

Novell Open Enterprise Server

www.novell.com

April, 2007

NOVELL CLUSTER SERVICES™ 1.8.2
ADMINISTRATION GUIDE FOR LINUX*



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. Please refer to 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 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

Novell, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.novell.com/company/legal/patents/> and one or more additional patents or pending patent applications in the U.S. and in other countries.

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

Online Documentation: To access the online documentation for this and other Novell products, and to get updates, see www.novell.com/documentation.

Novell Trademarks

ConsoleOne is a registered trademark of Novell, Inc. in the United States and other countries.

eDirectory is a trademark of Novell, Inc.

GroupWise is a registered trademark of Novell, Inc. in the United States and other countries.

NetWare is a registered trademark of Novell, Inc. in the United States and other countries.

NetWare Core Protocol and NCP are trademarks of Novell, Inc.

Novell is a registered trademark of Novell, Inc. in the United States and other countries.

Novell Authorized Reseller is a Service mark of Novell Inc.

Novell Cluster Services is a trademark of Novell, Inc.

Novell Directory Services and NDS are registered trademarks of Novell, Inc. in the United States and other countries.

Novell Storage Services is a trademark of Novell, Inc.

SUSE is a registered trademark of Novell, Inc. in the United States and other countries.

Third-Party Materials

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

Contents

About This Guide	7
1 Overview	9
1.1 Product Features	9
1.2 Product Benefits	9
1.3 Cluster Configuration	11
1.3.1 Cluster Components	13
2 What's New	15
3 Installation and Setup	17
3.1 Hardware Requirements	17
3.2 Software Requirements	17
3.3 Shared Disk System Requirements	18
3.4 Rules for Operating a Novell Cluster Services SAN	18
3.5 Installing Novell Cluster Services	18
3.5.1 Novell Cluster Services Licensing	19
3.5.2 Installing Novell Cluster Services during the OES Installation	19
3.5.3 Installing Novell Cluster Services after the OES Installation	20
3.5.4 Starting and Stopping Novell Cluster Services	21
3.6 Converting a NetWare Cluster to Linux	21
3.6.1 Changing Existing NetWare Cluster Nodes to Linux (Rolling Cluster Conversion)	21
3.6.2 Adding New Linux Nodes to Your NetWare Cluster	23
3.6.3 Mixed NetWare and Linux Clusters	24
3.6.4 Finalizing the Cluster Conversion	26
3.7 Setting Up Novell Cluster Services	27
3.7.1 Creating NSS Shared Disk Partitions and Pools	27
3.7.2 Creating NSS Volumes	29
3.7.3 Cluster Enabling NSS Pools and Volumes	31
3.7.4 Creating Traditional Linux Volumes on Shared Disks	32
3.7.5 Expanding EVMS Volumes on Shared Disks	35
3.7.6 Cluster Enabling Traditional Linux Volumes on Shared Disks	36
3.7.7 Creating Cluster Resource Templates	40
3.7.8 Creating Cluster Resources	41
3.7.9 Configuring Load Scripts	42
3.7.10 Configuring Unload Scripts	42
3.7.11 Setting Start, Failover, and Failback Modes	42
3.7.12 Assigning Nodes to a Resource	43
3.8 Configuration Settings	44
3.8.1 Editing Quorum Membership and Timeout Properties	44
3.8.2 Cluster Protocol Properties	45
3.8.3 Cluster IP Address and Port Properties	45
3.8.4 Resource Priority	46
3.8.5 Cluster E-Mail Notification	46
3.8.6 Cluster Node Properties	47
3.9 Additional Information	48

4	Managing Novell Cluster Services	49
4.1	Migrating Resources	49
4.2	Identifying Cluster and Resource States	49
4.3	Novell Cluster Services Console Commands	50
4.4	Customizing Cluster Services Management	54
4.5	Novell Cluster Services File Locations	55
4.6	Additional Cluster Operating Instructions	57
4.6.1	Connecting to an iSCSI Target	57
4.6.2	Adding a Node That Was Previously in the Cluster	58
4.6.3	Cluster Maintenance Mode	58
4.6.4	Shutting Down Linux Servers When Servicing Shared Storage	58
4.6.5	Preventing Cluster Node Reboot after Node Shutdown.	58
4.6.6	Problems Authenticating to Remote Servers during Cluster Configuration	59
4.6.7	Reconfiguring a Cluster Node	59
4.6.8	Device Name Required to Create a Cluster Partition.	59
4.6.9	Creating a Cluster Partition (SBD Partition) after Installation.	59
4.6.10	Mirroring SBD (Cluster) Partitions	60
A	Documentation Updates	61
A.1	December 23, 2005 (Open Enterprise Server SP2)	61

About This Guide

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

- ♦ Chapter 1, “Overview,” on page 9
- ♦ Chapter 2, “What’s New,” on page 15
- ♦ Chapter 3, “Installation and Setup,” on page 17
- ♦ Chapter 4, “Managing Novell Cluster Services,” on page 49
- ♦ Appendix A, “Documentation Updates,” on page 61

Audience

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

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comments feature at the bottom of each page of the online documentation, or go to www.novell.com/documentation/feedback.html and enter your comments there.

Documentation Updates

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

Documentation Conventions

In Novell documentation, a greater-than symbol (>) is used to separate actions within a step and items in a cross-reference path.

A trademark symbol (®, ™, etc.) denotes a Novell trademark. An asterisk (*) denotes a third-party trademark.

Overview

1

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

1.1 Product Features

Novell Cluster Services includes several important features to help you ensure and manage the availability of your network resources. These include:

- ◆ Support for shared SCSI, iSCSI or fibre channel storage area networks.
- ◆ Multinode all-active cluster (up to 32 nodes). Any server in the cluster can restart resources (applications, services, IP addresses, and file systems) from a failed server in the cluster.
- ◆ A single point of administration through the browser-based Novell iManager cluster configuration and monitoring GUI. iManager also lets you remotely manage your cluster.
- ◆ The ability to tailor a cluster to the specific applications and hardware infrastructure that fit your organization.
- ◆ Dynamic assignment and reassignment of server storage on an as-needed basis.
- ◆ The ability to automatically notify administrators through e-mail of cluster events and cluster state changes.

1.2 Product Benefits

Novell Cluster Services allows you to configure up to 32 Linux servers into a high-availability cluster, where resources can be dynamically switched or moved to any server in the cluster. Resources can be configured to automatically switch or be moved in the event of a server failure, or they can be moved manually to troubleshoot hardware or balance the workload.

Novell Cluster Services provides high availability from commodity components. Lower costs are obtained through the consolidation of applications and operations onto a cluster. The ability 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) are also important benefits of Novell Cluster Services.

An equally important benefit of implementing Novell Cluster Services is that you can reduce unplanned service outages as well as planned outages for software and hardware maintenance and upgrades.

Reasons you would want to implement Novell Cluster Services include the following:

- ◆ Increased availability
- ◆ Improved performance
- ◆ Low cost of operation
- ◆ Scalability
- ◆ Disaster recovery


- ◆ Data protection
- ◆ Server Consolidation
- ◆ Storage Consolidation

Shared disk fault tolerance can be obtained by implementing RAID on the shared disk subsystem.

An example of the benefits 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 subsystem connected to each of the servers in the cluster. The following figure 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 depending on Web Server 1 for Internet access, e-mail, and information lose their connections. The following figure 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 should a failure occur. In the previous example, 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

- ◆ 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 having to log in again.

Now suppose the problems with Web Server 1 are resolved, and Web Server 1 is returned to a normal operating state. Web Site A and Web Site B will automatically fail back, or be moved back to Web Server 1, and Web Server operation will return back to the way it was before Web Server 1 failed.

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


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

1.3 Cluster Configuration

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.


Typical resources might include data, applications, and services. The following figure shows how a typical fibre channel cluster configuration might look.

Figure 1-3 *Typical Fibre Channel Cluster Configuration*




Although fibre channel provides the best performance, you can also configure your cluster to use shared SCSI or iSCSI. The following figure shows how a typical shared SCSI cluster configuration might look.

Figure 1-4 *Typical Shared SCSI Cluster Configuration*



iSCSI is another option and is an alternative to fibre channel that can be used to create a low-cost SAN. The following figure shows how a typical iSCSI cluster configuration might look.

Figure 1-5 Typical iSCSI Cluster Configuration



1.3.1 Cluster Components

The following components make up a Novell Cluster Services cluster:

- ◆ From 2 to 32 Linux servers, each containing at least one local disk device.
- ◆ Novell Cluster Services software running on each Linux server in the cluster.
- ◆ A shared disk subsystem connected to all servers in the cluster (optional, but recommended for most configurations).
- ◆ High-speed fibre channel cards, cables, and switch or SCSI cards and cables used to connect the servers to the shared disk subsystem.

What's New

2

The following changes and enhancements were added to Novell® Cluster Services™ for Linux for Novell Open Enterprise Server (OES) Support Pack 2.

- ◆ It is now possible to choose a device for the SBD partition from a list rather than entering it manually. See [Section 3.5, “Installing Novell Cluster Services,”](#) on page 18.
- ◆ Some iManager cluster option names and locations have changed to make cluster configuration and management easier.
- ◆ It is now possible to upgrade a cluster node directly from NetWare 6.0 to OES Linux without first upgrading to NetWare 6.5. See [Section 3.6, “Converting a NetWare Cluster to Linux,”](#) on page 21.

Novell® Cluster Services™ can be installed during the Open Enterprise Server (OES) installation or after. OES is now part of the SUSE® Linux Enterprise Server (SLES) 9 installation. During the Novell Cluster Services part of the OES installation, you are prompted for configuration information that is necessary for Novell Cluster Services to function properly. This chapter contains information to help you install, set up, and configure Novell Cluster Services.

3.1 Hardware Requirements

The following list specifies hardware requirements for installing Novell Cluster Services. These requirements 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
- At least 512 MB of memory on each server in the cluster.

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

If you have a fibre channel SAN, the host bus adapters (HBAs) for each cluster node should be identical.

3.2 Software Requirements

Novell Cluster Services is installed as part of the OES installation. OES must be installed and running on each cluster server. In addition to OES, ensure that the following requirements are met:

- All servers in the cluster are configured with a static IP address and are on the same IP subnet
- There is an additional IP address for the cluster and for each cluster resource and cluster-enabled pool
- All servers in the cluster are in the same Novell eDirectory™ tree

If the servers in the cluster are in separate eDirectory containers, each server must have rights to the other server's containers and to the containers where any cluster enabled pool objects are stored. You can do this by adding trustee assignments for all cluster servers to a parent container of the containers where the cluster server objects reside. See [eDirectory Rights \(http://www.novell.com/documentation/edir873/edir873/data/fbachifb.html#fbachifb\)](http://www.novell.com/documentation/edir873/edir873/data/fbachifb.html#fbachifb) in the *eDirectory 8.7.3 Administration Guide* for more information.

- The browser that will be used to manage Novell Cluster Services is set to a supported language.

The iManager plug-in for Novell Cluster Services might not operate properly if the highest priority Language setting for your Web browser is set to a language other than one of the supported languages. To avoid problems, in your Web browser, click *Tools > Options > Languages*, and then set the first language preference in the list to a supported language.

3.3 Shared Disk System Requirements

A shared disk system (Storage Area Network, or SAN) is required for each cluster if you want data to be highly available. If a shared disk subsystem is used, ensure the following:

- ❑ At least 20 MB of free disk space on the shared disk system for creating a special cluster partition

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

- ❑ The shared disk system is properly set up and functional according to the manufacturer's instructions before installing Novell Cluster Services.
- ❑ We recommend that the disks contained in the shared disk system are configured to use mirroring or RAID to add fault tolerance to the shared disk system.
- ❑ If you are using iSCSI for shared disk system access, ensure you have configured iSCSI initiators and targets prior to installing Novell Cluster Services. See [Accessing iSCSI Targets on NetWare Servers from Linux Initiators \(http://www.novell.com/documentation/iscsi1_nak/iscsi/data/bswmaoa.html#bt8cyhf\)](http://www.novell.com/documentation/iscsi1_nak/iscsi/data/bswmaoa.html#bt8cyhf) for more information.

3.4 Rules for Operating a Novell Cluster Services SAN

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 disks, whether in the cluster or not, have access to all of the data on the shared storage space unless you specifically prevent such access. Novell Cluster Services arbitrates access to shared data for all cluster nodes, but cannot protect shared data from being corrupted by noncluster servers.

3.5 Installing Novell Cluster Services

It is necessary to install SLES 9/OES on every server you want to add to a cluster. You can install Novell Cluster Services and create a new cluster, or add a server to an existing cluster either during the SLES 9/OES installation or afterwards, using YaST.

If you are creating a new cluster, the YaST setup tool

- ◆ Creates a new Cluster object and Cluster Node object in eDirectory.
- ◆ Installs Novell Cluster Services software on the server.
- ◆ Creates a special cluster partition if you have a shared disk system.

If you are adding a server to an existing cluster, the YaST setup tool

- ◆ Creates a new Cluster Node object in eDirectory.
- ◆ Installs Novell Cluster Services software on the server.

You can install up to 32 nodes in each cluster.

3.5.1 Novell Cluster Services Licensing

You can add up to 32 nodes to a cluster. Novell Cluster Services for Linux includes licenses for two cluster nodes. You only need additional Cluster Server Licenses if you have a three-node or larger cluster. A paper license for additional cluster nodes can be obtained from Novell or from your Novell Authorized ResellerSM.

3.5.2 Installing Novell Cluster Services during the OES Installation

- 1 Start the SUSE Linux Enterprise Server 9 (SLES 9) installation and continue until you get to the Installation Settings screen, then click *Software*.

OES is part of the SLES 9 install.

The SLES 9/OES installation includes several steps not described here because they do not directly relate to Novell Cluster Services. For more detailed instructions on installing OES with SLES 9, see the *OES Linux Installation Guide*.

- 2 On the Software Selection screen, click *Detailed Selection*.
- 3 In the Selection window, click *Novell Cluster Services* and any other OES components that you want to install, then click *Accept*.

NSS is a required component for Novell Cluster Services and it is automatically selected when you select Novell Cluster Services. Installing NSS also allows you to create cluster-enabled NSS pools (virtual servers).

iManager is required to configure and manage Novell Cluster Services, and must be installed on at least one server.

- 4 Continue through the installation process until you reach the Installation Settings screen, then click the *Cluster Services* link.
 - 5 Choose whether eDirectory is installed locally or remotely, accept or change the Admin name and enter the Admin password, then click *Next*.
- eDirectory is automatically selected when NSS is selected.
- 6 Choose to either create a new cluster, configure Novell Cluster Services on a server that you will add to an existing cluster, or configure Novell Cluster Services later.
 - 7 Enter the fully distinguished name (FDN) of the cluster.

IMPORTANT: Use the dot format illustrated in the example. Do not use commas.

If you are creating a new cluster, this is the name you will give the new cluster and the eDirectory context where the new Cluster object will reside. You must specify an existing context. Specifying a new context does not create a new context.

If you are adding a server to an existing cluster, this is the name and eDirectory context of the cluster that you are adding this server to.

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

- 8 (Conditional) If you are creating a new cluster, enter a unique IP address for the cluster.

The cluster IP address is separate from the server IP address, is required to be on the same IP subnet as the other cluster servers, and is required for certain external network management

programs to get cluster status alerts. The cluster IP address provides a single point for cluster access, configuration, and management. A Master IP Address resource that makes this possible is created automatically during the Cluster Services installation.

The cluster IP address will be bound to the master node and will remain with the master node regardless of which server is the master node.

- 9 (Conditional) If you are creating a new cluster, select the device where the SBD partition will be created.

For example, the device might be something similar to `sdc`.

If you have a shared disk system or SAN attached to your cluster servers, Novell Cluster Services will create a small cluster partition on that shared disk system. This small cluster partition is referred to as the Split Brain Detector (SBD) partition. Specify the drive or device where you want the small cluster partition created.

If you do not have a shared disk system connected to your cluster servers, accept the default (none).

IMPORTANT: You must have at least 20 MB of free space on one of the shared disk drives to create the cluster partition. If no free space is available, the shared disk drives can't be used by Novell Cluster Services.

- 10 (Conditional) If you want to mirror the SBD partition for greater fault tolerance, select the device where you want to mirror to, then click *Next*.

You can also mirror SBD partitions after installing Novell Cluster Services. See [Section 4.6.10, "Mirroring SBD \(Cluster\) Partitions,"](#) on page 60.

- 11 Select the IP address 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 will use.

- 12 Choose whether to start Novell Cluster Services software after configuring it, then click *Next*.

This option applies only to installing Novell Cluster Services after the OES installation, because it starts automatically when the server reboots during the OES installation.

If you choose to not start Novell Cluster Services software, you need to either manually start it after the installation, or reboot the cluster server to automatically start it.

You can manually start Novell Cluster Services by going to the `/etc/init.d` directory and entering `./novell-ncs start` at the server console of the cluster server.

- 13 Continue through the rest of the OES installation.

3.5.3 Installing Novell Cluster Services after the OES Installation

If you did not install Novell Cluster Services during the OES installation, you can install it later by completing the following steps:

- 1 At the Linux server console, type `yast2 ncs`.

This installs the Novell Cluster Services software component and takes you to the cluster configuration screen.

You must be logged in as root to access the cluster configuration screen.

- 2 Continue by completing [Step 5 on page 19](#) through [Step 12 on page 20](#).

3.5.4 Starting and Stopping Novell Cluster Services

Novell Cluster Services automatically starts after it is installed. Novell Cluster Services also automatically starts when you reboot your OES Linux server. If you need to manually start Novell Cluster Services, go to the `/etc/init.d` directory and run `./novell-ncs start`. You must be logged in as root to run `novell-ncs start`.

IMPORTANT: If you are using iSCSI for shared disk system access, ensure you have configured iSCSI initiators and targets to start prior to starting Novell Cluster Services. You can do this by entering `chkconfig iscsi on` at the Linux server console.

To have Novell Cluster Services not start automatically after rebooting your OES Linux server, enter `chkconfig novell-ncs off` at the Linux server console before rebooting the server. You can also enter `chkconfig novell-ncs on` to again cause Novell Cluster Services to automatically start.

3.6 Converting a NetWare Cluster to Linux

This section covers the following information to help you convert and manage a mixed NetWare 6.5 and Linux cluster.

- ◆ [Section 3.6.1, “Changing Existing NetWare Cluster Nodes to Linux \(Rolling Cluster Conversion\),” on page 21](#)
- ◆ [Section 3.6.2, “Adding New Linux Nodes to Your NetWare Cluster,” on page 23](#)
- ◆ [Section 3.6.3, “Mixed NetWare and Linux Clusters,” on page 24](#)
- ◆ [Section 3.6.4, “Finalizing the Cluster Conversion,” on page 26](#)

If you have a NetWare 5.1 cluster, you must upgrade to a NetWare 6.5 cluster before adding new Linux cluster nodes or converting existing NetWare cluster nodes to Linux cluster nodes.

See [“Upgrading Novell Cluster Services”](#) in the *Novell Cluster Services Administration Guide* for information on upgrading Novell Cluster Services.

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

3.6.1 Changing Existing NetWare Cluster Nodes to Linux (Rolling Cluster Conversion)

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

During a rolling cluster conversion, one server is converted to Linux while the other servers in the cluster continue running NetWare 6.5. Then, if desired, another server can be converted to Linux,

and then another, until all servers in the cluster have been converted to Linux. You can also leave the cluster as a mixed NetWare and Linux cluster.

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

IMPORTANT: Mixed NetWare 6.5 and OES Linux clusters are supported, and mixed NetWare 6.0 and OES Linux clusters are also supported. Mixed clusters consisting of NetWare 6.0 servers, NetWare 6.5 servers, and OES Linux servers are not supported. All NetWare servers must be either version 6.5 or 6.0 in order to exist in a mixed NetWare and OES Linux cluster.

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

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

- 1** On the NetWare server you want to convert to Linux, run NWConfig and remove eDirectory.
You can do this by selecting the option in NWConfig to remove eDirectory from the server.
- 2** Bring down the NetWare server you want to convert to Linux.
Any cluster resources that were running on the server should fail over to another server in the cluster.
- 3** In eDirectory, remove (delete) the Cluster Node object, the Server object, and all corresponding objects relating to the downed NetWare server.
Depending on your configuration, there could be up to 10 or more objects that relate to the downed NetWare server.
- 4** Run DSRepair on another server in the eDirectory tree to fix any directory problems.
If DSRepair finds errors or problems, run it multiple times until no errors are returned.
- 5** Install SLES 9 and OES on the server, but do not install the Cluster Services component of OES.
You can use the same server name and IP address that were used on the NetWare server. This is suggested, but not required.
See the *OES Linux Installation Guide* for more information.
- 6** Set up and verify SAN connectivity for the Linux node.
Consult your SAN vendor documentation for SAN setup and connectivity instructions.
- 7** Install Cluster Services and add the node to your existing NetWare 6.5 cluster.
See [Section 3.5.3, “Installing Novell Cluster Services after the OES Installation,”](#) on page 20 for more information.
- 8** Enter `sbdutil -f` at the Linux server console to verify that the node can see the cluster (SBD) partition on the SAN.
`sbdutil -f` also tells you the device on the SAN where the SBD partition is located.
- 9** Start cluster software by going to the `/etc/init.d` directory and running `./novell-ncs start`.
You must be logged in as root to run `novell-ncs start`.

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

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

- ♦ The failback mode for the resources was set to Auto.
- ♦ You used the same node number for this Linux server as was used for the former NetWare server.

This only applies if this Linux server is the next server added to the cluster.

- ♦ This Linux server is the preferred node for the resources.

Simultaneously Changing Multiple NetWare Cluster Nodes to Linux

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

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

Another option if new cluster node IP addresses are required and new server hardware is being used is to shut down the NetWare nodes that are to be removed and then add the new Linux cluster nodes. After adding the new Linux cluster nodes, you can remove the NetWare cluster node-related objects as described in [Step 3 on page 22](#).

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

3.6.2 Adding New Linux Nodes to Your NetWare Cluster

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

- 1 Install SLES 9 and OES on the new node, but do not install the Cluster Services component of OES.

See the “[OES Linux Installation Guide](#)” for more information.

- 2 Set up and verify SAN connectivity for the new Linux node.

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

- 3 Install Cluster Services and add the new node to your existing NetWare 6.5 cluster.

See [Section 3.5.3, “Installing Novell Cluster Services after the OES Installation,” on page 20](#) for more information.

- 4 Enter `sbdutil -f` at the Linux server console to verify that the node can see the cluster (SBD) partition on the SAN.

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

- 5 Start cluster software by going to the `/etc/init.d` directory and running `novell-ncs start`.

You must be logged in as root to run `novell-ncs start`.

- 6 Add and assign cluster resources to the new Linux cluster node.

3.6.3 Mixed NetWare and Linux Clusters

Novell Cluster Services includes some specialized functionality to help NetWare and Linux servers coexist in the same cluster. This functionality is also beneficial as you migrate NetWare cluster servers to Linux. It automates the conversion of the Master IP Address resource and cluster-enabled NSS pool resource load and unload scripts from NetWare to Linux. The NetWare load and unload scripts are read from eDirectory, converted, and written into Linux load and unload script files. Those Linux load and unload script files are then searched for NetWare-specific command strings, and the command strings are then either deleted or replaced with Linux-specific command strings. Separate Linux-specific commands are also added, and the order of certain lines in the scripts is also changed to function with Linux.

Cluster resources that were originally created on Linux cluster nodes cannot be migrated or failed over to NetWare cluster nodes. Cluster resources that were created on NetWare cluster nodes and migrated or failed over to Linux cluster nodes can be migrated or failed back to NetWare cluster nodes. If you want resources that can run on both NetWare and Linux cluster nodes, create them on a NetWare server.

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

WARNING: Changing existing shared pools or volumes (storage reconfiguration) in a mixed NetWare/Linux cluster is not possible. If you need to make changes to existing pools or volumes, you must temporarily bring down either all Linux cluster nodes or all NetWare cluster nodes prior to making changes. Attempting to reconfigure shared pools or volumes in a mixed cluster can cause data loss.

The following table identifies some of the NetWare specific cluster load and unload script commands that are searched for and the Linux commands that they are replaced with (unless deleted).

Table 3-1 .Cluster Script Command Comparison

Action	NetWare Cluster Command	Linux Cluster Command
Replace	IGNORE_ERROR add secondary ipaddress	ignore_error add_secondary_ipaddress
Replace	IGNORE_ERROR del secondary ipaddress	ignore_error del_secondary_ipaddress
Replace	del secondary ipaddress	ignore_error del_secondary_ipaddress
Replace	add secondary ipaddress	exit_on_error add_secondary_ipaddress
Delete	IGNORE_ERROR NUDP (deletes entire line)	
Delete	IGNORE_ERROR HTTP (deletes entire line)	
Replace	nss /poolactivate=	nss /poolact=
Replace	nss /pooldeactivate=	nss /pooldeact=
Replace	mount <i>volume_name</i> VOLID= <i>number</i>	exit_on_error ncpcon mount <i>volume_name=number</i>

Action	NetWare Cluster Command	Linux Cluster Command
Replace	NUDP ADD <i>clusterservername</i> <i>ipaddress</i>	<code>exit_on_error ncpcon bind -- ncpservername=<i>ncpservername</i> -- ipaddress=<i>ipaddress</i></code>
Replace	NUDP DEL <i>clusterservername</i> <i>ipaddress</i>	<code>ignore_error ncpcon unbind -- ncpservername=<i>ncpservername</i> -- ipaddress=<i>ipaddress</i></code>
Delete	CLUSTER CVSBIND (deletes entire line)	
Delete	CIFS (deletes entire line)	

Unlike NetWare cluster load and unload scripts which are stored in eDirectory, the Linux cluster load and unload scripts are stored in files on Linux cluster servers. The files are automatically updated each time you make changes to resource load and unload scripts for NetWare cluster resources. The cluster resource name is used in the load and unload script filenames. The path to the files is `/etc/opt/novell/ncs/`.

The following examples provide a sample comparison between NetWare cluster load and unload scripts, and their corresponding Linux cluster load and load scripts.

Master IP Address Resource Load Script

NetWare

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

Linux

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

Master IP Address Resource Unload Script

NetWare

```
IGNORE_ERROR HTTPUNBIND 10.1.1.175
IGNORE_ERROR del secondary ipaddress 10.1.1.175
IGNORE_ERROR NUDP DEL BCCP_Cluster 10.1.1.175
IGNORE_ERROR CLUSTER CVSBIND DEL BCCP_Cluster 10.1.1.175
```

Linux

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

```
ignore_error del_secondary_ipaddress 10.1.1.175
exit 0
```

NSS Pool Resource Load Script

NetWare

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

Linux

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
exit_on_error nss /poolact=HOMES_POOL
exit_on_error ncpcon mount HOMES=254
exit_on_error add_secondary_ipaddress 10.1.1.180
exit_on_error ncpcon bind --ncpservname=BCC_CLUSTER_HOMES_SERVER --
ipaddress=10.1.1.180
exit 0
```

NSS Pool Resource Unload Script

NetWare

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

Linux

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
ignore_error ncpcon unbind --ncpservname=BCC_CLUSTER_HOMES_SERVER --
ipaddress=10.1.1.180
ignore_error del_secondary_ipaddress 10.1.1.180
exit_on_error nss /pooldeact=HOMES_POOL
exit 0
```

3.6.4 Finalizing the Cluster Conversion

If you have converted all nodes in a former NetWare cluster to Linux, you must finalize the conversion process by issuing the `cluster convert` command on one Linux cluster node. The `cluster convert` command moves cluster resource load and unload scripts from the files

where they were stored on Linux cluster nodes into eDirectory. This enables a Linux cluster that has been converted from NetWare to utilize eDirectory like the former NetWare cluster.

To finalize the cluster conversion:

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

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

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

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

Generating Linux Cluster Resource Templates

After converting all nodes in a former NetWare cluster to Linux, you might want to generate the cluster resource templates that are included with Novell Cluster Services for Linux. These templates are automatically created when you create a new Linux cluster, but are not created when you convert an existing NetWare cluster to Linux.

To generate or regenerate the cluster resource templates that are included with Novell Cluster Services for Linux, enter the following command on a Linux cluster server:

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

In addition to generating Linux cluster resource templates, this command deletes all NetWare cluster resource templates. Because of this, use this command only after all nodes in the former NetWare cluster are converted to Linux.

3.7 Setting Up Novell Cluster Services

If you created a new cluster, you now need to create and configure cluster resources. You might also need to create shared disk partitions and NSS pools if they do not already exist and, if necessary, configure the shared disk NSS pools to work with Novell Cluster Services. Configuring shared disk NSS pools to work with Novell Cluster Services can include cluster enabling the pools.

You must use iManager or NSSMU to cluster enable shared NSS disk pools. iManager must be used to create cluster resources.

3.7.1 Creating NSS Shared Disk Partitions and Pools

Before creating disk partitions and pools on shared storage (storage area network or SAN), Novell Cluster Services must be installed. You should carefully plan how you want to configure your shared storage prior to installing Novell Cluster Services. For information on configuring access to a NetWare server functioning as an iSCSI target, see [Accessing iSCSI Targets on NetWare Servers from Linux Initiators](http://www.novell.com/documentation/iscsi1_nak/iscsi/data/bswmaoa.html#bt8cyhf) (http://www.novell.com/documentation/iscsi1_nak/iscsi/data/bswmaoa.html#bt8cyhf).

To create NSS pools on shared storage, use either the server-based NSS Management Utility (NSSMU) or iManager. These tools can also be used to create NSS volumes on shared storage.

NSS pools can be cluster enabled at the same time they are created or they can be cluster enabled at a later time after they are created. To learn more about NSS pools, see “Pools” in the *Novell Storage Services Administration Guide*.

Creating Shared NSS Pools Using NSSMU

- 1 Start NSSMU by entering `nssmu` at the server console of a cluster server.
- 2 Select Devices from the NSSMU main menu and mark all shared devices as sharable for clustering.
On Linux, shared disks are not by default marked sharable for clustering. If a device is marked as sharable for clustering, all partitions on that device will automatically be sharable.
You can press F6 to individually mark devices as sharable.
- 3 From the NSSMU main menu, select Pools, press Insert, and then type a name for the new pool you want to create.
- 4 Select the device on your shared storage where you want the pool created.
Device names might be labelled something like `/dev/sdc`.
- 5 Choose whether you want the pool to be activated and cluster enabled when it is created.
The Activate on Creation feature is enabled by default. This causes the pool to be activated as soon as it is created. If you choose not to activate the pool, you will have to manually activate it later before it can be used.
The Cluster Enable on Creation feature is also enabled by default. If you want to cluster enable the pool at the same time it is created, accept the default entry (Yes) and continue with [Step 6](#). If you want to cluster enable the pool at a later date, change the default entry from Yes to No, select Create, and then go to “[Creating NSS Volumes](#)” on [page 29](#).
- 6 Specify the virtual server name, IP address, and advertising protocols.

NOTE: The CIFS and AFP check boxes can be checked, but CIFS and AFP functionality does not apply to Linux. Checking the checkboxes has no effect.

When you cluster enable a pool, a virtual Server object is automatically created and given the name of the Cluster object plus the cluster-enabled pool. For example, if the cluster name is `cluster1` and the cluster-enabled pool name is `pool1`, then the default virtual server name will be `cluster1_pool1_server`. You can edit the field to change the default virtual server name.

Each cluster-enabled NSS pool requires its own IP address. The IP address is used to provide access and failover capability to the cluster-enabled pool (virtual server). The IP address you assign to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.

You can select or deselect NCP. NCP™ is selected by default, and is the protocol used by Novell clients. Selecting NCP will cause commands to be added to the pool resource load and unload scripts to activate the NCP protocol on the cluster. This lets you ensure that the cluster-enabled pool you just created is highly available to Novell clients.

- 7 Select Create to create and cluster enable the pool.
Repeat the above steps for each additional pool you want to create on shared storage.
Continue with “[Creating NSS Volumes](#)” on [page 29](#).

Creating Shared NSS Pools Using iManager

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

The URL is `http://server_ip_address/nps/imanager.html`. Replace *server_ip_address* with the IP address or DNS name of a Linux server in the cluster that has iManager installed or with the IP address for Apache-based services.

- 2 Enter your username and password.
- 3 In the left column, locate *Storage*, then click the *Pools* link.
- 4 Enter a cluster server name or browse and select one, then click the *New* link.
- 5 Specify the new pool name, then click *Next*.
- 6 Check the box next to the device where you want to create the pool, then specify the size of the pool.
- 7 Choose whether you want the pool to be activated and cluster-enabled when it is created, then click *Next*.

The *Activate On Creation* check box is used to determine if the pool you are creating is to be activated as soon as it is created. The *Activate On Creation* check box is checked by default. If you uncheck the check box, you must manually activate the pool later before it can be used.

If you want to cluster enable the pool at the same time it is created, leave the *Cluster Enable on Creation* check box checked and continue with [Step 8 on page 29](#).

If you want to cluster enable the pool at a later date, uncheck the check box, click *Create*, and continue with [“Cluster Enabling NSS Pools and Volumes” on page 31](#).

- 8 Specify the virtual server name, pool IP address, and advertising protocols, then click *Finish*.

NOTE: The CIFS and AFP check boxes can be checked, but CIFS and AFP functionality does not apply to Linux. Checking the check boxes has no effect.

When you cluster-enable a pool, a virtual Server object is automatically created and given the name of the Cluster object plus the cluster-enabled pool. For example, if the cluster name is `cluster1` and the cluster-enabled pool name is `pool1`, then the default virtual server name will be `cluster1_pool1_server`. You can edit the field to change the default virtual server name.

Each cluster-enabled NSS pool requires its own IP address. The IP address is used to provide access and failover capability to the cluster-enabled pool (virtual server). The IP address you assign to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.

You can select or deselect *NCP*. *NCP* is selected by default, and is the protocol used by Novell clients. Selecting *NCP* causes commands to be added to the pool resource load and unload scripts to activate the *NCP* protocol on the cluster. This lets you ensure that the cluster-enabled pool you just created is highly available to Novell clients.

3.7.2 Creating NSS Volumes

If you plan on using a shared disk system in your cluster and need to create new NSS pools or volumes after installing Novell Cluster Services, the server used to create the volumes should already have NSS installed and running.

Using NSSMU

- 1 From the NSSMU main menu, select Volumes, then press Insert and type a name for the new volume you want to create.

Each shared volume in the cluster must have a unique name.

- 2 Select the pool where you want the volume to reside.
- 3 Review and change volume attributes as necessary.

You might want to enable the Flush Files Immediately feature. This will help ensure the integrity of volume data. Enabling the Flush Files Immediately feature improves file system reliability but hampers performance. You should consider this option only if necessary.

- 4 Either specify a quota for the volume or accept the default of 0 to allow the volume to grow to the pool size, then select Create.

The quota is the maximum possible size of the volume. If you have more than one volume per pool, you should specify a quota for each volume rather than allowing multiple volumes to grow to the pool size.

- 5 Repeat the above steps for each cluster volume you want to create.

Depending on your configuration, the new volumes will either mount automatically when resources that require them start or will have to be mounted manually on individual servers after they are up.

Using iManager

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

The URL is `http://server_ip_address/nps/imanager.html`. Replace *server_ip_address* with the IP address or DNS name of a Linux server in the cluster that has iManager installed or with the IP address for Apache-based services.

- 2 Enter your username and password.
- 3 In the left column, locate *Storage*, then click the *Volumes* link.
- 4 Enter a cluster server name or browse and select one, then click the *New* link.
- 5 Specify the new volume name, then click *Next*.

- 6 Check the box next to the cluster pool where you want to create the volume and either specify the size of the volume (Volume Quota) or check the box to allow the volume to grow to the size of the pool, then click *Next*.

The volume quota is the maximum possible size of the volume. If you have more than one volume per pool, you should specify a quota for each volume rather than allowing multiple volumes to grow to the pool size.

- 7 Review and change volume attributes as necessary.

The Flush Files Immediately feature helps ensure the integrity of volume data. Enabling the Flush Files Immediately feature improves file system reliability but hampers performance. You should consider this option only if necessary.

- 8 Choose whether you want the volume activated and mounted when it is created, then click *Finish*.

3.7.3 Cluster Enabling NSS Pools and Volumes

If you have a shared disk system that is part of your cluster and you want the pools and volumes on the shared disk system to be highly available to NCP clients, you will need to cluster enable those pools and volumes. Cluster enabling a pool or volume allows it to be moved or mounted on different servers in the cluster in a manner that supports transparent client reconnect.

Cluster-enabled volumes do not appear as cluster resources. NSS pools are resources, and load and unload scripts apply to pools and are automatically generated for them. Each cluster-enabled NSS pool requires its own IP address. This means that each cluster-enabled volume does not have an associated load and unload script or an assigned IP address.

NSS pools can be cluster enabled at the same time they are created. If you did not cluster enable a pool at creation time, the first volume you cluster enable in the pool automatically cluster enables the pool where the volume resides. After a pool has been cluster enabled, you need to cluster enable the other volumes in the pool if you want them to be mounted on another server during a failover.

When a server fails, any cluster-enabled pools being accessed by that server will fail over to other servers in the cluster. Because the cluster-enabled pool fails over, all volumes in the pool will also fail over, but only the volumes that have been cluster enabled will be mounted. Any volumes in the pool that have not been cluster enabled will have to be mounted manually. For this reason, volumes that aren't cluster enabled should be in separate pools that are not cluster enabled.

If you want each cluster-enabled volume to be its own cluster resource, each volume must have its own pool.

Some server applications don't require NCP client access to NSS volumes, so cluster enabling pools and volumes might not be necessary. Pools should be deactivated and volumes should be dismounted before being cluster enabled.

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

The URL is `http://server_ip_address/nps/imanager.html`. Replace *server_ip_address* with the IP address or DNS name of a Linux server in the cluster that has iManager installed or with the IP address for Apache-based services.

- 2 Enter your username and password.
- 3 In the left column, locate *Clusters*, then click the *Cluster Options* link.

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

- 4 Enter the cluster name or browse and select it, then click the *New* link.
- 5 Specify Pool as the resource type you want to create by clicking the *Pool* radio button, then click *Next*.
- 6 Enter the name of the pool you want to cluster-enable, or browse and select one.
- 7 (Optional) Change the default name of the virtual Server object.

When you cluster enable a pool, a Virtual Server object is automatically created and given the name of the Cluster object plus the cluster-enabled pool. For example, if the cluster name is `cluster1` and the cluster-enabled pool name is `pool1`, then the default virtual server name will be `cluster1_pool1_server`.

If you are cluster-enabling a volume in a pool that has already been cluster-enabled, the virtual Server object has already been created, and you can't change the virtual Server object name.

- 8 Enter an IP address for the pool.

Each cluster-enabled NSS pool requires its own IP address. The IP address is used to provide access and failover capability to the cluster-enabled pool (virtual server). The IP address assigned to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.

- 9 Select an advertising protocol.

NOTE: The CIFS and AFP check boxes can be checked, but CIFS and AFP functionality does not apply to Linux. Checking the check boxes has no effect.

You can select or deselect NCP. NCP™ is selected by default, and is the protocol used by Novell clients. Selecting NCP will cause commands to be added to the pool resource load and unload scripts to activate the NCP protocol on the cluster. This lets you ensure that the cluster-enabled pool you just created is highly available to Novell clients.

- 10 (Optional) Check the *Online Resource after Create* check box.

This causes the NSS volume to automatically mount when the resource is created.

- 11 Ensure that the *Define Additional Properties* check box is checked, then click *Next* and continue with [“Setting Start, Failover, and Failback Modes” on page 42](#).

NOTE: Cluster resource load and unload scripts are automatically generated for pools when they are cluster-enabled.

When the volume resource is brought online, the pool will automatically be activated. You don't need to activate the pool at the server console.

If you delete a cluster-enabled volume, Novell Cluster Services automatically removes the volume mount command from the resource load script. If you delete a cluster-enabled pool, Novell Cluster Services automatically removes the Pool Resource object and the virtual server object from eDirectory. If you rename a cluster-enabled pool, 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.

3.7.4 Creating Traditional Linux Volumes on Shared Disks

Although you can use the same Linux tools and procedures used to create partitions on local drives to create Linux file system partitions on shared storage, EVMS is the recommended tool. Using EVMS to create partitions, volumes, and file systems will help prevent data corruption caused by multiple nodes accessing the same data. You can create partitions and volumes using any of the journaled Linux file systems (EXT3, Reiser, etc.). To cluster enable Linux volumes, see [Section 3.7.6, “Cluster Enabling Traditional Linux Volumes on Shared Disks,” on page 36](#).

TIP: EVMS virtual volumes are recommended for Novell Cluster Services because they can more easily be expanded and failed over to different cluster servers than physical devices. You can enter `man evms` at the Linux server console to reference the `evms` man page, which provides additional instructions and examples for `evms`.

You can also enter `man mount` at the Linux server console to reference the `mount` man page, which provides additional instructions and examples for the `mount` command.

The following sections provide the necessary information for using EVMS to create a traditional Linux volume and file system on a shared disk:

- ♦ “Ensuring That the Shared Disk Is not a Compatibility Volume” on page 33
- ♦ “Removing Other Segment Managers” on page 33
- ♦ “Creating a Cluster Segment Manager Container” on page 34
- ♦ “Adding an Additional Segment Manager” on page 34
- ♦ “Creating an EVMS Volume” on page 35
- ♦ “Creating a File System on the EVMS Volume” on page 35

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

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

Ensuring That the Shared Disk Is not a Compatibility Volume

New EVMS volumes are by default configured as compatibility volumes. If any of the volumes on your shared disk (that you plan to use in your cluster) are compatibility volumes, you must delete them.

- 1 At the Linux server console, enter `evmsgui`.
- 2 Click the *Volumes* tab, then right-click the volume on the shared disk and select *Display details*.
- 3 Click the *Page 2* tab and determine from the Status field if the volume is a compatibility volume.

If the volume is a compatibility volume or has another segment manager on it, continue with **Step 3a** below.

- 3a Click the *Volumes* tab, right-click the volume, then select *Delete*.
- 3b Select the volume, then click *Recursive Delete*.
- 3c (Conditional) If a Response Required pop-up appears, click the *Write zeros* button.
- 3d (Conditional) If another pop-up appears, click *Continue* to write 1024 bytes to the end of the volume.

Removing Other Segment Managers

If any of the shared disks you plan to use with your cluster have other segment managers, you must delete them as well.

- 1 In `evmsgui`, click the *Disks* tab, then right-click the disk you plan to use for a cluster resource.
- 2 Select *remove segment manager from Object*.
This option only appears if there is another segment manager for the selected disk.
- 3 Select the listed segment manager and click *Remove*.

Creating a Cluster Segment Manager Container

To use a traditional Linux volume with EVMS as a cluster resource, you must use the Cluster Segment Manager (CSM) plug-in for EVMS to create a CSM container.

NOTE: CSM containers require Novell Cluster Services (NCS) to be running on all nodes that access the CSM container. Do not make to modifications to EVMS objects unless NCS is running.

CSM containers can provide exclusive access to shared storage.

- 1 In *evmsgui*, click *Actions*, select *Create*, then select *Container*.
- 2 Select *Cluster Segment Manager*, then click *Next*.
- 3 Select the disks (storage objects) you want to place in the container, then click *Next*.
- 4 On the Configuration Options page, select the node where you are creating the container, specify *Private* as the type, then choose a name for the container.

The name must be one word, must consist of standard alphanumeric characters, and must not be any of the following reserved words:

- ♦ Container
- ♦ Disk
- ♦ EVMS
- ♦ Plugin
- ♦ Region
- ♦ Segment
- ♦ Volume

- 5 Click *Save* to save your changes.

Adding an Additional Segment Manager

After creating a CSM container, you can optionally add an additional non-CSM segment manager container on top of the CSM container you just created. The benefit of this is that other non-CSM segment manager containers allow you to create multiple smaller EVMS volumes on your EVMS disk. You can then add additional EVMS volumes or expand or shrink existing EVMS volumes to utilize or create additional free space on your EVMS disk. In addition, this means that you can also have different file system types on your EVMS disk.

A CSM container uses the entire EVMS disk, which means that creating additional volumes or expanding or shrinking volumes is not possible. And, because only one EVMS volume is possible in the container, only one file system type is allowed in that container.

- 1 In *evmsgui*, click *Actions*, select *Add*, then select *Segment Manager to Storage Object*.
- 2 Choose the desired segment manager, then click *Next*.
Most of the segment manager will work. The DOS segment manager is added by default for some EVMS operations.
- 3 Choose the storage object (container) you want to add the segment manager to, then click *Next*.
- 4 Select the disk type (Linux is the default), click *Add*, then click *OK*.
- 5 Click *Save* to save your changes.

Creating an EVMS Volume

- 1 In *evmsgui*, click *Actions*, select *Create*, and then *EVMS Volume*.
- 2 Select the container you just created (either the CSM container or the additional segment manager container) and specify a volume name.
- 3 Click *Create*, then click *Save*.

Creating a File System on the EVMS Volume

- 1 In *evmsgui*, click the *Volumes* tab and right-click the volume you just created.
- 2 Select *Make File System*, choose a traditional Linux file system from the list, then click *Next*.
- 3 Specify a volume label, then click *Make*.
- 4 Save your changes by clicking *Save*.

3.7.5 Expanding EVMS Volumes on Shared Disks

As your storage needs increase, it might become necessary to add more disk space or drives to your shared storage system. EVMS provides features that allow you to expand or move existing volumes.

The two supported methods for creating additional space for an existing volume are:

- ♦ Expanding the volume to a separate disk
- ♦ Moving the volume to a larger disk

Expanding a Volume to a Separate Disk

- 1 Unmount the file system for the volume you want to expand.
- 2 In *evmsgui*, click the *Volumes* tab, right-click the volume you want to expand then select *Add Feature*.
- 3 Select *Drive Linking Feature*, then click *Next*.
- 4 Provide a name for the drive link, click *Add*, then save your changes.
- 5 Click *Actions*, select *Create*, and then click *Container*.
- 6 Select the *Cluster Segment Manager*, click *Next*, then select the disk you want to expand the volume to.
The entire disk is used for the expansion, so you must select a disk that does not have other volumes on it.
- 7 Provide the same settings information (name, type, etc.) as the existing container for the volume and save your changes.
- 8 Click the *Volumes* tab, right-click the volume, then click *Expand*.
- 9 Select the volume that you are expanding, then click *Next*.
- 10 Verify the current volume size and the size of the volume after it is expanded, then click *Next*.
The expanded volume size should include the size of the disk the volume is being expanded to.
- 11 Select the storage device the volume is being expanded to, select *Expand*, and save your changes.
- 12 Click *Save* and exit *evmsgui*.

Moving a Volume to a Larger Disk

- 1 Unmount the file system for the volume you want to move.
- 2 Add a larger disk to the CSM container.
 - 2a In evmsgui, click *Actions*, select *Create*, then click *Container*.
 - 2b Select the *Cluster Segment Manager*, then click *Next*.
 - 2c Select the larger disk you want to move the volume to.

The entire disk is used for the expansion, so you must select a disk that does not have other volumes on it.
 - 2d Provide the same settings information (name, type, etc.) as the existing container for the volume, then save your changes.
 - 2e Click *Save* and exit evmsgui.
- 3 Restart evmsgui, click the *Containers* tab, then expand the container so that the objects under the container appear.

The new disk should appear as part of the container.
- 4 Right-click the object for the disk where the volume resides and select *Replace*.
- 5 Select the object for the disk where the volume will be moved, then click *Next*.
- 6 Save your changes.

Saving your changes could take a while, depending on volume size and other factors.
- 7 Click *Save*, exit evmsgui, then restart evmsgui.
- 8 Click the *Volumes* tab, right-click the volume, then select *Check/Repair filesystem*.

This will run the repair process and ensure no problems exist on the moved volume.
- 9 Click the *Disks* tab, right-click the disk the volume was moved from, then select *Remove from container*.
- 10 Click *Save* and exit evmsgui.

3.7.6 Cluster Enabling Traditional Linux Volumes on Shared Disks

Cluster enabling a traditional Linux volume allows it to be moved or mounted on different servers in the cluster. This provides a way for clients to reconnect to the volume regardless of which server is hosting it.

EVMS containers are the unit of failover for traditional Linux volumes. Because the EVMS container is the unit of failover, all volumes in a container will also fail over, but only the volumes that are mounted through the cluster resource load script will be mounted. Any volumes in the container that are not mounted through the resource load script will have to be mounted manually.

The following sections contain information on cluster enabling a traditional Linux volume on a shared disk partition:

- ◆ [“Creating a Traditional Linux Volume Cluster Resource” on page 37](#)
- ◆ [“Configuring Traditional Linux Volume Load Scripts” on page 37](#)
- ◆ [“Configuring Traditional Linux Volume Unload Scripts” on page 39](#)

Creating a Traditional Linux Volume Cluster Resource

Creating a cluster resource for a traditional Linux volume allows it to be moved or mounted on different servers in the cluster.

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

The URL is `http://server_ip_address/nps/imanager.html`. Replace *server_ip_address* with the IP address or DNS name of an OES server in the cluster that has iManager installed or with the IP address for Apache-based services.

- 2 Enter your username and password.
- 3 In the left column, locate *Clusters*, then click the *Cluster Options* link.

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

- 4 Specify the cluster name or browse and select it, then click the *New* link.
- 5 Specify Resource as the resource type you want to create by clicking the *Resource* radio button, then click *Next*.
- 6 Specify the name of the resource you want to create.
This is the name you will assign the resource for the cluster-enabled volume.
- 7 In the Inherit From Template field, specify the `Generic_FS_Template`.
- 8 Select the *Define Additional Properties* check box, then continue with **Configuring Traditional Linux Volume Load Scripts** below.

Configuring Traditional Linux Volume Load Scripts

The resource load script specifies the commands to start the resource (including mounting the file system) on a server in the cluster, and is required for each Linux volume you cluster enable.

If you are creating a new cluster resource, the load script page should already be displayed. You can start with **Step 4**.

- 1 In the left column of the main iManager page, locate *Clusters*, then click the *Cluster Options* link.
- 2 Type the cluster name or browse and select it, select the box next to the resource whose load script you want to edit, then click the *Properties* link.
- 3 Click the *Scripts* tab, then click the *Load Script* link.

- 4 Edit or add the necessary commands to the script to load the resource on the server.

The generic file system template you specified in **Step 7** above contains a load script that you must edit to supply information specific to your file system resource.

The load script from the generic file system template should appear similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuns
# define the IP address
RESOURCE_IP=a.b.c.d
# define the file system type
MOUNT_FS=reiserfs
```

```

#define the container name
container_name=name
# define the device
MOUNT_DEV=/dev/evms/$container_name/volume_name
# define the mount point
MOUNT_POINT=/mnt/mount_point

#activate the container
exit_on_error activate_evms_container $container_name $MOUNT_DEV
# mount the file system
ignore_error mkdir -p $MOUNT_POINT
exit_on_error mount -t $MOUNT_FS $MOUNT_DEV $MOUNT_POINT
# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP
exit 0

```

The first section of the above load script example contains mount point, IP address, container name, and file system type/device variables that you must change to customize the script for your specific configuration.

- 5 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.

In the above example, if you specified

123.123.12.12 as the IP address
reiserfs as the file system type
cont1 as the container name
vol_one as the device
/mnt/vol_onemount as the mount point

your load script would appear like the script below.

```

#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
# define the IP address
RESOURCE_IP=123.123.12.12
# define the file system type
MOUNT_FS=reiserfs
#define the container name
container_name=cont1
# define the device
MOUNT_DEV=/dev/evms/$container_name/vol_one
# define the mount point
MOUNT_POINT=/mnt/vol_onemount

#activate the container
exit_on_error activate_evms_container $container_name $MOUNT_DEV

```

```
# mount the file system
ignore_error mkdir -p $MOUNT_POINT
exit_on_error mount -t $MOUNT_FS $MOUNT_DEV $MOUNT_POINT
# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP
exit 0
```

Configuring Traditional Linux Volume Unload Scripts

The resource unload script specifies the commands to stop the resource (including unmounting the file system) on a server in the cluster, and is also required for each Linux volume you cluster enable. If you are creating a new cluster resource, the unload script page should already be displayed. You can start with [Step 4](#).

- 1 In the left column of the main iManager page, locate *Clusters*, then click the *Cluster Options* link.
- 2 Type the cluster name or browse and select it, select the box next to the resource whose unload script you want to edit, then click the *Properties* link.
- 3 Click the *Scripts* tab, then click the *Unload Script* link.
- 4 Edit or add the necessary commands to the script to unload or stop the resource on the server.

The generic file system template you specified in [Step 7](#) above contains an unload script that you must edit to supply information specific to your file system resource.

The unload script from the generic file system template should appear similar to the following example:

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
# define the IP address
RESOURCE_IP=a.b.c.d
#define the container name
container_name=name
# define the mount point
MOUNT_POINT=/mnt/mount_point

#dismount the volume
exit_on_error ncs_dismount $MOUNT_POINT
# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP
# deport the container
exit_on_error deport_evms_container $container_name
# return status
exit 0
```

The first section of the above unload script example contains mount point, container name, and IP address variables that you must change to customize the unload script for your specific configuration.

- 5 Specify the *Unload Script Timeout* value, then click *Apply* to save the script or, if you are creating a new cluster resource, click *Next* and continue with [Section 3.7.11, “Setting Start, Failover, and Failback Modes,”](#) on page 42.

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.

In the above example, if you specified

123.123.12.12 as the IP address

cont1 as the container name

/mnt/vol_onemount as the mount point

Your unload script would appear like the script below.

```
#!/bin/bash
. /opt/novell/ncs/lib/ncsfuncs
# define the IP address
RESOURCE_IP=123.123.12.12
#define the container name
container_name=cont1
# define the mount point
MOUNT_POINT=/mnt/vol_onemount

#dismount the volume
exit_on_error ncs_dismount $MOUNT_POINT
# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP
# deport the container
exit_on_error deport_evms_container $container_name
# return status
exit 0
```

3.7.7 Creating Cluster Resource Templates

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

Novell Cluster Services provides the following cluster resource templates:

- ◆ DHCP
- ◆ DNS
- ◆ iFolder 2
- ◆ iPrint
- ◆ MYSQL
- ◆ Samba
- ◆ Generic IP SERVICE

This template can be modified to create cluster resources for certain server applications that run on your cluster.

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

The URL is `http://server_ip_address/nps/imanager.html`. Replace *server_ip_address* with the IP address or DNS name of an OES server in the cluster that has iManager installed, or the IP address for Apache-based services.

- 2 Enter your username and password.
- 3 In the left column, locate *Clusters*, then click the *Cluster Options* link.
iManager displays four links under Clusters that you can use to configure and manage your cluster.
- 4 Enter the cluster name or browse and select it, then click the *New* link.
- 5 Specify *Template* as the resource type you want to create by clicking the *Template* radio button, then click *Next*.
- 6 Enter the name of the template you want to create.
- 7 Ensure the *Define Additional Properties* check box is checked, then continue with [“Configuring Load Scripts” on page 42](#).

To finish creating a cluster resource template, you need to configure load and unload scripts, set failover and failback modes and, if necessary, change the node assignments for the resource template.

3.7.8 Creating Cluster Resources

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

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

The URL is `http://server_ip_address/nps/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 Enter your username and password.
- 3 In the left column, locate *Clusters*, then click the *Cluster Options* link.
iManager displays four links under Clusters that you can use to configure and manage your cluster.
- 4 Enter the cluster name, or browse and select it, then click the *New* link.
- 5 Specify *Resource* as the resource type you want to create by clicking the *Resource* radio button, then click *Next*.
- 6 Enter 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 will be converted to an underscore.

- 7 Check the *Define Additional Properties* check box.
- 8 Continue with [“Configuring Load Scripts” on page 42](#).

3.7.9 Configuring Load Scripts

A load script is required for each resource, service, or disk pool in your cluster. The load script specifies the commands to start the resource or service on a server.

If you are creating a new cluster resource, the load script page should already be displayed. You can start with [Step 4](#).

- 1 In the left column of the main iManager page, locate *Clusters*, then click the *Cluster Options* link.
- 2 Enter the cluster name or browse and select it, then check the box next to the resource whose load script you want to edit and click the *Properties* link.
- 3 Click the *Scripts* tab, then click the *Load Script* link.
- 4 Edit or add the necessary commands to the script to load the resource on the server.
You can then add any lines to the load script that are required to load needed services like Web servers, etc.
- 5 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.

3.7.10 Configuring Unload Scripts

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.

If you are creating a new cluster resource, the unload script page should already be displayed. You can start with [Step 4](#).

- 1 In the left column of the main iManager page, locate *Clusters*, then click the *Cluster Options* link.
- 2 Enter the cluster name or browse and select it, check the box next to the resource whose unload script you want to edit, then click the *Properties* link.
- 3 Click the *Scripts* tab, then click the *Unload Script* link.
- 4 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.
- 5 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.

3.7.11 Setting Start, Failover, and Failback Modes

You can configure the start, failover, and failback of cluster resources to happen manually or automatically. With the resource Start Mode set to AUTO, the resource automatically starts on a

server when the cluster is first brought up. If the resource Start Mode is set to MANUAL, you can manually start the resource on a server when you want, instead of having it automatically start when servers in the cluster are brought up.

With the resource Failover Mode set to AUTO, the resource automatically starts on the next server in the Assigned Nodes list in the event of a hardware or software failure. If the resource Failover Mode is set to MANUAL, you can intervene after a failure occurs and before the resource is moved to another node.

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

The preferred node is the first server in the list of the assigned nodes for the resource.

If you are creating a new cluster resource, the Resource Policies page should already be displayed. You can start with [Step 4](#).

- 1 In the left column of the main iManager page, locate *Clusters* and then click the *Cluster Options* link.
- 2 Enter the cluster name or browse and select it, check the box next to the resource whose start, failover, or failback modes you want to view or edit, then click the *Properties* link.
- 3 Click the *General* tab.
- 4 (Conditional) Check 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 will fail over to whichever node becomes the master.

- 5 (Conditional) Check 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.

Checking this box will ensure 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.

- 6 Choose 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.

- 7 Continue with [Assigning Nodes to a Resource](#), or if you are creating a new cluster resource, click Next, then continue with [Assigning Nodes to a Resource](#).

3.7.12 Assigning Nodes to a Resource

If you are creating a new cluster resource, the Preferred Nodes page should already be displayed. If you are assigning nodes for an existing resource, the Preferred Nodes page will be displayed as part of the Resource Policies page. You can start with [Step 4](#).

- 1 In the left column of the main iManager page, locate *Clusters*, then click the *Cluster Options* link.

- 2 Enter the cluster name or browse and select it, check the box next to the resource whose start, failover, or failback modes you want to view or edit, then click the *Properties* link.
- 3 Click the *General* tab.
- 4 From the list of unassigned nodes, select the server you want the resource assigned to, then click the right-arrow button to move the selected server to the Assigned Nodes list.
Repeat this step for all servers you want assigned to the resource. You can also use the left-arrow button to unassign servers from the resource.
- 5 Click the up-arrow and down-arrow buttons to change the failover order of the servers assigned to the resource or volume.
- 6 Click *Apply* or *Finish* to save node assignment changes.

3.8 Configuration Settings

Depending on your needs and cluster setup, some additional configuration might be required for you to effectively use Novell Cluster Services. This additional configuration might consist of changing the values on some of the properties for the Cluster object and the Cluster Node objects.

3.8.1 Editing Quorum Membership and Timeout Properties

You can edit Quorum Membership and Timeout properties using iManager.

- 1 In the left column of the main iManager page, locate *Clusters*, then click the *Cluster Options* link.
- 2 Enter the cluster name or browse and select it, then click the *Properties* button under the cluster name.
- 3 Click the *General* tab.

In iManager, the same page used to edit quorum membership and timeout is also used for the cluster IP address and port properties and for cluster e-mail notification.

Quorum Triggers (Number of Nodes)

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

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

Quorum Triggers (Timeout)

Timeout specifies the amount of time to wait for the number of servers defined in the Number of Nodes field to be up and running. If the timeout period elapses before the quorum membership reaches its specified number, resources will 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 will begin to load on the two servers that are up and running in the cluster.

3.8.2 Cluster Protocol Properties

You can use the Cluster Protocol property pages to view or edit the transmit frequency and tolerance settings for all nodes in the cluster, including the master node. The master node is generally the first node brought online in the cluster, but if that node fails, any of the other nodes in the cluster can become the master.

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

- 1 In the left column of the main iManager page, locate *Clusters*, then click the *Cluster Options* link.
- 2 Enter the cluster name or browse and select it, then click the *Properties* button under the cluster name.
- 3 Click the *Protocols* tab.

This page also lets you view the script used to configure the cluster protocol settings, but not change it. Changes made to the protocols setting will automatically update the script.

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, nonmaster nodes in the cluster send a signal that they are alive to the master node every second.

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.

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.

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 nonmaster 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.

Max Retransmits

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

3.8.3 Cluster IP Address and Port Properties

The Cluster IP address is assigned when you install Novell Cluster Services. The Cluster IP address normally does need to be changed, but can be if needed.

The default cluster port number is 7023, and is automatically assigned when the cluster is created. The cluster port number does not need to be changed unless a conflict is created by another resource using the same port number. If there is a port number conflict, change the Port number to any other value that doesn't cause a conflict.

- 1 In the left column of the main iManager page, locate *Clusters*, then click the *Cluster Options* link.
- 2 Enter the cluster name or browse and select it, then click the *Properties* button under the cluster name.
- 3 Click the *General* tab.

In iManager, the same page used to view or edit the cluster IP address and port properties is also used for quorum membership and timeout and for cluster e-mail notification.

3.8.4 Resource Priority

The Resource Priority allows you to control 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.

This is useful for ensuring that the most critical resources load first and are available to users before less critical resources.

- 1 In the left column of the main iManager page, locate *Clusters*, then click the *Cluster Options* link.
- 2 Enter the cluster name or browse and select it, then click the *Properties* button under the cluster name.
- 3 Click the *Priorities* tab.
- 4 To change the priority for a resource, select the resource in the list by clicking it, then click the up-arrow or down-arrow to move the resource up or down in the list. This lets you change the load order of the resource relative to other cluster resources on the same node.
- 5 Click the *Apply* button to save changes made to resource priorities.

3.8.5 Cluster E-Mail Notification

Novell Cluster Services can automatically send out e-mail messages for certain cluster events like cluster and resource state changes or nodes joining or leaving the cluster.

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

- 1 In the left column of the main iManager page, locate *Clusters*, then click the *Cluster Options* link.
- 2 Enter the cluster name or browse and select it, then click the *Properties* button under the cluster name.
- 3 Click the *General* tab.
- 4 Check or uncheck the *Enable Cluster Notification Events* check box to enable or disable e-mail notification.
- 5 If you enable e-mail notification, add the desired e-mail addresses in the field provided.

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

- 6 If you enable e-mail notification, specify the type of cluster events you want administrators to receive messages for.

To only receive notification of critical events like a node failure or a resource going comatose, click the *Receive Only Critical Events* radio button.

To receive notification of all cluster state changes including critical events, resource state changes, and nodes joining and leaving the cluster, click the *Verbose Messages* radio button.

To receive notification of all cluster state changes in XML format, choose the *XML Messages* option. XML format messages can be interpreted and formatted with a parser that lets you customize the message information for your specific needs.

- 7 Click the *Apply* button to save changes.

IMPORTANT: Novell Cluster Services uses Postfix to send e-mail alerts. If you have a cluster resource that uses SMTP, that resource might not work in the cluster unless you change the Postfix configuration. For example, GroupWise® uses SMTP and will not function as a cluster resource if Postfix uses the same port, which it does by default. In this case, Postfix must be configured to use a different port. You can do this by editing the `etc/postfix/main.cf` file and changing the values for the `inet_interfaces`, `mydestination`, and `mynetworks_style` lines. You also need to change the listen port for the `smtpd` process in the `etc/postfix/master.cf` file. See [the Postfix Web site \(http://www.postfix.org\)](http://www.postfix.org) for more information on configuring Postfix.

3.8.6 Cluster Node Properties

You can view or edit the cluster node number or IP address of the selected node or view the context for the Linux Server object.

- 1 In the left column of the main iManager page, locate *Clusters*, then click the *Cluster Options* link.
- 2 Enter the cluster name or browse and select it, check the box next to the cluster node whose properties you want to view or edit, then click the *Properties* link.
- 3 View or edit the IP address, then click *Apply* to update the information in eDirectory.

If the IP address changes for this server, the new information is not automatically updated in eDirectory.

(Node) Number+IP Address

Number+IP Address specifies the cluster node number and IP address for the selected node. If the cluster node number or IP address changes for the selected node, the new information is not automatically updated in eDirectory. Edit the information and click *Apply* to update the information in eDirectory.

Distinguished Name

The Distinguished Name is the eDirectory name and context for the Server object.

3.9 Additional Information

For additional information on managing Novell Cluster Services, see [Chapter 4, “Managing Novell Cluster Services,”](#) on page 49.

Managing Novell Cluster Services

4

After you have installed, set up, and configured Novell® Cluster Services™ for your specific needs, some additional information can be useful to help you effectively manage your cluster. This information consists of instructions for migrating resources, identifying cluster and resource states, and customizing cluster management.

4.1 Migrating Resources

You can migrate resources to different servers in your cluster without waiting for a failure to occur. You might want to migrate resources to lessen the load on a specific server, to free up a server so it can be brought down for scheduled maintenance, or to increase the performance of the resource or application by putting it on a faster machine.

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

- 1 In the left column of the main iManager page, locate *Clusters*, then click the *Cluster Manager* link.
- 2 Enter the cluster name or browse and select it, then check the box next to the resource you want to migrate and click *Migrate*.
A page appears, displaying a list of possible servers that you can migrate this resource to.
- 3 Select a server from the list to migrate the resource to, then click *OK* to migrate the resource to the selected server.

If you select a resource and click *Offline*, the resource will be unloaded from the server. It will not load on any other servers in the cluster and will remain unloaded until you load it again. This option is useful for editing resources because resources can't be edited while loaded or running on a server.

4.2 Identifying Cluster and Resource States

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

Cluster servers and resources display in different colors, depending on their operating state. When servers and resources are display with a green ball, they are in a normal (online or running) operating condition. When a server that has been part of the cluster has a red X in the icon, it has failed. When a resource is red, it is waiting for administrator intervention. When a server is gray with no break in the icon, either that server is not currently a member of the cluster or its state is unknown. When a resource is blank or has no colored icon, it is unassigned, offline, changing state, or in the process of loading or unloading.

The yellow diamond in the middle of the server icon designates the master server in the cluster. The master server is initially the first server in the cluster, but another server can become the master if the first server fails.

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

The following table identifies the different resource states and gives descriptions and possible actions for each state. In iManager, click Clusters, then click Cluster Manager and enter the name of the desired cluster. A list of resources and resource states will display.

Table 4-1 Cluster Resource States

Resource State	Description	Possible Actions
Alert	Either the Start, Failover, or Failback mode for the resource has been set to Manual. The resource is waiting to start, fail over, or fail back on the specified server.	Click the <i>Alert</i> status indicator. Depending on the resource state, you will be prompted to start, fail over, or fail back the resource.
Comatose	The resource is not running properly and requires administrator intervention.	Click the <i>Comatose</i> status indicator and bring the resource offline. After resource problems have been resolved, the resource can be brought back online (returned to the running state).
Loading	The resource is in the process of loading on a server.	None.
NDS_Sync	The properties of the resource have changed and the changes are still being synchronized in Novell eDirectory™.	None.
Offline	Offline status indicates the resource is shut down or is in a dormant or inactive state.	Click the <i>Offline</i> status indicator and, if desired, click the <i>Online</i> button to load the resource on the best node possible, given the current state of the cluster and the resource's preferred nodes list.
Quorum Wait	The resource is waiting for the quorum to be established so it can begin loading.	None.
Running	The resource is in a normal running state.	Click the <i>Running</i> status indicator and choose to either migrate the resource to a different server in your cluster or unload (bring offline) the resource.
Unassigned	There isn't an assigned node available that the resource can be loaded on.	Click the <i>Unassigned</i> status indicator and, if desired, offline the resource. Offlining the resource will prevent it from running on any of its preferred nodes should any of them join the cluster.
Unloading	The resource is in the process of unloading from the server it was running on.	None.

4.3 Novell Cluster Services Console Commands

Novell Cluster Services provides several server console commands to help you perform certain cluster-related tasks. The following table lists the cluster-related server console commands and gives a brief description of each command. 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 server console. You can also enter `cluster help` at the console prompt to get information on the commands and their functions. The functions of many of the commands can also be performed using iManager. See the other sections of this document for additional information.

Table 4-2 Cluster Console Commands

Cluster Console Command	Description
ALERT {resource}{YES NO}	The resource start, failover, or failback mode is set to manual and the resource is waiting to start on a node, or fail over or fail back to another node. Specify the resource name in the command and use the YES or NO switch to specify whether you want the resource to fail over, fail back, or start.
CONVERT {Preview, Commit}{Resource}	Finalizes the cluster conversion from NetWare® to Linux after all nodes in a mixed cluster have been converted to Linux. Specify a resource name with the Preview switch to view the resource load and unload script changes prior to finalizing the conversion. Use the Commit switch without specifying a resource to finalize the conversion for all cluster resources. The <code>CLUSTER CONVERT</code> command can only be executed on Linux cluster nodes.
DOWN	Removes all cluster nodes from the cluster. Has the same effect as executing the <code>CLUSTER LEAVE</code> command on every server in the cluster.
INFO {All, Basic, Notification, Priority, Protocol, Summary}	Displays information on cluster configuration. All displays a combination of Basic, Notification, Priority, and Protocol information. Basic displays IP address, port, and cluster quorum settings. Notification displays cluster e-mail notification settings. Priority displays the resource priority list. Protocol displays the cluster protocol settings. Summary displays the cluster protocol summary.
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 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.

Cluster Console Command	Description
MAINTENANCE {ON OFF}	Turning this switch on lets you temporarily suspend the cluster heartbeat while hardware maintenance is being performed. This is useful if you want to reset or power down the LAN switch without bringing the cluster servers down. Turning this switch on from one cluster server puts the entire cluster in maintenance mode.
MIGRATE {resource}{node name}	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.
OFFLINE {resource}	Unloads the specified resource from the node where it is currently running.
ONLINE {resource}{node name}	Starts the specified resource on the most preferred node that is currently active. You can start the resource on a different node by specifying that node in the command.
POOLS	Lists the NSS pools on the shared disk system that are accessible by Novell Cluster Services.
RESOURCES	Lists all resources that currently exist in the cluster. The resources do not need to be online or running.
RESTART {seconds}	Restarts Novell Cluster Services software on all servers in the cluster.

Cluster Console Command	Description
SET {Parameter} {Value}	<p>Sets cluster parameters individually for the cluster. See Section 3.8, "Configuration Settings," on page 44 for more information on cluster parameters.</p> <p>Specify one of the following parameters and a value for that parameter:</p> <p>IPADDRESS Sets the cluster IP address to the specified value. If you change the cluster IP address, you must restart cluster software on all cluster nodes.</p> <p>PORT Sets, or lets you change the cluster port number.</p> <p>QUORUMWAIT This is the amount of time in seconds that the cluster waits before resources start to load.</p> <p>QUORUM This is the number of nodes that must be running in the cluster before resources will start to load.</p> <p>HEARTBEAT This is the amount of time in seconds between transmits for all nodes in the cluster except the master.</p> <p>TOLERANCE This is the amount of time in seconds that the master node gives all other nodes in the cluster to signal that they are alive.</p> <p>MASTERWATCHDOC This is the amount of time in seconds between transmits for the master node in the cluster.</p> <p>SLAVEWATCHDOC This is the amount of time in seconds that the slave nodes give the master node in the cluster to signal that it is alive.</p> <p>MAXRETRANSMITS This is the maximum number of times transmits will be attempted between the master node and slave nodes.</p> <p>ENABLEEMAIL Enables and disables e-mail notification. You can set the value to OFF to disable e-mail notification, or either CRITICAL or VERBOSE to enable e-mail notification.</p> <p>EMAILADDRESSES Lets you specify the e-mail addresses used for e-mail notification. The addresses should be separated by spaces. Using this parameter without specifying any addresses will clear existing addresses that have been set previously.</p> <p>EMAILOPTIONS Sets the e-mail notification options. Specify XML as the value to receive e-mail notification in XML format. Not specifying any value with this parameter will turn notification in XML format off.</p>
STATS {Display, Clear}	<p>Reports the node number, node name, and heartbeat information. You must switch to the log console screen to see cluster statistics.</p>

Cluster Console Command	Description
STATUS {resource}	Reports the status of the specified resource. This includes the number of times the resource has been migrated or failed over to another server, the resource state, and the node where the resource is currently running.
VIEW	Displays the node name, cluster epoch number, master node name, and a list of nodes that are currently members of the cluster.

4.4 Customizing Cluster Services Management

Some portions of Novell Cluster Services management can be performed and customized using virtual XML files that exist on the `_admin` volume.

The cluster-related virtual XML files (management access points) are created on each server's `_admin` volume. 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 `\\cluster/_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.

The following table lists the cluster-related virtual XML files and gives a brief description of each.

Table 4-3 Cluster-Related Virtual XML Files

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

Virtual XML Filename	Description
<code>Node.cmd</code>	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.
<code>Cluster.cmd</code>	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.
<code>Resource.cmd</code>	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.

4.5 Novell Cluster Services File Locations

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. The following table lists the path and purpose for some of the files that are part of Novell Cluster Services (NCS).

Table 4-4 *Novell Cluster Services File Locations*

NCS File Name and Path	Purpose
<code>/etc/init.d/novell-ncs</code>	LSB Compliant Service
<code>/etc/opt/novell/ncs/nodename</code>	This node's name
<code>/lib/evms/2.3.3/ncs-1.0.0.so</code>	EVMS snap-in
<code>/opt/novell/ncs/bin/ClusterCli.pl</code>	Cluster CLI Engine
<code>/opt/novell/ncs/bin/ClusterCliSnapinInterface.pm</code>	Cluster CLI Engine
<code>/opt/novell/ncs/bin/ClusterCliUtils.pm</code>	Cluster CLI Engine
<code>/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Alert.pm</code>	Cluster CLI Command
<code>/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Down.pm</code>	Cluster CLI Command
<code>/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Info.pm</code>	Cluster CLI Command
<code>/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Join.pm</code>	Cluster CLI Command
<code>/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Leave.pm</code>	Cluster CLI Command
<code>/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Maintenance.pm</code>	Cluster CLI Command
<code>/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Migrate.pm</code>	Cluster CLI Command

NCS File Name and Path	Purpose
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Offline.pm	Cluster CLI Command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Online.pm	Cluster CLI Command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Pools.pm	Cluster CLI Command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Resources.pm	Cluster CLI Command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Restart.pm	Cluster CLI Command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Set.pm	Cluster CLI Command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Stats.pm	Cluster CLI Command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_Status.pm	Cluster CLI Command
/opt/novell/ncs/bin/Snapins/ClusterCliSnapin_View.pm	Cluster CLI Command
/opt/novell/ncs/bin/adminfs	Cluster management (iManager and CLI)
/opt/novell/ncs/bin/ldncs	Loads NCS, used by Cluster Start command
/opt/novell/ncs/bin/ncs-configd.py	Cluster configuration daemon
/opt/novell/ncs/bin/ncs-emaild	Cluster e-mail daemon
/opt/novell/ncs/bin/ncs-resourced.py	Daemon used to run load and unload scripts
/opt/novell/ncs/bin/ncstempl.py	Used to install cluster resource templates
/opt/novell/ncs/bin/sbdutil	SBD partition utility
/opt/novell/ncs/bin/uldncs (not yet implemented)	Unloads NCS, Used by Cluster Stop command
/opt/novell/ncs/lib/ncs-1.0.0.so	EVMS snap-in
/opt/novell/ncs/lib/ncsfuncs	Shared library commands for load/unload scripts
/opt/novell/ncs/schema/ncs.ldif	NCS Schema file
/opt/novell/ncs/schema/ncs.sch	NCS Schema file
/usr/include/ncssdk.h	NCS SDK
/usr/lib/libncssdk.so	NCS SDK
/usr/lib/libncssdk.so.1.0.0	NCS SDK
/usr/sbin/rcnovell-ncs	Link to etc/init.d/novell-ncs
/usr/share/man/man7/sbdutil.7.gz	SBDUTIL Man page

NCS File Name and Path	Purpose
<code>/var/opt/novell/ncs/hangcheck-timer.conf</code>	hang check option (comparable to CPU hog)
<code>/lib/modules/kernel_dir/ncs/clstrlib.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/cma.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/cmsg.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/crm.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/css.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/cvb.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/gipc.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/sbd.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/sbdlib.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/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.

4.6 Additional Cluster Operating Instructions

The following instructions provide additional information for operating Novell Cluster Services.

4.6.1 Connecting to an iSCSI Target

For instructions on configuring an OES Linux server as an iSCSI initiator and connecting to an iSCSI target, go to “[Accessing iSCSI Targets on NetWare Servers from Linux Initiators](#)” in the *iSCSI 1.1.3 Administration Guide for NetWare 6.5*.

If you are connecting to an iSCSI target that already has NSS partitions and pools created on it, you may not be able to access those NSS partitions and pools until you either reboot the Linux initiator server or run the `evms_activate` command at the Linux server console. This is required for each Linux initiator server that will access the iSCSI target.

4.6.2 Adding a Node That Was Previously in the Cluster

- 1 If necessary, install NetWare and Novell Cluster Services, including the latest Service Pack on the server using the same node name and IP address.

If your SAN is not configured, install Novell Cluster Services after configuring the SAN.

- 2 If the Cluster object for the server is still present, use ConsoleOne® to delete the object.

You can do this by going to the Cluster container, selecting the node in the right frame, and pressing Delete.

- 3 Run the Novell Cluster Services installation.

The node will assume its former identity.

4.6.3 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. See [Section 4.3, “Novell Cluster Services Console Commands,” on page 50](#) for more information.

If the master server in the cluster goes down while the cluster is in cluster maintenance mode, you must enter `cluster maintenance off` on all remaining cluster servers to bring the cluster out of maintenance mode. This is only necessary if the master server in the cluster goes down. If the master server in the cluster is up, you can enter `cluster maintenance off` on one server in the cluster to bring the entire cluster out of maintenance mode.

4.6.4 Shutting Down Linux Servers When Servicing Shared Storage

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

4.6.5 Preventing Cluster Node Reboot after Node Shutdown

If LAN connectivity is lost between a cluster node and the other nodes in the cluster, it is possible that the lost node will be automatically shut down by the other cluster nodes. This is normal cluster operating behavior, and it prevents the lost node from trying to load cluster resources because it cannot detect the other cluster nodes.

By default, cluster nodes are configured to reboot after an automatic shutdown. On certain occasions, you might want to prevent a downed cluster node from rebooting so you can troubleshoot problems. To do this, edit the `opt/novell/ncs/bin/ldncs` file and find the following line:

```
echo -n $TOLERANCE > /proc/sys/kernel/panic
```

Replace `$TOLERANCE` with `0` to cause the server to not automatically reboot after a shutdown. After editing the `ldncs` file, you must reboot the server to cause the change to take effect.

4.6.6 Problems Authenticating to Remote Servers during Cluster Configuration

If, during the OES cluster installation and configuration, you choose Remote System on the NCS LDAP Configuration page and you have LDAP configured to point to a NetWare 6.0 or earlier NetWare server, the cluster configuration will fail. To work around this problem, you must edit the `etc/openldap/ldap.conf` file and 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.

4.6.7 Reconfiguring a Cluster Node

If you want to remove a cluster node from one cluster and add it to another cluster, you must reconfigure the node.

- 1 Bring down the cluster node.
- 2 In eDirectory, delete the cluster node object from the cluster container.
- 3 In eDirectory, find the NCP Server object and then go to Properties for that object.
This is not the cluster node object in the cluster container. It is the server object for the server.
- 4 Click *Other*, then find and delete the NCS:NetWare Cluster attribute.
- 5 Reconfigure Novell Cluster Services by following the procedure outlined in [Section 3.5.3, “Installing Novell Cluster Services after the OES Installation,”](#) on page 20.

4.6.8 Device Name Required to Create a Cluster Partition

If you are planning to work with shared-disk NSS Pools and Volumes, you must install a shared-disk cluster by entering a device name for the cluster Split Brain Detector (SBD) partition at cluster creation (new cluster) time. If you don't enter a device name, you won't be able to cluster enable NSS Pools.

4.6.9 Creating a Cluster Partition (SBD Partition) after Installation

If you did not create a cluster partition during the Novell Cluster Services installation, you can create one later using the `SBDUTIL` utility. You must have a shared disk system (a Storage Area Network or SAN) connected to your cluster nodes before attempting to create a cluster partition. See [Section 3.3, “Shared Disk System Requirements,”](#) on page 18 for more information.

Before creating an SBD partition, you should first make sure one does not already exist on your shared disk system. To do this, enter `sbduutil -f` at the server console of a cluster node.

If a cluster partition already exists, do not create another one. If a cluster partition does not exist, enter `sbduutil -c -d device_name` at the server console. Replace `device_name` with the name of the device where you want to create the cluster partition.

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

```
sbduutil -c -d /dev/sda
```

See the man page for `sbduutil` for more information on how to use it.

After creating the SBD partition, you must edit the Cluster object in eDirectory and enable the Shared Disk Flag attribute. You must then save changes and reboot the cluster. To do this

- 1 Start iManager, click *eDirectory Administration*, then click *Modify Object*.
- 2 Enter the Cluster object name, or browse and select it, then click *OK*.
- 3 Under Valued Attributes, click *NCS:Shared Disk Flag*, then click *Edit*.
- 4 Check the *NCS:Shared Disk Flag* check box, then click *OK*.
- 5 Click *Apply* to save changes, then reboot the cluster.

4.6.10 Mirroring SBD (Cluster) Partitions

To achieve a greater level of fault tolerance, you can mirror SBD partitions. You must use the `evmsgui` utility to create and mirror SBD partitions. If an SBD partition was created either during the Novell Cluster Services installation or later using the `sbdutil` command, you must delete that partition prior to creating and mirroring SBD partitions using `evmsgui`. To see if an SBD partition already exists, enter `sbdutil -f` at the server console of a Linux cluster server. (See [Step 9 on page 20](#) for more information on SBD partitions.)

If an SBD partition was created during the Novell Cluster Services installation or later using the `sbdutil` command, delete it.

- 1 Enter `cluster down` at the server console of one cluster server.
This will cause all cluster servers to leave the cluster.
- 2 Delete the SBD partition.
You can use `nssmu`, `evmsgui`, or other utilities to delete the SBD partition.

To create an SBD partition using `evmsgui`:

- 1 At the Linux server console of a cluster server, enter `evmsgui` to start the `evmsgui` utility.
- 2 Click *Action*, then click *Create*.
- 3 Click *Segment*, choose the *NetWare Segment Manager*, then click *Next*.
- 4 Select *Free Space Storage Object*, then click *Next*.
- 5 Specify 8 MB as the size of the cluster partition, then choose *SBD* as the partition type.
- 6 Enter the name of your cluster as the Label, then click *Create*.

If necessary, repeat the above steps to create a second SBD partition.

To mirror SBD partitions:

- 1 At the Linux server console of a cluster server, enter `evmsgui` to start the `evmsgui` utility.
- 2 Click *Segments*.
- 3 Locate one SBD partition and right-click it.
- 4 Select *Mirror Segment*, then click *OK*.
- 5 Reboot all cluster nodes.

Documentation Updates

A

This *Novell Cluster Services 1.8.2 Administration Guide for Linux* has been updated with the following information on December 23, 2005:

A.1 December 23, 2005 (Open Enterprise Server SP2)

Location	Change
Entire guide.	Page design reformatted to comply with revised Novell® documentation standards.
Section 3.5, "Installing Novell Cluster Services," on page 18	It is now possible to choose a device for the SBD partition from a list rather than entering it in manually.
Sections covering IManager cluster configuration and management	Some iManager cluster option names and locations have changed. These changes are reflected in several locations in the documentation.
Section 3.6, "Converting a NetWare Cluster to Linux," on page 21	It is now possible to upgrade a cluster node directly from NetWare® 6.0 to OES Linux without first upgrading to NetWare 6.5. Information about this has been added to this section.
