

The benefits of the new master-election algorithm are:

- ♦ The new algorithm elects the same master that would be determined with the old algorithm.
- ♦ The time to elect a master node is deterministic, regardless of the size of the cluster or the nature of the problem.
- ♦ The traditional master-election logic is mostly preserved.
- ♦ The first master request goes out while waiting for nodes to respond to pings, or to time out.
- ♦ The second master request goes to a node that is known to be able to communicate, and is almost certain to succeed.
- ♦ Election time is substantially decreased in some cases.

Figure C-2 Master-Election Algorithm for OES 11 SP1 (and Fully Patched OES 11 and OES 2 SP3)



D Comparing Novell Cluster Services for Linux and NetWare

[Table D-1](#) compares the features and capabilities of Novell Cluster Services on Novell Open Enterprise Server (OES) 11 SP1 to Novell Cluster Services on NetWare 6.5 SP8.

Table D-1 Comparison of Novell Cluster Services for OES 11 SP1 and NetWare 6.5 SP8

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
Operating system	OES 11 SP1	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
Guest servers on KVM VMs as cluster nodes	Yes	Yes
Business Continuity Clustering support	BCC 2.0 is planned for OES 11 SP1. 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 5.7, “Configuring Additional Administrators,” on page 96.	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 74.</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 76.</p>	The user who installs Novell Cluster Services must have schema extension rights.
Forward migration for Novell Cluster Services	<p>OES 2 SP3 to OES 11 SP1. 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 7, “Upgrading Clusters from OES 2 SP3 to OES 11x,” on page 105.</p>	<p>NetWare 6.5 SP8 to OES 11 SP1. 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	NetWare 6.5 SP8 to OES 11 SP1	Not applicable
mixed-mode Linux and NetWare clusters	Supported only for rolling cluster conversions from NetWare to Linux. For information, see OES 11 SP1: Novell Cluster Services 2.1 NetWare to Linux Conversion Guide .	Supported only for rolling cluster conversions from NetWare to Linux. NSS pools created on NetWare can fail over to Linux nodes.
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.

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
Mirrored SBD	<p>During the install.</p> <p>After the install and before adding a second node to the cluster by using the sbdutil and specifying two devices.</p> <p>Mirroring an existing SBD by using NLVM.</p>	<p>During the install on the first server in the cluster, or by using the sbdutil after the install.</p>
Shared disks	<p>Fibre Channel SAN LUNs</p> <p>iSCSI SAN LUNs</p> <p>SCSI disks (shared external drive arrays)</p>	<p>Fibre Channel SAN LUNs</p> <p>iSCSI SAN LUNs</p> <p>SCSI disks (shared external drive arrays)</p>
Cluster-aware shared devices	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.	<p>NSS is the default file system on NetWare.</p> <p>Novell Cluster Services is not supported on NetWare traditional volumes.</p>
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	<p>Yes</p> <p>For information, see Chapter 12, “Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes,” on page 217.</p> <p>Shareable for Clustering</p> <p>Multiple-Server Activation Prevention (MSAP)</p> <p>Cluster volume broker; Linux kernel module handles NSS pool events.</p>	<p>Yes</p> <p>For information, see “Setting Up Cluster Resources for Novell Cluster Services” in the <i>NW6.5 SP8: Novell Cluster Services 1.8.5 Administration Guide</i>.</p> <p>Shareable for Clustering</p> <p>Multiple-Server Activation Prevention (MSAP)</p> <p>Cluster volume broker</p>

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 13, “Configuring and Managing Cluster Resources for Shared LVM Volume Groups,” on page 305</p> <p>For information about managing resources migrated from OES 2 SP3 to OES 11 SP1, see Chapter 14, “Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers,” on page 371.</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 SP1: 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 Volume Pairs with Novell Cluster Services” in the <i>OES 11 SP1: 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 15.2, “Virtual Machines as Cluster Resources,” on page 392.</p>	Not applicable
iManager 2.7.5	Yes	Yes
Clusters plug-in and Storage Management plug-in for iManager	Yes	Yes
My Clusters and My Resources feature of the Clusters plug-in	Yes	Can be used to manage NetWare 6.5 SP8 servers, but cannot be run on NetWare servers.
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
Creating cluster enabled NSS pools by using NLVM commands	Yes	Not applicable

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
Creating cluster enabled LVM volume groups by using NSSMU	Yes	Not applicable
Creating cluster enabled LVM volume groups by using NLVM commands	Yes	Not applicable
Command line interface	Yes; using the terminal console as the <code>root</code> user	Yes; using the terminal console
XML-based API	Yes; same as NetWare except that it uses the <code>/_adminfs</code> path on Linux.	Yes; same as for Linux except that it uses the <code>_admin</code> volume on NetWare.
Load, unload, and monitor scripts	Yes Script commands differ. Scripts are automatically translated from NetWare commands to Linux commands during the cluster conversion from NetWare to Linux. For a comparison of script commands, see “ Planning the Conversion of Load and Unload Scripts ” in the <i>OES 11 SP1: Novell Cluster Services 2.1 NetWare to Linux Conversion Guide</i> .	Load and unload scripts; no monitoring
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
NCP support for accessing files on Dynamic Storage Technology volumes made up of two NSS volumes	Yes	DST not available
Novell AFP support for accessing files on shared NSS pools	Yes; cross-protocol file locking available with NCP and Novell CIFS, or with NCP and Novell Samba	Yes; with cross-protocol file locking
Novell CIFS support for accessing files on shared NSS pools	Yes; cross-protocol file locking available with NCP and Novell AFP, or with NCP and Novell Samba	Yes; with cross-protocol file locking
Novell CIFS support for accessing files on Dynamic Storage Technology volumes made up of two NSS volumes	Yes; cross-protocol file locking available in NCP	DST not available
Linux Samba/CIFS support for accessing files on shared NSS pools on Linux	Yes; with cross-protocol file locking available in NCP Requires users to be Linux-enabled with Linux User Management. Requires Universal Password.	Not applicable

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
Linux Samba/CIFS support for accessing files on shared Linux POSIX file systems	Yes Requires users to be enabled with Linux User Management. Requires Universal Password.	Not applicable
Leverage Heartbeat 2 resource agents	Yes	Not applicable
LAN fault tolerance	Channel bonding For information, see <code>/usr/src/linux/Documentation/bonding.txt</code>	NIC teaming For information, see “ NIC Teaming ” in the <i>NW 6.5 SP8: TCP/IP Administration Guide</i>
Multipath I/O	Device Mapper - Multipath I/O, or third-party MPIO solutions For information, see “ Managing Multipath I/O for Devices ” (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)	Media Manager Multipath I/O, or third-party MPIO solutions For information, see “ Managing Multipath I/O to Devices (NetWare) ” in the <i>NW 6.5 SP8: NSS File System Administration Guide</i> .
Monitor script	Yes	No
Resource Mutual Exclusion Groups	Yes	No
STONITH (Feature Preview)	Yes	No

E Comparing Resources Support for Linux and NetWare

[Table E-1](#) compares clustering support for Novell Open Enterprise Server (OES) 11 SP1 services when using Novell Cluster Services on OES 11 SP1 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 SP1: Novell Cluster Services 2.1 NetWare to Linux Conversion Guide](#).

Table E-1 Comparison of Clustering Support for OES 11 Services on Linux and NetWare

Service	NetWare 6.5 SP8	OES 11 SP1	For conversion information, see
AFP (Apple Filing Protocol)	Yes See “Setting Up for Macintosh” in the <i>NW 6.5 SP8: AFP, CIFS, and NFS (NFAP) Administration Guide</i> .	Yes See “Configuring AFP with Novell Cluster Services for an NSS File System” in the <i>OES 11 SP1: Novell AFP for Linux Administration Guide</i> .	“Novell AFP” in the <i>OES 11 SP1: Novell Cluster Services 2.1 NetWare to Linux Conversion Guide</i>
Apache Web Server	Yes See “Apache with Novell Cluster Services” in the <i>NW6.5 SP8: Novell Cluster Services 1.8.5 Resource Configuration Guide</i> .	Yes; use the standard Apache Web Server for Linux.	“Apache HTTP Server” in the <i>OES 11 SP1: Novell Cluster Services 2.1 NetWare to Linux Conversion Guide</i>
Archive and Version Services (Novell)	Yes Uses a MySQL database. See “Installing and Configuring New Archive Servers” in the <i>NW 6.5 SP8: Novell Archive and Version Services 2.1 Administration Guide</i> .	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 <i>OES 11 SP1: Novell Archive and Version Services 2.1 Administration Guide</i> .	“Novell Archive and Version Services” in the <i>OES 11 SP1: Novell Cluster Services 2.1 NetWare to Linux Conversion Guide</i>

Service	NetWare 6.5 SP8	OES 11 SP1	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 SP1: Novell CIFS for Linux Administration Guide</i> .	“ Novell CIFS ” in the <i>OES 11 SP1: Novell Cluster Services 2.1 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 SP1: Novell Cluster Services 2.1 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 SP1: Novell DNS/DHCP Services for Linux Administration Guide</i> .	“ DHCP Server ” in the <i>OES 11 SP1: Novell Cluster Services 2.1 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 SP1: Novell DNS/DHCP Services for Linux Administration Guide</i> .	“ DNS Server ” in the <i>OES 11 SP1: Novell Cluster Services 2.1 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 SP1 node and you set the global server-level parameters to be the same on each one.
eDirectory 8.8x	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 SP1	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 3.3.8 Administration Guide .	“eDirectory Server Certificates” in the OES 11 SP1: Novell Cluster Services 2.1 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 Novell FTP (Pure-FTPd) solution. See “ Cluster Enabling Pure-FTPd in an OES 11 SP1 Environment ” in the OES 11 SP1: 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 SP1: iPrint Linux Administration Guide .	“Novell iPrint” in the OES 11 SP1: Novell Cluster Services 2.1 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 SP1: Novell Cluster Services 2.1 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 SP1	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 SP1: 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 3.9.1 Administration Guide</i> .	“ Migration ” in the <i>OES 11 SP1: 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 SP1: Novell QuickFinder Server 5.0 Administration Guide</i> .	“ QuickFinder Server ” in the <i>OES 11 SP1: Novell Cluster Services 2.1 NetWare to Linux Conversion Guide</i>
Storage, DST shadow volume pairs	Not applicable	Yes See “ Configuring DST Shadow Volume Pairs with Novell Cluster Services ” in the <i>OES 11 SP1: 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 SP1	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 SP1: 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 12, “Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes,” on page 217. For a feature comparison, see “Cross-Platform Issues for NSS” in the <i>OES 11 SP1: NSS File System Administration Guide for Linux</i> .	“Novell Storage Services Pools” in the <i>OES 11 SP1: Novell Cluster Services 2.1 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 SP1 nodes can be used in OES 11 SP1 clusters. Nodes can be any combination of virtual and physical servers.	See Chapter 15, “Configuring Novell Cluster Services in a Virtualization Environment,” on page 391.
Xen virtual machines on the host server	Not applicable	Yes; use the Xen and XenLive templates.	See Section 15.2, “Virtual Machines as Cluster Resources,” on page 392.

F Clusters Plug-In Changes for Novell iManager 2.7.5

The Clusters plug-in for Novell iManager 2.7.5 or later supports the management of OES and NetWare clusters and resources. The availability of Novell Cluster Services features depends on the version of Novell Cluster Services and the server platform that are installed on the cluster being managed. You can use either the old plug-in or the new plug-in to manage an OES 2 SP3 cluster and its resources.

Understanding the Cluster Plug-In's New Interface

The Clusters plug-in for Novell iManager 2.7.5 or later has been reorganized. In *Roles and Tasks* under *Clusters*, the *Cluster Manager*, *BCC Manager*, *Cluster Event Log*, and *Cluster Options* menu items have been replaced with two menu options:

- ♦ **My Clusters:** The logged-in cluster administrator can set up a personalized list of clusters to manage. This allows an administrator to view at a glance the status of multiple clusters. An administrator can also customize the display to sort the entries, modify the columns, or filter the entries. The list of clusters and display preferences persist between the administrator's logins to iManager on the same server. For information, see "[Setting Up a Personalized List of Clusters to Manage](http://www.novell.com/documentation/oes11/clus_admin_lx/data/myclusters.html)" (http://www.novell.com/documentation/oes11/clus_admin_lx/data/myclusters.html) in the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).
- ♦ **My Resources:** The logged-in cluster administrator can set up a personalized list of cluster resources to manage. This allows an administrator to view at a glance the status of multiple cluster resources for multiple clusters. An administrator can also customize the display to sort the entries, modify the columns, or filter the entries. The list of cluster resources and display preferences persist between the administrator's logins to iManager on the same server. For information, see "[Setting Up a Personalized List of Resources to Manage](http://www.novell.com/documentation/oes11/clus_admin_lx/data/myresources.html)" (http://www.novell.com/documentation/oes11/clus_admin_lx/data/myresources.html) in the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

From a customized My Clusters page or a My Resources page, you can click the name link of a cluster to manage the cluster. The old plug-in's *Clusters* menu options are available here as tabs: *Cluster Manager*, *BCC Manager*, *Cluster Event Log*, and *Cluster Options*.

Upgrading the Clusters Plug-In

If you use Role-Based Services (RBS), upgrading the Clusters plug-in does not automatically update the RBS settings. The RBS Configuration page reports that the Clusters plug-in is out-of-date. The plug-in must be reinstalled on the RBS Configuration page in order to pick up the *My Clusters* and *My Resources* menu options. For information, see "[Updating Role-Based Services for the Clusters Plug-in for OES 11 SP1](http://www.novell.com/documentation/oes11/clus_admin_lx/data/inst_clusters.html#inst_clusters_rb)" (http://www.novell.com/documentation/oes11/clus_admin_lx/data/inst_clusters.html#inst_clusters_rb) in the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

Comparing Tasks in the Old and New Clusters Plug-Ins

Use the tables in this section to understand how to perform tasks in the old Clusters plug-in for iManager 2.7.4 and the new Clusters plug-in for iManager 2.7.5.

- ♦ [Table F-1, “Selecting a Cluster to Manage,” on page 458](#)
- ♦ [Table F-2, “Managing a Cluster,” on page 459](#)
- ♦ [Table F-3, “Accessing Cluster Properties,” on page 460](#)
- ♦ [Table F-4, “Managing a Cluster Resource or Resource Properties,” on page 461](#)

[Table F-1](#) describes how to select the cluster you want to manage.

Table F-1 *Selecting a Cluster to Manage*

Clusters Plug-In for iManager 2.7.4 and Earlier	Clusters Plug-In for iManager 2.7.5 and Later
<ol style="list-style-type: none">1. In <i>Roles and Tasks</i>, select <i>Clusters</i>, then select <i>Cluster Manager</i>, <i>BCC Manager</i>, <i>Cluster Event Log</i>, or <i>Cluster Options</i>.2. Specify the cluster name, or browse and select the Cluster object. You must repeat the selection process to manage a different cluster. There is no capability to concurrently view the status of multiple clusters.	<ol style="list-style-type: none">1. In <i>Roles and Tasks</i>, select <i>Clusters > My Clusters</i>. The list is initially empty.2. On the My Clusters page, click <i>Add</i> to open the Novell eDirectory browser pop-up window.3. Browse the tree where you are currently logged in to locate and select one or more clusters that you want to manage, then click <i>OK</i>. The selected clusters are added to the list. Your personalized list is displayed each time you log in to this iManager server and return to the My Clusters page. Use the <i>Add</i> and <i>Remove</i> options to modify your list as needed.4. Do any of the following from this page:<ul style="list-style-type: none">♦ Concurrently view the status of multiple clusters.♦ Click the name link of a cluster to manage it. The Cluster Manager page opens. You can also select the other tabs from this page: <i>BCC Manager</i>, <i>Cluster Event Log</i>, or <i>Cluster Options</i>.♦ Select the check box next to the cluster, then select <i>Actions > Edit Properties</i> to manage the cluster properties.♦ Select the check box next to the cluster, then select <i>Actions > Run Report</i>.

[Table F-2](#) describes how to access the cluster management pages: Cluster Manager, BCC Manager, Cluster Event Log, and Cluster Options. The instructions for the Clusters plug-in for iManager 2.7.5 and later assume that you have created a personalized list of clusters, as described in [Table F-1](#).

Table F-2 *Managing a Cluster*

Clusters Plug-In for iManager 2.7.4 and Earlier	Clusters Plug-In for iManager 2.7.5 and Later
Cluster Manager <ol style="list-style-type: none"> 1. In <i>Roles and Tasks</i>, select <i>Clusters > Cluster Manager</i>. 2. Specify the cluster name, or browse and select the Cluster object. 	Cluster Manager <ol style="list-style-type: none"> 1. In <i>Roles and Tasks</i>, select <i>Clusters > My Clusters</i>. 2. Click the name link of the cluster. 3. Select the <i>Cluster Manager</i> tab.
BCC Manager <ol style="list-style-type: none"> 1. In <i>Roles and Tasks</i>, select <i>Clusters > BCC Manager</i>. 2. Specify the cluster name, or browse and select the Cluster object. 	BCC Manager <ol style="list-style-type: none"> 1. In <i>Roles and Tasks</i>, select <i>Clusters > My Clusters</i>. 2. Click the name link of the cluster. 3. Select the <i>BCC Manager</i> tab.
Cluster Event Log <ol style="list-style-type: none"> 1. In <i>Roles and Tasks</i>, select <i>Clusters > Cluster Event Log</i>. 2. Specify the cluster name, or browse and select the Cluster object. 	Cluster Event Log <ol style="list-style-type: none"> 1. In <i>Roles and Tasks</i>, select <i>Clusters > My Clusters</i>. 2. Click the name link of the cluster. 3. Select the <i>Cluster Event Log</i> tab.
Cluster Options <ol style="list-style-type: none"> 1. In <i>Roles and Tasks</i>, select <i>Clusters > Cluster Options</i>. 2. Specify the cluster name, or browse and select the Cluster object. 	Cluster Options <ol style="list-style-type: none"> 1. In <i>Roles and Tasks</i>, select <i>Clusters > My Clusters</i>. 2. Click the name link of the cluster. 3. Select the <i>Cluster Options</i> tab.

Table F-3 describes how to access the cluster properties pages: Policies, Priorities, Protocols, Resource Mutual Exclusion, and Business Continuity. The instructions for the Clusters plug-in for iManager 2.7.5 and later assume that you have created a personalized list of clusters, as described in Table F-1.

Table F-3 *Accessing Cluster Properties*

Clusters Plug-In for iManager 2.7.4 and Earlier	Clusters Plug-In for iManager 2.7.5 and Later
	From My Resources 1. In <i>Roles and Tasks</i> , select <i>Clusters > My Clusters</i> . 2. Select the check box next to the cluster, then click <i>Actions > Edit Properties</i> .
From Cluster Options 1. In <i>Roles and Tasks</i> , select <i>Clusters > Cluster Options</i> . 2. Specify the cluster name, or browse and select the Cluster object. 3. Click <i>Properties</i> .	From Cluster Options 1. In <i>Roles and Tasks</i> , select <i>Clusters > My Clusters</i> . 2. Click the name link of the cluster. 3. Select the <i>Cluster Options</i> tab. 4. Click <i>Properties</i> .

Table F-4 describes how to view or manage cluster resources, including their properties: Policies, Monitoring, Preferred Nodes, Scripts (Load, Unload, and Monitor), and Business Continuity.

Table F-4 Managing a Cluster Resource or Resource Properties

Clusters Plug-In for iManager 2.7.4 and Earlier	Clusters Plug-In for iManager 2.7.5 and Later
<p>There is no capability to concurrently view the status of cluster resources across multiple clusters.</p>	<p>From My Resources:</p> <ol style="list-style-type: none"> 1. In <i>Roles and Tasks</i>, select <i>Clusters > My Resources</i>. <p>The initial list is empty.</p> <ol style="list-style-type: none"> 2. On the My Resources page, click <i>Add</i> to open the Novell eDirectory browser pop-up window. 3. Browse the tree where you are currently logged in to locate and select the cluster resources that you want to manage for one or more clusters, then click <i>OK</i>. <p>The selected cluster resources are added to the list.</p> <p>Your personalized list is displayed each time you log in to this iManager server and return to the My Resources page. Use the <i>Add</i> and <i>Remove</i> options to modify your list as needed.</p> <ol style="list-style-type: none"> 4. Do any of the following from this page: <ul style="list-style-type: none"> ◆ Concurrently view the status of multiple resources across multiple clusters. ◆ Click the name link of a resource to open its <i>Resource Properties</i>. ◆ Click the name link of the resource's cluster to manage the cluster.
<p>From Cluster Manager:</p> <ol style="list-style-type: none"> 1. In <i>Roles and Tasks</i>, select <i>Clusters > Cluster Manager</i>. 2. Specify the cluster name, or browse and select the Cluster object. 3. Do any of the following: <ul style="list-style-type: none"> ◆ View the status of resources on this cluster. ◆ Click the name link of a resource to open its <i>Resource Properties</i>. ◆ Select the check box next to a resource, then click one of the following options: <ul style="list-style-type: none"> Online Offline Migrate Respond to Alert 	<p>From Cluster Manager:</p> <ol style="list-style-type: none"> 1. In <i>Roles and Tasks</i>, select <i>Clusters > My Clusters</i>. 2. Click the name link of the cluster. 3. Select the <i>Cluster Manager</i> tab. 4. Do any of the following: <ul style="list-style-type: none"> ◆ View the status of resources on this cluster. ◆ Click the name link of a resource to open its <i>Resource Properties</i>. ◆ Select the check box next to a resource, then click one of the following options: <ul style="list-style-type: none"> Online Offline Migrate Respond to Alert

Clusters Plug-In for iManager 2.7.4 and Earlier	Clusters Plug-In for iManager 2.7.5 and Later
<p>From Cluster Options:</p> <ol style="list-style-type: none"> 1. In <i>Roles and Tasks</i>, select <i>Clusters > Cluster Options</i>. 2. Specify the cluster name, or browse and select the Cluster object. 3. Click the name link of a resource to open its <i>Resource Properties</i>. <p>You can also select the check box next to the resource, then click the <i>Details</i> link.</p>	<p>From Cluster Options:</p> <ol style="list-style-type: none"> 1. In <i>Roles and Tasks</i>, select <i>Clusters > My Clusters</i>. 2. Click the name link of the cluster. 3. Select the <i>Cluster Options</i> tab. 4. Click the name link of a cluster resource to open its <i>Resource Properties</i>. <p>You can also select the check box next to the resource, then click the <i>Details</i> link.</p>

G 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 G.1, “January 6, 2014,” on page 464](#)
- ♦ [Section G.2, “December 6, 2013,” on page 464](#)
- ♦ [Section G.3, “September 2013 Scheduled Maintenance,” on page 464](#)
- ♦ [Section G.4, “September 21, 2013,” on page 466](#)
- ♦ [Section G.5, “September 6, 2013,” on page 467](#)
- ♦ [Section G.6, “August 16, 2013,” on page 468](#)
- ♦ [Section G.7, “August 6, 2013,” on page 468](#)
- ♦ [Section G.8, “July 30, 2013,” on page 468](#)
- ♦ [Section G.9, “July 11, 2013,” on page 470](#)
- ♦ [Section G.10, “May 2013 Scheduled Maintenance,” on page 470](#)
- ♦ [Section G.11, “May 21, 2013,” on page 472](#)
- ♦ [Section G.12, “May 7, 2013,” on page 473](#)
- ♦ [Section G.13, “April 2013 Scheduled Maintenance,” on page 473](#)
- ♦ [Section G.14, “April 18, 2013,” on page 474](#)
- ♦ [Section G.15, “April 12, 2013,” on page 475](#)
- ♦ [Section G.16, “January 2013 Scheduled Maintenance,” on page 477](#)
- ♦ [Section G.17, “November 2012 Scheduled Maintenance,” on page 478](#)
- ♦ [Section G.18, “September 2012 Scheduled Maintenance,” on page 479](#)
- ♦ [Section G.19, “September 12, 2012,” on page 480](#)
- ♦ [Section G.20, “August 28, 2012 \(OES 11 SP1\),” on page 481](#)

G.1 January 6, 2014

Updates were made to the following section. The changes are explained below.

G.1.1 Upgrading OES 11 Clusters

Location	Change
Section 6.2, "Adding OES 11 SP1 Nodes to an OES 11 Cluster (Rolling Cluster Upgrade)," on page 102	If autostart is configured for the node, to prevent NCS from automatically loading during upgrade, you need to run <code>chkconfig -s novell-ncs off</code> after NCS has been stopped (before the upgrade), and <code>chkconfig -s novell-ncs on</code> after NCS has been started (after the upgrade).

G.2 December 6, 2013

The new site for iManager 2.7.6 and earlier is at <https://www.netiq.com/documentation/imanager27>.

Updates were made to the following sections. The changes are explained below.

G.2.1 Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 12.14, "Renaming the Mount Point Path for a Clustered NSS Volume," on page 275	The NSS volume attribute <i>Allow renaming mount point</i> is enabled by default.

G.2.2 Troubleshooting Novell Cluster Resources

Location	Change
Section 16.7, "iManager: Error 500 Occurs after Installing or Updating the Clusters Plug-In for iManager 2.7.5," on page 408	This problem is resolved in the Clusters plug-in for iManager 2.7.6 and later.

G.3 September 2013 Scheduled Maintenance

Updates were made to the following sections. The changes are explained below.

- [Section G.3.1, "Configuring and Managing Cluster Resources," on page 465](#)
- [Section G.3.2, "Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes," on page 465](#)
- [Section G.3.3, "Managing Clusters," on page 465](#)

- ♦ [Section G.3.4, “Planning for Novell Cluster Services,” on page 466](#)
- ♦ [Section G.3.5, “What’s New or Changed in Novell Cluster Services,” on page 466](#)

G.3.1 Configuring and Managing Cluster Resources

Location	Change
Section 10.15, “Deleting Cluster Resources, or Disabling Clustering for a Pool, LVM Volume Group, or Service,” on page 209	NSS tools have been modified so that it is not necessary to stop NCS in order to disable the Shareable for Clustering setting. The device should not contain partitions that are used by shared pools, clustered pools, or SBDs.

G.3.2 Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 12.19, “Disabling Clustering for a Pool,” on page 296	NSS tools have been modified so that it is not necessary to stop NCS in order to disable the Shareable for Clustering setting. The device should not contain partitions that are used by shared pools, clustered pools, or SBDs.
Section 12.21, “Deleting a Clustered Pool,” on page 301	

G.3.3 Managing Clusters

Location	Change
Section 9.1.2, “Stopping Novell Cluster Services,” on page 140	Novell Cluster Services is a kernel-space application. As with any kernel-space application, unloading it is a best effort and is not guaranteed to stop the modules. If Novell Cluster Services is busy providing services from the kernel when you attempt to stop it, the process might not stop, and a reboot might be required.
Section 9.11, “Joining a Node to the Cluster (Rejoining a Cluster),” on page 155	While the node is not active in the cluster, any changes in storage objects (such as creating, deleting, or expanding clustered pools) are not updated automatically for the node. After you rejoin the node to the cluster, use the <code>nlvm rescan</code> command to recognize the storage changes.

G.3.4 Planning for Novell Cluster Services

Location	Change
Section 4.4, "Hardware Requirements," on page 53	In addition to Ethernet NICs, Novell Cluster Services supports VLAN on NIC bonding in OES 11 SP1 with the latest patches applied. No modifications to scripts are required. You can use ethx or vlanx interfaces in any combination for nodes in a cluster.

G.3.5 What's New or Changed in Novell Cluster Services

Location	Change
Section 2.1, "What's New (OES 11 SP1 September 2013 Patches)," on page 29	This section is new.

G.4 September 21, 2013

Modified Novell eDirectory 8.8 SP7 links to point to the [NetIQ eDirectory 8.8 SP7 Web site \(http://www.netiq.com/documentation/edir887\)](http://www.netiq.com/documentation/edir887).

Updates were made to the following section. The changes are explained below.

G.4.1 Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes

Location	Change
"Location of Cluster, Server, and Storage Objects" on page 220	If the objects are in different contexts, you might need to cluster migrate the pool cluster resource back to the context where the pool was created in order to modify the pool or volume, or to perform other tasks like setting up Distributed File Services junctions or home directories. You receive an eDirectory error if the operation cannot find the information that it needs in the same context.
Section 12.12, "Renaming a Clustered NSS Pool," on page 252	This section was updated for clarity.
Section 12.12.3, "Renaming a Shared Pool and Its Resource Objects," on page 258	This section is new.
Section 12.12.4, "Renaming a Shared Pool in a DST Cluster Resource," on page 262	This section is new.
Section 12.13, "Renaming a Clustered NSS Volume," on page 264	This section is new.

G.5 September 6, 2013

Updates were made to the following sections. The changes are explained below.

- ♦ [Section G.5.1, “Configuring and Managing Cluster Resources,” on page 467](#)
- ♦ [Section G.5.2, “Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes,” on page 467](#)

G.5.1 Configuring and Managing Cluster Resources

Location	Change
Section 10.1.2, “Using Parameter Values with Spaces in a Cluster Script,” on page 178	In cluster scripts, if a parameter value contains spaces, you must enclose the value with both double-quotes (") and a single-quotes ('), such as "'quoted text'".
Section 10.1.3, “Using Double Quotation Marks in a Cluster Script,” on page 178	This section is new.
Section 10.9, “Configuring Preferred Nodes and Failover Order for a Resource,” on page 198	<p>By default, a new resource is configured to run on all nodes, including future nodes. If its preferred node list is ever changed (directly or indirectly) by administrators, future nodes are not automatically added to the resource's preferred nodes list.</p> <p>IMPORTANT: Ensure that you prepare the node for the services in the resource before you migrate or fail over the resource to it.</p>

G.5.2 Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 12.10, “Adding NFS Export for a Clustered Pool Resource,” on page 247	Before you use the <code>exportfs(8)</code> command in resource scripts, ensure that you install the <code>nfs-kernel-server</code> package on all of the cluster nodes in the resource's preferred nodes list. This contains the support utilities for the kernel <code>nfsd</code> daemon. It is not installed by default. You can use the <i>YaST > Software > Software Management</i> tool to find and install the package.
Section 12.10.1, “Sample NFS Load Script,” on page 247	<p>Add the following command in the load script before the <code>exportfs</code> command:</p> <pre># Start the NFS service exit_on_error rcnfsserver start</pre>

G.6 August 16, 2013

Updates were made to the following section. The changes are explained below.

G.6.1 Planning for Novell Cluster Services

Location	Change
Section 4.8.1, "Shared Devices," on page 67	All devices that contribute space to a clustered pool must be able to fail over with the pool cluster resource. You must use the device exclusively for the clustered pool; do not use space on it for other pools or for Linux volumes. A device must be marked as Shareable for Clustering before you can use it to create or expand a clustered pool.
Section 4.8.4, "Shared RAID Devices," on page 69	Any NSS software RAID0/5 device that is used for a clustered pool must contribute space exclusively to that pool; it cannot be used for other pools. This allows the device to fail over between nodes with the pool cluster resource. Ensure that its component devices are marked as Shareable for Clustering before you use a RAID0/5 device to create or expand a clustered pool

G.7 August 6, 2013

Updates were made to the following section. The changes are explained below.

- ♦ [Section G.7.1, "Configuring Cluster Policies and Priorities," on page 468](#)

G.7.1 Configuring Cluster Policies and Priorities

Location	Change
Section 8.6, "Configuring Cascade Failover Prevention," on page 126	You can create a configuration file in the <code>/etc/modprobe.d</code> file that contains the flag setting to enable and disable the Cascade Failover Prevention function for the server.

G.8 July 30, 2013

Updates were made to the following section. The changes are explained below.

- ♦ [Section G.8.1, "Configuring and Managing Cluster Resources," on page 469](#)
- ♦ [Section G.8.2, "Configuring and Managing Cluster Resources for Shared LVM Volume Groups," on page 469](#)

- ♦ [Section G.8.3, “Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes,” on page 469](#)
- ♦ [Section G.8.4, “Configuring Cluster Policies and Priorities,” on page 469](#)

G.8.1 Configuring and Managing Cluster Resources

Location	Change
Section 10.15, “Deleting Cluster Resources, or Disabling Clustering for a Pool, LVM Volume Group, or Service,” on page 209	This section was revised for clarity.

G.8.2 Configuring and Managing Cluster Resources for Shared LVM Volume Groups

Location	Change
Section 13.11, “Deleting a Clustered LVM Volume (Created in NSSMU or NLVM),” on page 366	This section is new.

G.8.3 Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 12.19, “Disabling Clustering for a Pool,” on page 296	This section is new.
Section 12.20, “Deleting and Re-Creating a Pool Cluster Resource (Disabling and Re-Enabling Clustering for a Pool),” on page 298	This section is new.
Section 12.21, “Deleting a Clustered Pool,” on page 301	This section was updated for clarity and moved to the NSS chapter.

G.8.4 Configuring Cluster Policies and Priorities

Location	Change
Section 8.6.1, “Understanding the Cascade Failover Prevention Quarantine,” on page 126	A resource might be quarantined if it is likely responsible for several consecutive server failures, unrelated to interference from failures of other resources.

G.9 July 11, 2013

Novell Business Continuity Clustering 2.0 for OES 11 SP1 was released in July 2013. For purchasing information, see the [Novell Business Continuity Clustering product page](http://www.novell.com/products/businesscontinuity/) (<http://www.novell.com/products/businesscontinuity/>).

Updates were made to the following sections. The changes are explained below.

- ♦ [Section G.9.1, “Planning for Novell Cluster Services,” on page 470](#)
- ♦ [Section G.9.2, “Upgrading Clusters from OES 2 SP3 to OES 11x,” on page 470](#)

G.9.1 Planning for Novell Cluster Services

Location	Change
Section 4.6.6, “Clusters Plug-in for iManager,” on page 60	Novell Business Continuity Clustering 2.0 for OES 11 SP1 is supported in the Clusters plug-in for OES 11 SP1 with the latest patches applied. See the BCC 2.0: Administration Guide for OES 11 SP1 (http://www.novell.com/documentation/bcc/bcc20_admin_lx/data/bookinfo.html).

G.9.2 Upgrading Clusters from OES 2 SP3 to OES 11x

Location	Change
Section 7.6, “Upgrading a Cluster with DST Resources from OES 2 SP3 to OES 11x,” on page 112	If you replace an OES 2 SP3 node with an OES 11x node instead of using an in-place upgrade to OES 11x, you must manually configure the global policies and move the <i>All Shadowed Volumes</i> policies data to the OES 11x node.

G.10 May 2013 Scheduled Maintenance

Updates were made to the following sections. The changes are explained below.

- ♦ [Section G.10.1, “Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes,” on page 471](#)
- ♦ [Section G.10.2, “Console Commands for Novell Cluster Services,” on page 471](#)
- ♦ [Section G.10.3, “Managing Clusters,” on page 471](#)
- ♦ [Section G.10.4, “Security Considerations,” on page 471](#)
- ♦ [Section G.10.5, “Planning for Novell Cluster Services,” on page 472](#)
- ♦ [Section G.10.6, “What’s New or Changed in Novell Cluster Services,” on page 472](#)

G.10.1 Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes

Location	Change
“Extending the Size of a LUN” on page 286 “Creating and Sharing a New LUN Device” on page 291	WARNING: In EMC PowerPath environments, do not use the <code>rescan-scsi-bus.sh</code> utility provided with the operating system or the HBA vendor scripts for scanning the SCSI buses. To avoid potential file system corruption, EMC requires that you follow the procedure provided in the vendor documentation for EMC PowerPath for Linux.

G.10.2 Console Commands for Novell Cluster Services

Location	Change
“SCAN FOR NEW DEVICES” on page 422	You can use the Linux <code>rescan-scsi-bus.sh</code> utility to scan for new devices. WARNING: In EMC PowerPath environments, do not use the <code>rescan-scsi-bus.sh</code> utility provided with the operating system or the HBA vendor scripts for scanning the SCSI buses. To avoid potential file system corruption, EMC requires that you follow the procedure provided in the vendor documentation for EMC PowerPath for Linux.

G.10.3 Managing Clusters

Location	Change
Section 9.5, “Viewing the Cluster Event Log in iManager,” on page 146	This section is new.
Section 9.6, “Viewing the Cluster Event Log at the Command Line,” on page 149	This section is new.

G.10.4 Security Considerations

Location	Change
Section 17.4, “Log Files,” on page 416	Added links to procedures to view the Event Log.

G.10.5 Planning for Novell Cluster Services

Location	Change
Section 4.6.4, "SLP," on page 58	The SLP refresh interval is now based on the eDirectory advertise-life-time setting.

G.10.6 What's New or Changed in Novell Cluster Services

Location	Change
Section 2.2, "What's New (OES 11 SP1 May 2013 Patches)," on page 29	This section is new.
Section 2.3, "What's New (OES 11 May 2013 Patches)," on page 30	This section is new.

G.11 May 21, 2013

Updates were made to the following sections. The changes are explained below.

- ♦ [Section G.11.1, "Configuring and Managing Cluster Resources," on page 472](#)
- ♦ [Section G.11.2, "Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes," on page 473](#)
- ♦ [Section G.11.3, "What's New or Changed in Novell Cluster Services," on page 473](#)

G.11.1 Configuring and Managing Cluster Resources

Location	Change
Section 10.9, "Configuring Preferred Nodes and Failover Order for a Resource," on page 198	When you change the failover order of nodes in a resource's Preferred Nodes list, the change takes effect immediately when you click <i>OK</i> or <i>Apply</i> if the resource is running. Otherwise, it takes effect the next time the resource is brought online.
The sections noted were moved from this chapter to Chapter 8, "Configuring Cluster Policies, Protocols, and Properties," on page 113	Section 8.6, "Configuring Cascade Failover Prevention," on page 126
	Section 8.7, "Configuring the Cluster Node Reboot Behavior," on page 129
	Section 8.8, "Configuring STONITH," on page 130
	Section 8.9, "Viewing or Modifying the NCS Proxy User Assignment in the NCS_Management_Group," on page 132
	Section 8.11, "Moving a Cluster, or Changing the Node IP Addresses, LDAP Servers, or Administrator Credentials for a Cluster," on page 134

G.11.2 Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 12.14, "Renaming the Mount Point Path for a Clustered NSS Volume," on page 275	This section is new.

G.11.3 What's New or Changed in Novell Cluster Services

Location	Change
Section 2.4, "What's New (OES 11 SP1 April 2013 Patches)," on page 30	This section is new.

G.12 May 7, 2013

Updates were made to the following section. The changes are explained below.

G.12.1 Troubleshooting Novell Cluster Services

Location	Change
"Migration Authentication Error Occurs when Connecting to a New Cluster Resource on the Target Server" on page 412	This section is new.

G.13 April 2013 Scheduled Maintenance

Updates were made to the following sections. The changes are explained below.

- [Section G.13.1, "Managing Clusters," on page 473](#)
- [Section G.13.2, "Troubleshooting Novell Cluster Services," on page 474](#)
- [Section G.13.3, "What's New or Changed in Novell Cluster Services," on page 474](#)

G.13.1 Managing Clusters

Location	Change
Section 8.6, "Configuring Cascade Failover Prevention," on page 126	This section is new. Cascading failover prevention is supported in OES 11 and later.

G.13.2 Troubleshooting Novell Cluster Services

Location	Change
"Cannot Find the Prevent Cascading Failover Option"	This section is was removed. Cascading failover is supported in OES 11 and later.

G.13.3 What's New or Changed in Novell Cluster Services

Location	Change
Section 2.5, "What's New (OES 11 April 2013 Patches)," on page 31	This section is new.
"Cascading Failover Prevention" in Section 2.11, "What's New (OES 11)," on page 37	This section is new.

G.14 April 18, 2013

Updates were made to the following sections. The changes are explained below.

- [Section G.14.1, "Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes," on page 474](#)
- [Section G.14.2, "Troubleshooting Novell Cluster Services," on page 475](#)
- [Section G.14.3, "Upgrading Clusters from OES 2 SP3 to OES 11x," on page 475](#)
- [Section G.14.4, "What's New or Changed in Novell Cluster Services," on page 475](#)

G.14.1 Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes

Location	Change
"Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes" on page 217	<p>All sample load scripts and unload scripts for NSS pool cluster resources were updated to reflect the new default sequence of commands for OES 11 SP1. Newly created pool cluster resources use the following sequence for commands in a load script. The order is reversed in an unload script.</p> <p>NSS [NFS] (If NFS is used, you manually add the command here.) Secondary IP address NCP [AFP] [CIFS]</p>

G.14.2 Troubleshooting Novell Cluster Services

Location	Change
“Error NFS Stale Handle Reported on Host when Cluster Failover Occurs” on page 410	This section is new.

G.14.3 Upgrading Clusters from OES 2 SP3 to OES 11x

Location	Change
“Upgrading Clusters from OES 2 SP3 to OES 11x” on page 105	<p>IMPORTANT: You can upgrade a cluster from OES 2 SP3 to OES 11 or from OES 2 SP3 to OES 11 SP1. The same release version of OES 11 must be installed and running on each new or upgraded node in the cluster.</p> <p>The procedures in this section refer to OES 11x. A specific version of OES 11 is mentioned only if the procedure applies only to that version of the platform.</p>
“Adding OES 11x Nodes to an OES 2 SP3 Cluster (Rolling Cluster Upgrade)” on page 108	<p>During a rolling cluster upgrade, one OES 11x server is added at a time, either by adding a new node or upgrading an OES 2 SP3 node:</p> <ul style="list-style-type: none">♦ Add an OES 11x node and migrate resources to it.♦ Remove (cluster leave command) an OES 2 SP3 node, upgrade it to OES 11x, return it to the cluster, and migrate resources to it.

G.14.4 What’s New or Changed in Novell Cluster Services

Location	Change
“Change in the Command Sequence for Pool Cluster Resource Scripts” in Section 2.10, “What’s New (OES 11 SP1),” on page 33.	This section is new.

G.15 April 12, 2013

Updates were made to the following sections. The changes are explained below.

- ♦ [Section G.15.1, “Configuring and Managing Cluster Resources for Shared LVM Volume Groups,” on page 476](#)
- ♦ [Section G.15.2, “Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes,” on page 476](#)

G.15.1 Configuring and Managing Cluster Resources for Shared LVM Volume Groups

Location	Change
“Configuring an LVM Volume Group Cluster Resource with NSS Management Tools” on page 312	<p>Ensure that you log in to the master node to create clustered LVM volume groups with NSSMU and NLVM commands.</p> <p>Typically, 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.</p> <p>When an LVM logical volume group is clustered, CLVM manages the share state, not the device. NSSMU reports the device as Not Shareable for Clustering. NSSMU reports the clustered volume is Shareable for Clustering.</p>
“Creating a Generic File System Cluster Resource for an LVM Volume Group” on page 335	<p>Ensure that the SAN device is attached to all of the nodes in the cluster.</p> <p>Ensure that the existing LVM volume you want to cluster-enable is deactive on all nodes in the cluster.</p> <p>Be mindful that Linux path names are case sensitive.</p> <p>The mount path must already exist.</p> <p>If the resource goes comatose, typical issues are that the target LVM volume is still active locally on a node, the path names were not entered as case-sensitive, the IP address is not unique in the network, or the mount path does not exist.</p>

G.15.2 Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes

Location	Change
“Creating a Cluster-Enabled Pool and Volume with iManager” on page 227	<p>Ensure that the SAN device is attached to all of the nodes in the cluster.</p>
“Creating a Cluster-Enabled Pool and Volume with NSSMU” on page 230	<p>Ensure that you log in to the master node to create clustered pools and volumes.</p> <p>Typically, the pool creation takes less than a minute. 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.</p>

Location	Change
Chapter 12, "Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes," on page 217	In the <code>ncpcon mount</code> command, place the <code>/opt</code> switch before the volume information. <code>ncpcon mount /opt=ns=long <VOLUMENAME>=<VOLUMEID></code>

G.16 January 2013 Scheduled Maintenance

Updates were made to the following sections. The changes are explained below.

- [Section G.16.1, "Clusters Plug-In Changes for Novell iManager 2.7.5," on page 477](#)
- [Section G.16.2, "Configuring and Managing Cluster Resources," on page 477](#)
- [Section G.16.3, "Upgrading Clusters from OES 2 SP3 to OES 11 SP1," on page 477](#)
- [Section G.16.4, "What's New or Changed in Novell Cluster Services," on page 477](#)

G.16.1 Clusters Plug-In Changes for Novell iManager 2.7.5

The [Appendix F, "Clusters Plug-In Changes for Novell iManager 2.7.5," on page 457](#) section is new. It describes the changes to the Clusters plug-in for Novell iManager 2.7.5, and compares how to perform common tasks in the old and new versions of the plug-in.

G.16.2 Configuring and Managing Cluster Resources

Location	Change
Step 8 in Section 10.9, "Configuring Preferred Nodes and Failover Order for a Resource," on page 198	If you remove nodes from the Assigned list in the Edit view, the removed nodes are automatically added to the Unassigned list when you save.

G.16.3 Upgrading Clusters from OES 2 SP3 to OES 11 SP1

Location	Change
Section 7.6, "Upgrading a Cluster with DST Resources from OES 2 SP3 to OES 11x," on page 112	This section is new.

G.16.4 What's New or Changed in Novell Cluster Services

Location	Change
"What's New (OES 11 SP1 January 2013 Patches)" on page 31	This section is new.

Location	Change
“What’s New (OES 11 January 2013 Patches)” on page 32	This section is new.

G.17 November 2012 Scheduled Maintenance

Updates were made to the following sections. The changes are explained below.

- [Section G.17.1, “Configuring Cluster Policies, Protocols, and Properties,” on page 478](#)
- [Section G.17.2, “Configuring and Managing Cluster Resources,” on page 478](#)
- [Section G.17.3, “Configuring and Managing Cluster Resources for Shared LVM Volume Groups,” on page 478](#)
- [Section G.17.4, “Managing Clusters,” on page 479](#)
- [Section G.17.5, “Upgrading OES 11 Clusters,” on page 479](#)
- [Section G.17.6, “What’s New or Changed in Novell Cluster Services,” on page 479](#)

G.17.1 Configuring Cluster Policies, Protocols, and Properties

Location	Change
“Personalizing the My Clusters Display” on page 119	If scrolling is needed to view the complete list, you can use the <i>Refresh</i> option to set a longer refresh rate for the page to allow enough time to view the status for all items.

G.17.2 Configuring and Managing Cluster Resources

Location	Change
“Personalizing the My Resources Display” on page 182	If scrolling is needed to view the complete list, you can use the <i>Refresh</i> option to set a longer refresh rate for the page to allow enough time to view the status for all items.

G.17.3 Configuring and Managing Cluster Resources for Shared LVM Volume Groups

Location	Change
“Configuring an LVM Volume Group Cluster Resource with NSS Management Tools” on page 312	If you enable NCP for the volume, this name must comply with the “Naming Conventions for NCP Volumes” on page 307 . Lowercase letters are automatically changed to uppercase for the NCP volume name.

G.17.4 Managing Clusters

Location	Change
“Monitoring Cluster and Resource States” on page 142	If scrolling is needed to view the complete list, you can use the <i>Refresh</i> option to set a longer refresh rate for the page to allow enough time to view the status for all items.

G.17.5 Upgrading OES 11 Clusters

Location	Change
“Upgrade Issues for OES 11 SP1” on page 104	This section is new. The Section 6.3.1, “Updating the Clusters Plug-in and Role-Based Services in Novell iManager,” on page 104 was moved to this section.
“DHCP Cluster Resource: Path Change for dhcpd.pid” on page 104	This issue is new.

G.17.6 What’s New or Changed in Novell Cluster Services

Location	Change
“What’s New (OES 11 SP1 November 2012 Patches)” on page 33	This section is new.
“DHCP Cluster Resource: Path Change for dhcpd.pid” on page 36	This section is new.

G.18 September 2012 Scheduled Maintenance

Updates were made to the following sections. The changes are explained below.

- ♦ [Section G.18.1, “Configuring and Managing Cluster Resources,” on page 479](#)
- ♦ [Section G.18.2, “Planning for Novell Cluster Services,” on page 480](#)
- ♦ [Section G.18.3, “Troubleshooting Novell Cluster Services,” on page 480](#)
- ♦ [Section G.18.4, “What’s New or Changed in Novell Cluster Services,” on page 480](#)

G.18.1 Configuring and Managing Cluster Resources

Location	Change
Section 10.2.5, “Updating the My Resources List after Renaming or Deleting a Resource,” on page 184	This section is new.

G.18.2 Planning for Novell Cluster Services

Location	Change
Section 4.8.2, "SBD Partitions," on page 67	NLVM and other NSS management tools need the SBD to detect if a node is a member of the cluster and to get exclusive locks on physically shared storage. Because an SBD partition ends on a cylinder boundary, the partition size might be slightly smaller than the size you specify. When you use an entire device for the SBD partition, you can use the 'max' option as the size, and let the software determine the size of the partition.

G.18.3 Troubleshooting Novell Cluster Services

Location	Change
Section 16.18, "Error 20897 - This node is not a cluster member," on page 411	This section is new.
Section 16.8, "iManager: Problems with Clusters Plug-In After Update for Novell iManager 2.7.5," on page 408	This section is new.

G.18.4 What's New or Changed in Novell Cluster Services

Location	Change
"What's New (OES 11 September 2012 Patches)" on page 33	This section is new.

G.19 September 12, 2012

Updates were made to the following sections. The changes are explained below.

- ♦ [Section G.19.1, "Configuring and Managing Cluster Resources," on page 481](#)
- ♦ [Section G.19.2, "Configuring Cluster Policies, Protocols, and Priorities," on page 481](#)
- ♦ [Section G.19.3, "Installing and Configuring Novell Cluster Services," on page 481](#)
- ♦ [Section G.19.4, "Troubleshooting Novell Cluster Services," on page 481](#)

G.19.1 Configuring and Managing Cluster Resources

Location	Change
Section 10.2, "Setting Up a Personalized List of Resources to Manage," on page 179	Your personalized My Resources list and display settings are saved on the iManager server.

G.19.2 Configuring Cluster Policies, Protocols, and Priorities

Location	Change
Section 8.2, "Setting Up a Personalized List of Clusters to Manage," on page 117	Your personalized My Clusters list and display settings are saved on the iManager server.

G.19.3 Installing and Configuring Novell Cluster Services

Location	Change
Section 5.8.2, "Installing the Clusters Plug-In," on page 97 Step 5 in Section 5.8.3, "Uninstalling and Reinstalling the Clusters Plug-In," on page 97	Verify that the <code>/var/opt/novell/iManager/nps/WEB-INF/lib</code> directory contains the class loader file <code>commons-lang-2.6.jar</code> (or later version), then remove all earlier versions of the file. You must remove the earlier versions (such as <code>commons-lang.jar</code> or <code>commons-lang-2.4.jar</code>) to ensure that iManager loads the 2.6 or later version of the class loader file.

G.19.4 Troubleshooting Novell Cluster Services

Location	Change
Section 16.7, "iManager: Error 500 Occurs after Installing or Updating the Clusters Plug-In for iManager 2.7.5," on page 408	This section is new.

G.20 August 28, 2012 (OES 11 SP1)

Procedures throughout the guide were updated to use the newly reorganized Clusters plug-in for Novell iManager. In addition, updates were made to the following sections. The changes are explained below.

- ♦ [Section G.20.1, "Console Commands for Novell Cluster Services," on page 482](#)
- ♦ [Section G.20.2, "Comparing Novell Cluster Services for Linux and NetWare," on page 482](#)
- ♦ [Section G.20.3, "Configuring and Managing Cluster Resources," on page 482](#)
- ♦ [Section G.20.4, "Configuring and Managing Cluster Resources for Shared Linux LVM Volume Groups," on page 483](#)

- ♦ [Section G.20.5, “Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes,” on page 483](#)
- ♦ [Section G.20.6, “Configuring Cluster Policies, Protocols, and Priorities,” on page 484](#)
- ♦ [Section G.20.7, “Configuring Novell Cluster Services in a Virtualization Environment,” on page 484](#)
- ♦ [Section G.20.8, “Electing a Master Node,” on page 484](#)
- ♦ [Section G.20.9, “Files for Novell Cluster Services,” on page 485](#)
- ♦ [Section G.20.10, “Installing and Configuring Novell Cluster Services on OES 11,” on page 485](#)
- ♦ [Section G.20.11, “Managing Clusters,” on page 485](#)
- ♦ [Section G.20.12, “Planning for Novell Cluster Services,” on page 485](#)
- ♦ [Section G.20.13, “Security Considerations,” on page 486](#)
- ♦ [Section G.20.14, “Troubleshooting Novell Cluster Services,” on page 486](#)
- ♦ [Section G.20.15, “Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers,” on page 486](#)
- ♦ [Section G.20.16, “Upgrading Clusters from OES 2 SP3 to OES 11 SP1,” on page 486](#)
- ♦ [Section G.20.17, “Upgrading OES 11 Clusters,” on page 487](#)
- ♦ [Section G.20.18, “What’s New or Changed in Novell Cluster Services,” on page 487](#)

G.20.1 Console Commands for Novell Cluster Services

Location	Change
“NCPCON Mount Command for NCP and NSS Volumes” on page 435	This section is new.

G.20.2 Comparing Novell Cluster Services for Linux and NetWare

Updated for new or changed OES 11 SP1 features.

G.20.3 Configuring and Managing Cluster Resources

Location	Change
Section 10.1.2, “Using Parameter Values with Spaces in a Cluster Script,” on page 178	This section is new.
Section 10.2, “Setting Up a Personalized List of Resources to Manage,” on page 179	This section is new.
Section 10.13, “Changing the IP Address of a Cluster Resource,” on page 206	This procedures in this section are new.

G.20.4 Configuring and Managing Cluster Resources for Shared Linux LVM Volume Groups

Location	Change
“NCP File Access with Novell NCP Server” on page 307	For information about using antivirus software with NCP volumes, see “ Security ” in the <i>OES 11 SP1: Planning and Implementation Guide</i> .
Section 13.3, “Configuring an LVM Volume Group Cluster Resource with NSS Management Tools,” on page 312	This section was expanded to include nlvm commands and NCP enabling the clustered LVM logical volume.
Section 13.6, “Enabling NCP File Access for a Clustered LVM Volume,” on page 348	This section is new.
Section 13.7, “Modifying the LVM Resource Scripts for Other Services,” on page 355	This section is new.

G.20.5 Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes

Location	Change
“Volume” on page 219	For information about using antivirus software with NSS volumes, see “ Antivirus Support for NSS ” in the <i>OES 11 SP1: NSS File System Administration Guide for Linux</i> and “ Security ” in the <i>OES 11 SP1: Planning and Implementation Guide</i> .
“NCP Server for Linux” on page 220	NCP Server does not support cross-protocol locks across a cluster migration or failover of the resource. Hyphens are used in the default NCP virtual server name. Underscores can be used in customized names.
“Novell CIFS for Linux” on page 221	The default name suggested for the CIFS virtual server name is based on the NCP virtual server name.
“Initializing and Sharing a Device with NLVM” on page 225	This section is new.
“Cluster-Enabling an Existing NSS Pool and Its Volumes” on page 233	The default name suggested for the CIFS virtual server name is based on the NCP virtual server name.
“Configuring a Load Script for the Shared NSS Pool” on page 240	To customize the mount point for a cluster enabled volume, add the <code>path</code> option to the <code>ncpcon mount</code> command.
“Configuring a Monitor Script for the Shared NSS Pool” on page 244	If AFP is used as an advertising protocol, the <code>exit_on_error afpstat</code> command is added to the monitor script.
“Adding NFS Export for a Clustered Pool Resource” on page 247	This section is new.

Location	Change
Section 12.12, "Renaming a Clustered NSS Pool," on page 252	Because renaming involves changing information for the Pool object in eDirectory, you must ensure that the shared pool is active on a cluster node that has its NCP server object in the same context as the Pool object of the pool you are going to rename. Added details for the renaming procedure.
"Changing the Name Space for a Clustered NSS Volume" on page 281	This section is new
"Expanding the Size of a Clustered Pool" on page 284	Corrected the procedures for an online pool expansion.

G.20.6 Configuring Cluster Policies, Protocols, and Priorities

Location	Change
Section 8.4, "Configuring Cluster Protocols," on page 122	Modify the Master Watchdog and Slave Watchdog parameter settings only when supervised by Novell Technical Support (http://www.novell.com/support) .
Section 8.5, "Configuring Cluster Event Email Notification," on page 124	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.
Section 8.10, "Viewing or Modifying the Cluster IP Address or Port," on page 133	A cluster restart is required before iManager can continue to manage the cluster.

G.20.7 Configuring Novell Cluster Services in a Virtualization Environment

Location	Change
Chapter 15, "Configuring Novell Cluster Services in a Virtualization Environment," on page 391	To get started with KVM virtualization, see the Virtualization with KVM documentation (http://www.suse.com/documentation/sles11/book_kvm/data/book_kvm.html) .

G.20.8 Electing a Master Node

The section [Appendix C, "Electing a Master Node," on page 441](#) is new.

G.20.9 Files for Novell Cluster Services

Location	Change
Table B-1, "Location and Purpose of Novell Cluster Services Files," on page 437	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.
"/opt/novell/ncs/bin/NCS_STONITH_SCRIPT file" on page 438	You create and code the /opt/novell/ncs/bin/NCS_STONITH_SCRIPT script file to enable and control the STONITH feature.

G.20.10 Installing and Configuring Novell Cluster Services on OES 11

Location	Change
Section 5.8, "Installing or Updating the Clusters Plug-in for iManager," on page 96	Expanded to explain the updates for OES 11 SP1.

G.20.11 Managing Clusters

Location	Change
Section 9.3, "Monitoring Cluster and Resource States," on page 142	Added information about the <code>Resource not found</code> condition for resources on the My Resources page of the Clusters plug-in.
Section 9.4, "Responding to a Start, Failover, or Failback Alert," on page 145	This section is new.
Table 9-2, "Cluster Resource States," on page 144	Updated the Alert entry to identify alert types: Start alert, Failover alert, and Failback alert.
Section 8.8, "Configuring STONITH," on page 130	This section is new.

G.20.12 Planning for Novell Cluster Services

Location	Change
Section 4.3, "Volume ID Requirements," on page 52	Volume IDs for BCC-enabled clustered volumes must be unique across all nodes in every peer cluster.
"NCP Server for Linux" on page 63	NCP Server does not support cross-protocol locks across a cluster migration or failover of the resource.

Location	Change
Section 4.10.3, "Modifying the Polling Interval, No Path Retry, and Failback Settings in the multipath.conf File," on page 71	Added information about the <code>polling_interval</code> .
Section 4.6.8, "SFCB and CIMOM," on page 61	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.

G.20.13 Security Considerations

Location	Change
"Ports" on page 415	Email Notification uses Postfix, port 25.
"Log Files" on page 416	NCS configuration writes log messages to <code>/var/log/YaST2/y2log</code> .
"Configuration Files" on page 416	This section is new.

G.20.14 Troubleshooting Novell Cluster Services

Location	Change
"iManager: Error 500 Occurs while Connecting to a Cluster" on page 408	This section is new.

G.20.15 Upgrading and Managing Cluster Resources for Linux POSIX Volumes with CSM Containers

Location	Change
"Configuring and Adding OES 11 SP1 Nodes to the OES 2 SP3 Cluster" on page 385	The path for the NLVM configuration file is <code>/etc/opt/novell/nss/nlvm.conf</code> .

G.20.16 Upgrading Clusters from OES 2 SP3 to OES 11 SP1

Location	Change
"Supported Upgrade Paths from OES 2 SP3 to OES 11x" on page 107	Upgrade is supported from OES 2 SP3 on SLES 10 SP3 or SLES 11 SP4 to OES 11 SP1 on SLES 11 SP2.
Section 7.3, "Requirements and Guidelines for Upgrading Clusters from OES 2 SP3," on page 108	Upgrade OES 2 SP3 cluster nodes with the higher IP addresses first.
Multiple locations	Procedures have been updated for OES 11 SP1.

G.20.17 Upgrading OES 11 Clusters

This section is new.

G.20.18 What's New or Changed in Novell Cluster Services

Location	Change
"What's New (OES 11 SP1)" on page 33	This section is new.

