Novell Cluster Services™ 1.8.8 Administration Guide for Linux
Open Enterprise Server 2 SP3

December 2013

Novell®

Contents

Contents   **7**

About This Guide

This guide describes how to install, upgrade, configure, and manage Novell Cluster Services for Novell Open Enterprise Server (OES) 2 Linux or later. It is divided into the following sections:

Audience

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

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

Feedback

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

Documentation Updates

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

Additional Documentation

For information about converting clusters from NetWare to Linux, see the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*.

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

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

1 Overview of Novell Cluster Services

Novell Cluster Services is a server clustering system that ensures high availability and manageability of critical network resources including data, applications, and services. It is a multi-node clustering product for Novell Open Enterprise Server 2 Linux. It stores information in Novell eDirectory about the cluster and its resources. It supports failover, failback, and cluster migration of individually managed cluster resources. You can cluster migrate a resource to a different server in order to perform rolling cluster maintenance or to balance the resource load across member nodes.

1.1 Why Should I Use Clusters?

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

Each server in the cluster runs the same operating system and applications that are needed to provide the application, service, or data resources to clients. Shared devices are connected to and mounted on only one server at a time. Clustering software monitors the health of each of the member servers by listening for its heartbeat, a simple message that lets the others know it is alive.

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

1.2 Benefits of Novell Cluster Services

Novell Cluster Services provides high availability for data and services running on OES 2 servers. You can configure up to 32 OES 2 Linux servers in a high-availability cluster, where resources can be dynamically relocated to any server in the cluster. Each cluster resource can be configured to

automatically fail over to a preferred server in the event of a failure on the server where it is currently running. In addition, costs are lowered through the consolidation of applications and operations onto a cluster.

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

Novell Cluster Services clusters provide the following benefits over standalone servers:

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

1.3 Product Features

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

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

1.4 Clustering for High-Availability

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

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

◆ Equipment to connect servers to the shared disk subsystem, such as one of the following:

- ◆ High-speed Fibre Channel cards, cables, and switches for a Fibre Channel SAN
- ◆ Ethernet cards, cables, and switches for an iSCSI SAN
- ◆ SCSI cards and cables for external SCSI storage arrays

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

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

Figure 1-1   *Three-Server Cluster*
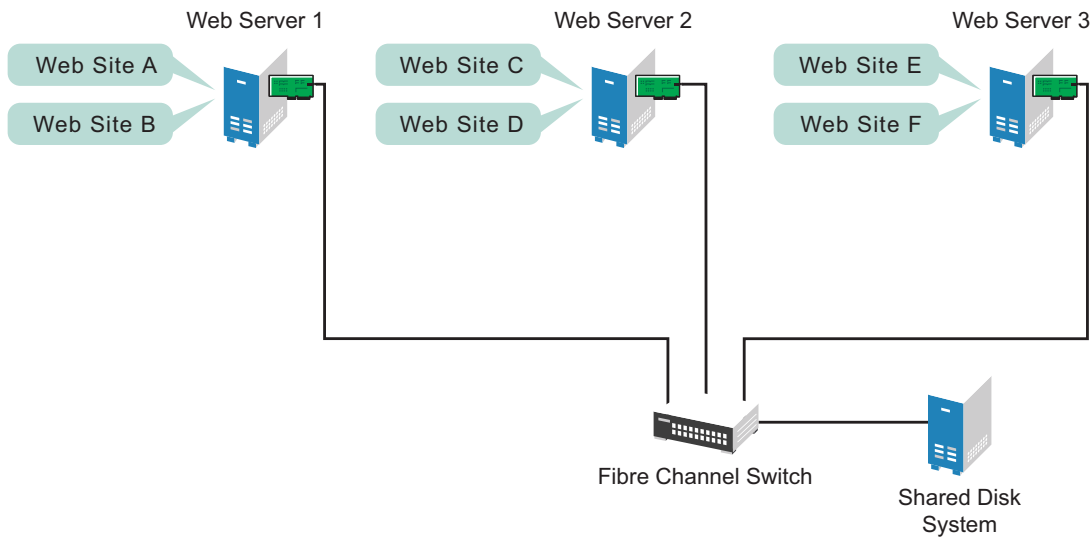
During normal cluster operation, each server is in constant communication with the other servers in the cluster and performs periodic polling of all registered resources to detect failure.

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

Figure 1-2   *Three-Server Cluster after One Server Fails*
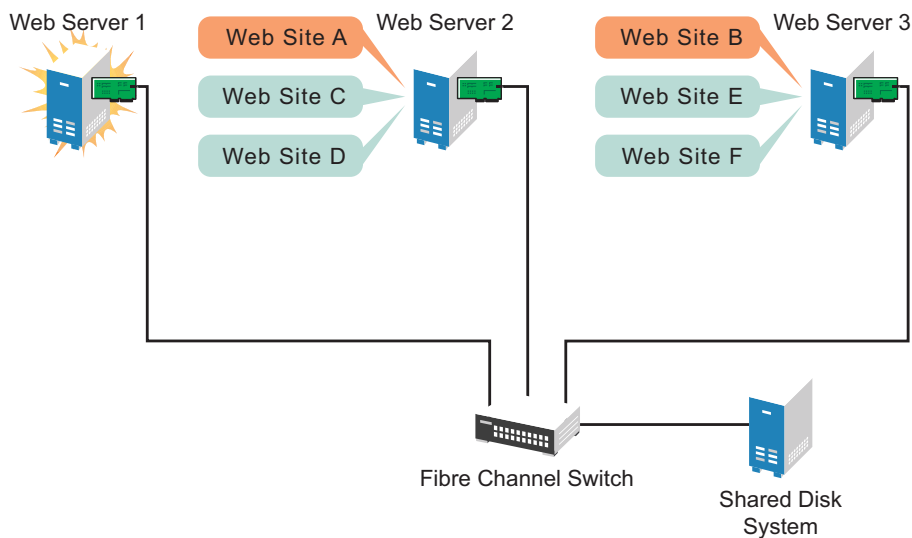
Web Site A moves to Web Server 2 and Web Site B moves to Web Server 3. IP addresses and certificates also move to Web Server 2 and Web Server 3.

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

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

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

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

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

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

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

1.5 Shared Disk Scenarios

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

1.5.1 Using Fibre Channel Storage Systems

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

Figure 1-3 *Typical Fibre Channel Cluster Configuration*
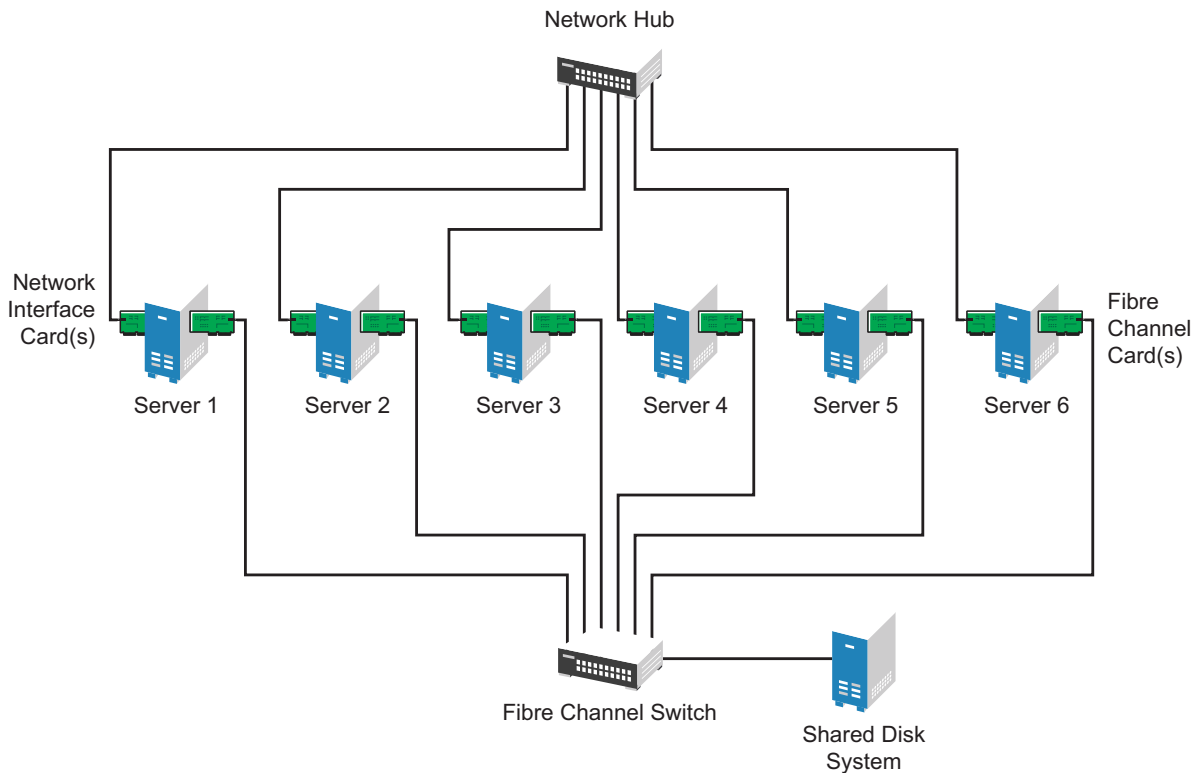
1.5.2 Using iSCSI Storage Systems

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

Figure 1-4   *Typical iSCSI Cluster Configuration*
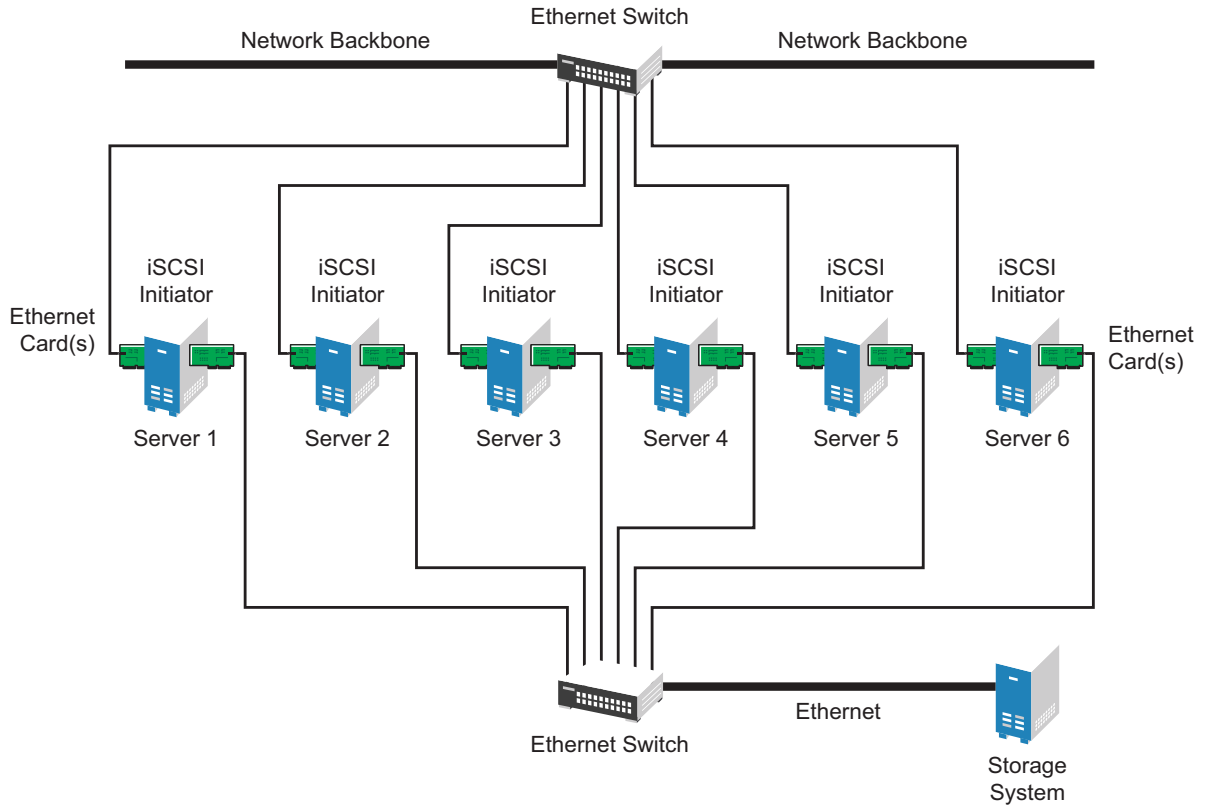
1.5.3 Using Shared SCSI Storage Systems

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

Figure 1-5 *Typical Shared SCSI Cluster Configuration*
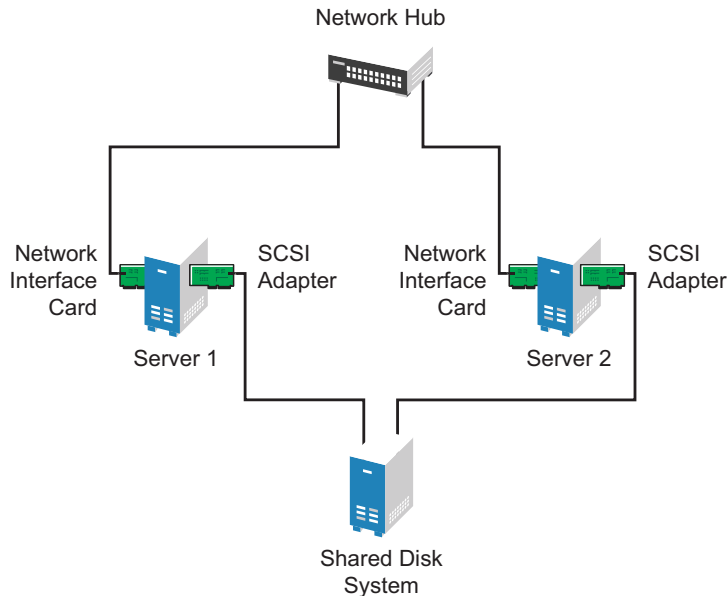
1.6 Terminology

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

1.6.1 The Cluster

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

Cluster IP Address

The unique static IP address for the cluster.

Server IP Address

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

Master Node

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

Slave Node

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

Split-Brain Detector (SBD)

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

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

Shared Storage

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

1.6.2 Cluster Resources

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

-
-

Resource IP Address

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

NCS Virtual Server

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

Resource Templates

A resource template contains the default load, unload, and monitoring scripts and default settings for service or file system cluster resources. Resource templates are available for the following OES 2 Services and file systems.

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

Personalized templates can also be created. For information, see .

Service Cluster Resource

An application or OES 2 service that has been cluster-enabled. The application or service is installed on all nodes in the cluster where the resource can be failed over. The cluster resource includes scripts for loading, unloading, and monitoring. The resource can also contain the configuration information for the application or service.

Pool Cluster Resource

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

Linux POSIX Volume Cluster Resource

A cluster-enabled Linux POSIX volume. The volume is bound to the resource IP address. This provides location independent access to data on the volume via native Linux protocols such as Samba or FTP.

NCP Volume Cluster Resource

An NCP volume (or share) that has been created on top of a cluster-enabled Linux POSIX volume. The NCP volume is re-created by a command in the resource load script whenever the resource is brought online. The NCP volume is bound to an NCS Virtual Server object (`NCS:NCP Server`) and to the resource IP address. This provides location independent access to the data on the volume for NCP clients in addition to the native Linux protocols such as Samba or FTP.

DST Volume Cluster Resource

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

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

Cluster Resource Scripts

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

1.6.3 Failover Planning

Heartbeat

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

Quorum

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

Preferred Nodes

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

Resource Priority

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

Resource Mutual Exclusion Groups

Administrator-specified groups of resources that should not be allowed to run on the same node at the same time. This capability is available in OES 2 SP3 and later.

Failover

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

Fan-Out Failover

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

Failback

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

Cluster Migrate

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

Leave a Cluster

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

Join a Cluster

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

2 What's New or Changed for Novell Cluster Services

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

2.1 What's New (May 2013 Patches)

In addition to bug fixes, Novell Cluster Services provides the following enhancement and behavior changes in the June 2013 Scheduled Maintenance for OES 2 SP3:

SLP Refresh Interval for Cluster Resource Virtual NCP Servers

NCP Server was modified to refresh its OpenSLP registration of cluster resource virtual NCP servers based on the setting for the eDirectory advertise-life-time (`n4u.nds.advertise-life-time`) parameter. The `n4u.nds.advertise-life-time` parameter is set by default to 3600 seconds (1 hour) and has a valid range of 1 to 65535 seconds. Previously, NCP Server re-registered the virtual NCP servers with SLP every 30 minutes regardless of the eDirectory advertise-life-time setting. For information about setting the eDirectory advertise-life-time parameter in a cluster, see Section 4.5.3, "SLP," on page 61.

Updated the DHCP PID File Location in the DHCP_Template

In the DHCP_Template for DHCP cluster resources, the PID file location was changed to `/var/lib/dhcp/var/run/dhcpd.pid`. The change applies automatically to any newly created DHCP cluster resources. For information about configuring DHCP cluster resources, see "Configuring DHCP with Novell Cluster Services for the Linux File System" in the *OES 2 SP3: Novell DNS/DHCP Administration Guide*.

2.2 What's New (April 2013 Patches)

Upgrade to eDirectory 8.8.7

An upgrade to Novell eDirectory 8.8 SP7 is available in the April 2013 Scheduled Maintenance for OES 2 SP3. For information about the eDirectory upgrade, see TID 7011599 (http://www.novell.com/support/kb/doc.php?id=7011599) in the Novell Knowledgebase.

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

2.3 What's New (January 2013 Patches)

Upgrade to Novell iManager 2.7.6

The January 2013 Scheduled Maintenance for OES 2 SP3 includes a channel upgrade from Novell iManager 2.7.5 to Novell iManager 2.7.6.

Novell iManager 2.7.6 provides the following enhancements:

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

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

iManager 2.7.6 documentation is available on the Web (https://www.netiq.com/documentation/imanager/). For earlier iManager versions, see "Previous Releases" (https://www.netiq.com/documentation/imanager27/#prev).

Novell Client Support for Windows 8 and Server 2012

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

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

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

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

Clusters Plug-In Changes for Novell iManager 2.7.5 or Later

The Clusters plug-in for Novell iManager 2.7.5 or later was released in OES 11 SP1. It supports the management of OES and NetWare clusters and resources. The availability of Novell Cluster Services features depends on the version of Novell Cluster Services and the server platform that are installed on the cluster being managed. You can use either the old plug-in or the new plug-in to manage an OES 2 SP3 cluster and its resources.

Understanding the Cluster Plug-In's New Interface

The Clusters plug-in for Novell iManager 2.7.5 or later has been reorganized. In *Roles and Tasks* under *Clusters*, the *Cluster Manager*, *BCC Manager*, *Cluster Event Log*, and *Cluster Options* menu items have been replaced with two menu options:

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

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

From a customized My Clusters page or a My Resources page, you can click the name link of a cluster to manage the cluster. The old plug-in's *Clusters* menu options are available here as tabs: *Cluster Manager*, *BCC Manager*, *Cluster Event Log*, and *Cluster Options*.

Upgrading the Clusters Plug-In

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

Comparing Tasks in the Old and New Clusters Plug-Ins

Use the tables in this section to understand how to perform tasks in the old Clusters plug-in for iManager 2.7.4 and the new Clusters plug-in for iManager 2.7.5.

- Table 2-1, "Selecting a Cluster to Manage," on page 32

Table 2-1 describes how to select the cluster you want to manage.

Table 2-1 *Selecting a Cluster to Manage*

Clusters Plug-In for iManager 2.7.4 and Earlier	Clusters Plug-In for iManager 2.7.5 and Later

1. In *Roles and Tasks*, select *Clusters*, then select *Cluster Manager*, *BCC Manager*, *Cluster Event Log*, or *Cluster Options*.
 	
2. Specify the cluster name, or browse and select the Cluster object.
 	
You must repeat the selection process to manage a different cluster. There is no capability to concurrently view the status of multiple clusters.
	1. In *Roles and Tasks*, select *Clusters > My Clusters*.
	
	The list is initially empty.
	
	2. On the My Clusters page, click *Add* to open the Novell eDirectory browser pop-up window.
	
	3. Browse the tree where you are currently logged in to locate and select one or more clusters that you want to manage, then click *OK*.
	
	The selected clusters are added to the list.
	
	Your personalized list is displayed each time you log in to this iManager server and return to the My Clusters page. Use the *Add* and *Remove* options to modify your list as needed.
	
	4. Do any of the following from this page:
	
	   ◆ Concurrently view the status of multiple clusters.
	
	   ◆ Click the name link of a cluster to manage it.
	
	     The Cluster Manager page opens. You can also select the other tabs from this page: *BCC Manager*, *Cluster Event Log*, or *Cluster Options*.
	
	   ◆ Select the check box next to the cluster, then select *Actions > Edit Properties* to manage the cluster properties.
	
	   ◆ Select the check box next to the cluster, then select *Actions > Run Report*.

Table 2-2 describes how to access the cluster management pages: Cluster Manager, BCC Manager, Cluster Event Log, and Cluster Options. The instructions for the Clusters plug-in for iManager 2.7.5 and later assume that you have created a personalized list of clusters, as described in Table 2-1.

Table 2-2 *Managing a Cluster*

Clusters Plug-In for iManager 2.7.4 and Earlier	Clusters Plug-In for iManager 2.7.5 and Later
Cluster Manager 1. In *Roles and Tasks*, select *Clusters > Cluster Manager*. 2. Specify the cluster name, or browse and select the Cluster object.	**Cluster Manager** 1. In *Roles and Tasks*, select *Clusters > My Clusters*. 2. Click the name link of the cluster. 3. Select the *Cluster Manager* tab.
BCC Manager 1. In *Roles and Tasks*, select *Clusters > BCC Manager*. 2. Specify the cluster name, or browse and select the Cluster object.	**BCC Manager** 1. In *Roles and Tasks*, select *Clusters > My Clusters*. 2. Click the name link of the cluster. 3. Select the *BCC Manager* tab.
Cluster Event Log 1. In *Roles and Tasks*, select *Clusters > Cluster Event Log*. 2. Specify the cluster name, or browse and select the Cluster object.	**Cluster Event Log** 1. In *Roles and Tasks*, select *Clusters > My Clusters*. 2. Click the name link of the cluster. 3. Select the *Cluster Event Log* tab.
Cluster Options 1. In *Roles and Tasks*, select *Clusters > Cluster Options*. 2. Specify the cluster name, or browse and select the Cluster object.	**Cluster Options** 1. In *Roles and Tasks*, select *Clusters > My Clusters*. 2. Click the name link of the cluster. 3. Select the *Cluster Options* tab.

Table 2-3 describes how to access the cluster properties pages: Policies, Priorities, Protocols, Resource Mutual Exclusion, and Business Continuity. The instructions for the Clusters plug-in for iManager 2.7.5 and later assume that you have created a personalized list of clusters, as described in Table 2-1.

Table 2-3 *Accessing Cluster Properties*

Clusters Plug-In for iManager 2.7.4 and Earlier	Clusters Plug-In for iManager 2.7.5 and Later
	From My Resources 1. In *Roles and Tasks*, select *Clusters > My Clusters*. 2. Select the check box next to the cluster, then click *Actions > Edit Properties*.
From Cluster Options 1. In *Roles and Tasks*, select *Clusters > Cluster Options*. 2. Specify the cluster name, or browse and select the Cluster object. 3. Click *Properties*.	**From Cluster Options** 1. In *Roles and Tasks*, select *Clusters > My Clusters*. 2. Click the name link of the cluster. 3. Select the *Cluster Options* tab. 4. Click *Properties*.

Table 2-4 describes how to view or manage cluster resources, including their properties: Policies, Monitoring, Preferred Nodes, Scripts (Load, Unload, and Monitor), and Business Continuity.

Table 2-4   *Managing a Cluster Resource or Resource Properties*

Clusters Plug-In for iManager 2.7.4 and Earlier	Clusters Plug-In for iManager 2.7.5 and Later

There is no capability to concurrently view the status of cluster resources across multiple clusters.	**From My Resources:**
	
	1. In *Roles and Tasks*, select *Clusters > My Resources*.
	
	   The initial list is empty.
	
	2. On the My Resources page, click *Add* to open the Novell eDirectory browser pop-up window.
	
	3. Browse the tree where you are currently logged in to locate and select the cluster resources that you want to manage for one or more clusters, then click *OK*.
	
	   The selected cluster resources are added to the list.
	
	   Your personalized list is displayed each time you log in to this iManager server and return to the My Resources page. Use the *Add* and *Remove* options to modify your list as needed.
	
	4. Do any of the following from this page:
	
	   ◆ Concurrently view the status of multiple resources across multiple clusters.
	
	   ◆ Click the name link of a resource to open its *Resource Properties*.
	
	   ◆ Click the name link of the resource's cluster to manage the cluster.

From Cluster Manager:
 	
1. In *Roles and Tasks*, select *Clusters > Cluster Manager*.
 	
2. Specify the cluster name, or browse and select the Cluster object.
 	
3. Do any of the following:
 	
   ◆ View the status of resources on this cluster.
 	
   ◆ Click the name link of a resource to open its *Resource Properties*.
 	
   ◆ Select the check box next to a resource, then click one of the following options:
 	
   Online
   Offline
   Migrate
   Respond to Alert
	From Cluster Manager:
	
	1. In *Roles and Tasks*, select *Clusters > My Clusters*.
	
	2. Click the name link of the cluster.
	
	3. Select the *Cluster Manager* tab.
	
	4. Do any of the following:
	
	   ◆ View the status of resources on this cluster.
	
	   ◆ Click the name link of a resource to open its *Resource Properties*.
	
	   ◆ Select the check box next to a resource, then click one of the following options:
	
	   Online
	   Offline
	   Migrate
	   Respond to Alert

Clusters Plug-In for iManager 2.7.4 and Earlier	Clusters Plug-In for iManager 2.7.5 and Later
From Cluster Options: 1. In *Roles and Tasks*, select *Clusters > Cluster Options*. 2. Specify the cluster name, or browse and select the Cluster object. 3. Click the name link of a resource to open its *Resource Properties*. You can also select the check box next to the resource, then click the *Details* link.	**From Cluster Options:** 1. In *Roles and Tasks*, select *Clusters > My Clusters*. 2. Click the name link of the cluster. 3. Select the *Cluster Options* tab. 4. Click the name link of a cluster resource to open its *Resource Properties*. You can also select the check box next to the resource, then click the *Details* link.

2.4 What's New (September 2012 Patches)

In addition to bug fixes, the OES 2 SP3 January 2012 Scheduled Maintenance patch provides the following enhancement and behavior changes for Novell Cluster Services:

Master Election Process

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

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

Updating to the Clusters Plug-In for Novell iManager 2.7.5

The Clusters plug-in for Novell iManager 2.7.5 in OES 11 SP1 supports management of OES 2 SP3 clusters. The plug-in was reorganized to display two tasks in the left panel: *My Clusters* and *My Resources*. Ensure that you perform the following tasks after updating the plug-in:

- Step 5 in Section 5.7.3, "Uninstalling and Reinstalling the Clusters Plug-In," on page 93
- Section 5.7.4, "Updating Role-Based Services for the Clusters Plug-In for OES 11 SP1," on page 94

For information about using the new Clusters plug-in, see the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide* (http://www.novell.com/documentation/oes11/clus_admin_lx/data/h4hgu4hs.html).

2.5 What's New (January 2012 Patches)

In addition to bug fixes, the OES 2 SP3 January 2012 Scheduled Maintenance patch provides the following changes for Novell Cluster Services:

- **Case Insensitive Resource Name for the Monitor Command:** The `cluster monitor` command was modified so that the cluster resource name that you provide is treated as case insensitive. Resource names should be case insensitive for all Novell Cluster Services commands.

For information about using `cluster` commands, see Section A.1, "Cluster Management Commands," on page 263.

◆ **Editing Preferred Nodes:** Edits are not allowed to be made to the Preferred Nodes list for the `Master_IP_Address_Resource`. The Edit button is now disabled on the master resource's Preferred Nodes page, which is in keeping with the disabled up, down, left, and right arrows on the page.

For information, see Section 10.8, "Configuring Preferred Nodes for a Resource," on page 160.

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

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

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

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

For information, see Section 12.4, "Cluster-Enabling an Existing NSS Pool and Its Volumes," on page 184.

◆ **NetWare to Linux Conversion:** The NetWare to Linux conversion of Novell Cluster Services clusters is supported from NetWare 6.5 SP8 (with the latest patches applied) to OES 2 SP3 on the SUSE Linux Enterprise 10 SP4 operating system.

For information, see *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*.

2.6   What's New (November 2011 Patches)

In addition to bug fixes, the OES 2 SP3 November 2011 Scheduled Maintenance patch provides the following change for Novell Cluster Services:

◆ Novell CIFS provides a `monitor` command option in this patch for OES 2 SP3 that provides a restart capability if the `cifsd` daemon goes down. If you create a new pool cluster resource with CIFS enabled as an advertising protocol, the following line is added to the resource's monitor script.

```
exit_on_error rcnovell-cifs monitor
```

Previously, the CIFS `status` command was used. You can modify the line in existing pool cluster resources to take advantage of the CIFS restart capability of the `monitor` command. For information, see Section 12.7, "Configuring a Monitor Script for the Shared NSS Pool," on page 191.

2.7   What's New (August 2011 Patches)

In addition to bug fixes, Novell Cluster Services added support for OES 2 SP3 services and file systems on the SUSE Linux Enterprise Server (SLES) 10 SP4 operating system. You can upgrade to SLES 10 SP4 by using the `move-to-sles10-sp4 patch` in the SLES patch channel.

2.8 What's New (May 2011 Patches)

In addition to bug fixes, the May 2011 Scheduled Maintenance patches for OES 2 SP2 and OES 2 SP3 provide the following change for Novell Cluster Services:

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

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

2.9 What's New (April 2011 Patches)

In addition to bug fixes, the April 2011 Scheduled Maintenance patches for OES 2 SP2 and OES 2 SP3 provide the following changes for Novell Cluster Services:

 Section 2.9.1, "Script Changes for New NSS Pool Cluster Resources," on page 38
 Section 2.9.2, "Run Report Opens in a New Window," on page 38

2.9.1 Script Changes for New NSS Pool Cluster Resources

For newly created clustered NSS pools and for newly cluster-enabled existing NSS pools, the default order in the load script has changed to add the secondary IP address before the NSS pool activation. For example:

```
exit_on_error add_secondary_ipaddress <ip_address>
exit_on_error nss /poolact=<pool_name>
```

The changed order is reflected in the sample unload script in Section 12.5, "Configuring a Load Script for the Shared NSS Pool," on page 189.

For newly created clustered NSS pools and for newly cluster-enabled existing NSS pools, the default order in the unload script has changed to deactivate the NSS pool before deleting the secondary IP address. For example:

```
ignore_error nss /pooldeact=<pool_name>
ignore_error del_secondary_ipaddress <ip_address>
```

The changed order is reflected in the sample unload script in Section 12.6, "Configuring an Unload Script for the Shared NSS Pool," on page 190.

2.9.2 Run Report Opens in a New Window

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

2.10 What's New (OES 2 SP3)

In addition to bug fixes and the changes from previously released patches, the following changes and enhancements were made for Novell Cluster Services for Linux in OES 2 SP3:

2.10.1 Support for the OES Common Proxy User

Support was added for the OES Common Proxy User feature in eDirectory 8.8.6. For information, see "OES Common Proxy User" on page 57.

2.10.2 NCS_Management Group

The `NCS_Management` group for a cluster contains the user name of the cluster administrator. This can be the OES Common Proxy user, an LDAP Admin user, or another administrator user that you specify for the cluster. Typically, the cluster administrator user is specified during the initial cluster configuration, as described in Step 3 of Section 5.5.5, "Configuring a New Cluster," on page 88. For information about configuring this user for an existing cluster, see Section 8.7, "Changing the NCS Proxy User Assignment in the NCS_Management_Group," on page 114.

2.10.3 Viewing and Configuring the LDAP Server Information

The following LDAP-related improvements were made for the handling of LDAP server information for a cluster:

- **Viewing the LDAP Server:** A new option was added to show which LDAP server would be used if Novell Cluster Services initiated an LDAP transaction now. It also shows information about how Novell Cluster Services settled on the LDAP server (that is, the servers tried and results).

- **Preferred List for LDAP Servers:** Novell Cluster Services switches to a different LDAP server if a dead server is encountered in the binding process. A list of LDAP servers that the cluster can use can be found in the `/etc/opt/novell/ncs/clstrlib.conf` file. Novell Cluster Services switches to the next LDAP server according to the order specified in the list.

2.10.4 Resource Mutual Exclusion Groups

The Resource Mutual Exclusion (RME) Groups feature allows you to configure sets of resources that cannot be assigned to the same node at the same time. For information, see Section 10.10, "Configuring Resource Mutual Exclusion Groups," on page 162.

2.10.5 Reboot Option for Failure Actions in Resource Monitoring

A new reboot option is available as a failure action in the Resource Monitoring setup. If a failure action initiates, this option reboots the operating system on the hosting node without synchronizing or unmounting the disks. Each of the resources on the hosting node fails over to the next server in its *Assigned Nodes* list because of the reboot. This is a hard reboot, not a graceful one. The reboot option is normally used only for a mission-critical cluster resource that must remain available. For information, see "Failure Action" on page 156.

2.10.6 Modifying the Preferred Nodes List as Text

A text editor environment is provided for ordering nodes listed in the Preferred Nodes page of the Clusters plug-in. Previously, you could modify the order of nodes in the list only by using the arrows to move one node up or down at a time. For information, see Step 7 in Section 10.8, "Configuring Preferred Nodes for a Resource," on page 160.

2.10.7 Modifying Resource Priorities List as Text

A text editor environment is provided for ordering resources in the Resource Priorities page of the Clusters plug-in. Previously, you could modify the order of resources in the list only by using the arrows to move one resource up or down at a time. For information, see Step 5 in Section 10.9, "Configuring Resource Priorities for Load Order," on page 161.

2.10.8 Cluster Report Information

In the Clusters plug-in for iManager, the *Clusters > Cluster Manager > Run Report* output has been modified to include the following information

- Monitoring script for each cluster resource where Monitoring is enabled
- An *RME Groups* section that lists the member resources of each group. For information, see "Cluster Report" on page 165.

2.10.9 Informative Subject Line for Email Notifications

The subject line of email messages sent by Novell Cluster Services now provides information about the cluster name, resource name, action taken, and node name. For example:

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

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

2.10.10 Confirming a Cluster Down or Cluster Restart Command

The `cluster down` command and `cluster restart` command now prompt for a confirmation (Yes/No) before bringing down the cluster.

2.10.11 Handling Reboot Preferences with the Linux Kernel Panic Setting

Novell Cluster Services has been modified to handle panic situations for a node according to the Linux kernel panic setting on the node. This setting determines whether the server is automatically rebooted after an automatic node shutdown. For information, see Section 9.11, "Preventing a Cluster Node Reboot after a Node Shutdown," on page 134.

2.10.12 Previewing All Cluster Resources Before the Final Conversion from NetWare 6.5 to OES 2 Linux

You can preview all resources before finalizing a cluster conversion by using the `cluster convert preview` command without specifying a resource name. Previously, you could preview only one resource at a time. For information, see "Finalizing the Cluster Conversion" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*.

2.10.13 Providing the Password for ncs_install.py

The `ncs_install.py` script prompts for a password when it is run from the command line. Previously, the password was specified in the script.

2.10.14 Renaming a Pool Cluster Resource

The new `cluster rename` command allows you to rename a pool cluster resource.

`cluster rename <resource_name> <new_resource_name>`

For information, see Section 10.13, "Renaming a Cluster Resource," on page 166.

2.10.15 Controlling Monitoring for Resources

The new `cluster monitor` command allows you to start, stop, or view the status of monitoring for a specified cluster resource. Monitoring must be enabled for the cluster resource in order to use this command.

`cluster monitor <resource_name> {start | stop | status}`

For information, see Section 10.11, "Controlling Resource Monitoring," on page 166.

2.10.16 Cascading Failover Prevention

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

2.11 What's New (June 2010 Patches)

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

◆ **Email:** The subject line of email messages sent by Novell Cluster Services now provides information about the cluster name, resource name, action taken, and node name. For example:

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

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

◆ **Cluster Report:** The *Clusters > Cluster Manager > Run Report* output has been modified to include the Monitoring script for each cluster resource.

◆ **Panic:** Novell Cluster Services has been modified to handle panic situations for a node according to the Linux kernel panic setting on the node. For information, see Section 9.11, "Preventing a Cluster Node Reboot after a Node Shutdown," on page 134.

2.12 What's New (January 2010 Patches)

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

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

2.13 What's New (OES 2 SP2)

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

◆ Section 2.13.1, "Improved Error Reporting," on page 42
◆ Section 2.13.2, "Improved Time Calculations," on page 43
◆ Section 2.13.3, "Specifying the Size of the SBD Partition," on page 43
◆ Section 2.13.4, "Customizing Translation Syntax," on page 43
◆ Section 2.13.5, "Migration Tool Support for Cluster Configurations," on page 43
◆ Section 2.13.6, "New iFolder Resource Template," on page 43
◆ Section 2.13.7, "Removed MPK Calls from the Code," on page 43
◆ Section 2.13.8, "Cluster Restart Is No Longer Required in a Rolling Cluster Upgrade," on page 44

2.13.1 Improved Error Reporting

This release provides improved error reporting for file protocol errors.

2.13.2 Improved Time Calculations

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

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

For example:

```
cluster stats display

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

2.13.3 Specifying the Size of the SBD Partition

When you configure the SBD partition during the cluster configuration of a new cluster (as described in "Configuring a New Cluster" on page 88), you can now specify the size of the partition.

2.13.4 Customizing Translation Syntax

Beginning in OES 2 SP2, Novell Cluster Services allows you to customize the translation syntax that is used for load and unload scripts in mixed-platform situations by defining new syntax translations in the /var/opt/novell/ncs/customized_translation_syntax file that you create. The clstrlib.py script reads the additional translation syntax from the syntax file, and processes them in addition to the normal translations in the Cluster Translation Library. For information, see "Customizing the Translation Syntax for Converting Load and Unload Scripts" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*.

2.13.5 Migration Tool Support for Cluster Configurations

Support was added for migrating services and data in cluster configurations by using the OES 2 SP2 Migration Tool. For instructions on using the Migration Tool to migrate services and data, see the *OES 2 SP3: Migration Tool Administration Guide*.

2.13.6 New iFolder Resource Template

The Novell iFolder 3.*x* resource template has been modified for OES 2 SP2. For information about using the new template, see Section 6.4.1, "Updating the iFolder Resource Template," on page 99.

2.13.7 Removed MPK Calls from the Code

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

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

In the OES 2 SP1 release, a cluster restart was required at the end of a rolling cluster upgrade in order to properly update the names of the nodes, as described in Section 2.14.3, "Behavior Change for Adding a Node," on page 44. This issue was resolved in OES 2 SP2. The rolling cluster upgrade process no long requires a cluster restart.

2.14 What's New (OES 2 SP1)

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

- Section 2.14.1, "Schema Extension," on page 44
- Section 2.14.2, "Installation by Container Administrator," on page 44
- Section 2.14.3, "Behavior Change for Adding a Node," on page 44
- Section 2.14.4, "Attribute NCS: GIPC Config Is No Longer Maintained," on page 45
- Section 2.14.5, "Support for Novell AFP for Linux," on page 45
- Section 2.14.6, "Support for Novell CIFS for Linux," on page 45
- Section 2.14.7, "Support for Domain Services for Windows," on page 45

2.14.1 Schema Extension

The administrator of a Novell eDirectory tree can now extend the schema for cluster objects before clusters are installed in the tree. This allows container administrators to install Novell Cluster Services without needing tree-level administrator rights. See Section 5.2, "Extending the eDirectory Schema to Add Cluster Objects," on page 78.

2.14.2 Installation by Container Administrator

Container administrators can install Novell Cluster Services without needing tree-level administrator rights. Ensure that you have the rights needed for the install. See Section 5.3, "Assigning Install Rights for Container Administrators (or Non-Administrator Users)," on page 80.

2.14.3 Behavior Change for Adding a Node

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

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

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

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

or

```
rcnovell-ncs restart
```

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

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

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

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

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

2.14.4 Attribute NCS: GIPC Config Is No Longer Maintained

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

2.14.5 Support for Novell AFP for Linux

This release supports Novell AFP (Apple Filing Protocol) for Linux in combination with Novell Storage Services (NSS) volumes on OES 2 SP1 Linux. See "Novell AFP" on page 69.

2.14.6 Support for Novell CIFS for Linux

This release supports Novell CIFS for Linux in combination with NSS volumes on OES 2 SP1 Linux. See "Novell CIFS" on page 69.

2.14.7 Support for Domain Services for Windows

This release supports using clusters in Domain Services for Windows contexts for OES 2 SP1 Linux. See "Novell Domain Services for Windows" on page 69.

2.15 What's New (OES 2)

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

❒ **Resource Monitoring:** See Section 10.6, "Enabling Monitoring and Configuring the Monitor Script," on page 155.

❒ **Support for Xen Virtualization:** See Chapter 14, "Configuring Novell Cluster Services in a Xen Virtualization Environment," on page 237.

3 Planning for a Cluster

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

IMPORTANT: For information about the system requirements for installing and using Novell Cluster Services, see Chapter 4, "Planning for Novell Cluster Services," on page 55.

- Section 3.1, "Determining Your Design Criteria," on page 47
- Section 3.2, "Using Cluster Best Practices," on page 47
- Section 3.3, "Planning the LAN Connectivity," on page 48
- Section 3.4, "Planning the Shared Storage Connectivity," on page 50
- Section 3.5, "Planning the Shared Storage Solution," on page 50
- Section 3.6, "Planning the eDirectory Deployment," on page 50
- Section 3.7, "Planning for Shared Storage as Cluster Resources," on page 51
- Section 3.8, "Planning for OES Services as Cluster Resources," on page 52

3.1 Determining Your Design Criteria

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

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

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

3.2 Using Cluster Best Practices

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

- Ensure that eDirectory is stable before implementing a cluster.

- Ensure that you have full Read/Write replicas of the entire eDirectory tree co-located in the data center where you are setting up the cluster.
- Ensure that IP addresses are unique.
- IP address assignments should be consistently applied for each cluster and its cluster resources.
- IP address changes for the cluster and cluster resources should only be made by using the procedure described in .

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

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

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

3.3 Planning the LAN Connectivity

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

3.3.1 VLAN

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

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

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

3.3.2 Channel Bonding

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

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

3.3.3 Spanning Tree Protocol

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

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

3.3.4 IP Addresses

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

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

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

3.3.5 Name Resolution

Ensure that SLP is properly configured for name resolution. For information, see Section 4.5.3, "SLP," on page 61.

3.4 Planning the Shared Storage Connectivity

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

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

- Use host-based multipath I/O management. For information, see the following resources:
    - Section 4.9, "Multipath I/O Configuration Requirements," on page 73
    - "Managing Multipath I/O for Devices" (http://www.suse.com/documentation/sles10/stor_admin/data/multipathing.html) in the *SLES 10 SP4 Storage Administration Guide* (http://www.suse.com/documentation/sles10/stor_admin/data/bookinfo.html)
- Connect each node via two fabrics to the storage area network (SAN).
- Use redundant SAN connections to provide fault-tolerant connectivity between the cluster nodes and the shared storage devices.
- Use LUN masking to exclusively assign each LUN to one or more host connections. For information, see Section 4.8, "SAN Rules for LUN Masking," on page 72.

3.5 Planning the Shared Storage Solution

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

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

    A LUN cannot be concurrently accessed by servers belonging to different clusters. This means that all resources on a given LUN can be active only in a given cluster at any given time.

- We recommend that you use only one LUN per pool, and only one volume per pool. If you use multiple LUNs for a given shared NSS pool, all LUNs must fail over together.

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

3.6 Planning the eDirectory Deployment

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

- Section 3.6.1, "Object Location," on page 50
- Section 3.6.2, "Cluster OU Context," on page 51
- Section 3.6.3, "Cluster OU Partitioning and Replication," on page 51

3.6.1 Object Location

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

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

3.6.2   Cluster OU Context

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

Figure 3-1   *Example: Cluster1 Container and Its Objects*

```
[ROOT]
   NOVELL
      LOC_1
         RESOURCES
            CLUSTER1
               CL1
               NODE101
               NODE102
```
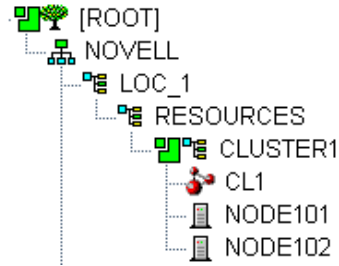
3.6.3   Cluster OU Partitioning and Replication

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

If you do not want to put a replica of eDirectory on the node, you must configure one or multiple LDAP servers for the node to use. The LDAP servers must have a master replica or a Read/Write replica of eDirectory. For information about how to modify the LDAP server list that is used by a cluster, see Section 8.9.1, "Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster," on page 116.

3.7   Planning for Shared Storage as Cluster Resources

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

File System or Storage Solution	For information, see
Novell Storage Services (NSS) pools	Chapter 12, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 173
Linux POSIX volumes	Chapter 13, "Configuring and Managing Cluster Resources for Shared Linux POSIX Volumes," on page 213
NCP (NetWare Control Protocol) volumes (NCP shares on cluster-enabled Linux POSIX volumes)	"Configuring NCP Volumes with Novell Cluster Services" in the *OES 2 SP3: NCP Server for Linux Administration Guide*
Dynamic Storage Technology (DST) volumes (NSS volumes configured in a shadow volume pair)	"Configuring DST Shadow Volume Pairs with Novell Cluster Services" in the *OES 2 SP3: Dynamic Storage Technology Administration Guide*

3.8 Planning for OES Services as Cluster Resources

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

Service	For information, see
Apache Web Server	"Apache HTTP Server" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*
Apple Filing Protocol (Novell AFP)	"Configuring AFP with Novell Cluster Services for an NSS File System" in the *OES 2 SP3: Novell AFP For Linux Administration Guide*
Archive and Version Services	"Configuring Archive and Version Service for Novell Cluster Services" in the *OES 2 SP3: Novell Archive and Version Services 2.1 Administration Guide for Linux*
Certificate Server (eDirectory Server Certificates)	"eDirectory Server Certificates" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*
CIFS (Novell CIFS)	"Configuring CIFS with Novell Cluster Services for an NSS File System" in the *OES 2 SP3: Novell CIFS for Linux Administration Guide*
DFS VLDB (Novell Distributed File Services Volume Location Database)	"Clustering Novell Distributed File Services" in the *OES 2 SP3: Novell Distributed File Services Administration Guide for Linux*
DHCP Server on a Linux POSIX volume	"Configuring DHCP with Novell Cluster Services for the Linux File System" in the *OES 2 SP3: Novell DNS/DHCP Administration Guide*
DHCP Server on an NSS volume	"Configuring DHCP with Novell Cluster Services for the NSS File System" in the *OES 2 SP3: Novell DNS/DHCP Administration Guide*
DNS Server	"Configuring DNS with Novell Cluster Services" in the *OES 2 SP3: Novell DNS/DHCP Administration Guide*
iFolder 3.*x*	"Clustering iFolder 3.8 Servers with Novell Cluster Services for Linux" in the *Novell iFolder 3.8 Administration Guide*
iPrint	"Configuring iPrint with Novell Cluster Services" in the *OES 2 SP3: iPrint for Linux Administration Guide*
MySQL	"High Availability and Scalability" in the *MySQL 5.x Reference Manual* (http://dev.mysql.com/doc/refman/5.5/en/ha-overview.html) "MySQL" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*
NetStorage	"Configuring NetStorage with Novell Cluster Services" in the *OES 2 SP3: NetStorage Administration Guide*.
PureFTP	"Cluster Enabling Pure-FTPd in an OES 2 Environment" in the *OES 2 SP3: Planning and Implementation Guide*

Service	For information, see
QuickFinder 5.0.x (Server Synchronization Feature)	"QuickFinder Server" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*
Samba (Novell Samba)	"Configuring Samba for Novell Cluster Services" in the *OES2 SP3: Samba Administration Guide*

4 Planning for Novell Cluster Services

This section describes the requirements for installing and using Novell Cluster Services on Novell Open Enterprise Server (OES) 2 SP3 for Linux servers.

IMPORTANT: For information about designing your cluster and cluster environment, see Chapter 3, "Planning for a Cluster," on page 47.

4.1 Cluster Administration Requirements

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

4.1.1 Cluster Installation Administrator

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

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

eDirectory Schema Administrator

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

For instructions, see Section 5.2, "Extending the eDirectory Schema to Add Cluster Objects," on page 78.

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

Container Administrator

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

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

For instructions, see Section 5.3, "Assigning Install Rights for Container Administrators (or Non-Administrator Users)," on page 80

4.1.2   NCS Proxy User

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

In OES 2 SP2 Linux and earlier, the NCS Proxy User was typically the LDAP administrator user name that was specified when setting up the server. You could alternately specify a user with sufficient rights to perform the installation and configuration tasks as specified in "Container Administrator" on page 56. You used this identity until the cluster and resources were configured, then set up a different user identity to use as the NCS Proxy User instead of continuing to use the LDAP administrator identity.

Beginning in OES 2 SP3, Novell Cluster Services supports the OES Common Proxy User enablement feature of Novell eDirectory 8.8.6. If the Common Proxy user is enabled in eDirectory when you configure the cluster, you can specify whether to use the Common Proxy user, the LDAP Admin user, or another administrator user. The specified user is automatically assigned to the NCS_Management group that resides in the Cluster object container. This accommodates the server-specific common user for each of the nodes. As a group member, the assigned user has the necessary rights for configuring the cluster and cluster resources and for exchanging information with eDirectory.

You can modify this default administrator user name or password for the user identity that is assigned as the NCS Proxy User after the install by following the procedure in Section 8.9, "Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster," on page 116.

Consider the following caveats for the three proxy user options:

- "OES Common Proxy User" on page 57
- "LDAP Admin User" on page 57
- "Another Administrator User" on page 57

OES Common Proxy User

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

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

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

For information about enabling or disabling the OES Common Proxy User, see the *OES 2 SP3: Installation Guide*. For caveats and troubleshooting information for the OES Common Proxy user, see the *OES 2 SP3: Planning and Implementation Guide*.

LDAP Admin User

If you specify the LDAP Admin user as the NCS Proxy User, you typically continue using this identity while you set up the cluster and cluster resources. After the cluster configuration is completed, you create another user identity to use for this purpose, and grant that user sufficient adminstrator rights as specified in "Cluster Administrator or Administrator-Equivalent User" on page 57.

Another Administrator User

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

4.1.3 Cluster Administrator or Administrator-Equivalent User

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

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

4.2 IP Address Requirements

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

❏ You need additional unique static IP addresses for the cluster and for each cluster resource and cluster-enabled pool.

❏ All IP addresses used by the master cluster IP address, its cluster servers, and its cluster resources must be on the same IP subnet. They do not need to be contiguous addresses.

4.3 Hardware Requirements

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

❏ A minimum of two Linux servers, and not more than 32 servers in a cluster

❏ At least 512 MB of additional memory on each server in the cluster

❏ One non-shared device on each server to be used for the operating system

❏ At least one network card per server in the same IP subnet

In addition, each server must meet the requirements for Novell Open Enterprise Server 2 Linux. For information, see "Meeting All Server Software and Hardware Requirements" in the *OES 2 SP3: Installation Guide*.

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

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

4.4 Xen Virtualization Environments

Xen virtualization software is included with SUSE Linux Enterprise Server. Novell Cluster Services supports using Xen virtual machine (VM) guest servers as nodes in a cluster. You can install Novell Cluster Services on the guest server just as you would a physical server. All templates except the Xen and XenLive templates can be used on a VM guest server. For examples, see Chapter 14, "Configuring Novell Cluster Services in a Xen Virtualization Environment," on page 237.

Novell Cluster Services is supported to run on a Xen host server where it can be used to cluster the virtual machine configuration files on Linux POSIX file systems. Only the Xen and XenLive templates are supported for use in the XEN host environment. For information about setting up Xen and XenLive cluster resources, see Section 14.2, "Virtual Machines as Cluster Resources," on page 238.

4.5 Software Requirements for Cluster Services

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

- Section 4.5.1, "Novell Open Enterprise Server 2 Linux," on page 59
- Section 4.5.2, "Novell eDirectory 8.8.6," on page 59
- Section 4.5.3, "SLP," on page 61
- Section 4.5.4, "Novell Cluster Services," on page 62
- Section 4.5.5, "Novell iManager 2.7.4," on page 63
- Section 4.5.6, "Clusters Plug-in for iManager," on page 63
- Section 4.5.7, "Storage-Related Plug-Ins for iManager," on page 64
- Section 4.5.8, "EVMS," on page 65
- Section 4.5.9, "OpenWBEM and CIMOM," on page 65
- Section 4.5.10, "CASA," on page 66
- Section 4.5.11, "Web Browser," on page 66

4.5.1 Novell Open Enterprise Server 2 Linux

For new installs and migrations, Novell Cluster Services 1.8.8 for Linux supports Novell Open Enterprise Server 2 SP3 services and file systems running on SUSE Linux Enterprise Server 10 SP4. Novell Cluster Services is one of the OES 2 Services patterns for OES 2 Linux.

We recommend having uniform nodes in the cluster. The same release of OES 2 Linux must be installed and running on each node in the cluster. Mixed-node clusters running 32-bit and 64-bit versions of the same OES 2 Linux release are supported. Wherever possible when working with mixed platforms, you should set up the preferred nodes for each resource so that the resource fails over only between 32-bit nodes or only between 64-bit nodes.

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

Upgrading from	For information, see:
OES 2 Linux	Chapter 6, "Upgrading OES 2 Linux Clusters," on page 97
OES 1 Linux	Chapter 7, "Upgrading OES 1 Linux Clusters to OES 2 Linux," on page 101
NetWare 6.5 SP8	*OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*

4.5.2 Novell eDirectory 8.8.6

Novell eDirectory 8.8.6 (or later) is required for managing the Cluster object and Cluster Node objects for Novell Cluster Services in OES 2 SP3. eDirectory must be installed and running in the same tree where you create the cluster. eDirectory can be installed on any node in the cluster, on a separate server, or in a separate cluster. You can install an eDirectory master replica or replica in the cluster, but it is not required to do so for Novell Cluster Services.

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

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

- "eDirectory Tree" on page 60
- "eDirectory Context" on page 60
- "Cluster Object Container" on page 60
- "Cluster Objects Stored in eDirectory" on page 61
- "LDAP Server List" on page 61

eDirectory Tree

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

eDirectory Context

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

Cluster Object Container

We recommend that the Cluster object and all of its member Server objects and Storage objects be located in the same OU. Multiple Cluster objects can co-exist in the same eDirectory container.

Figure 4-1   *Same Container for Cluster Object and Server Objects*
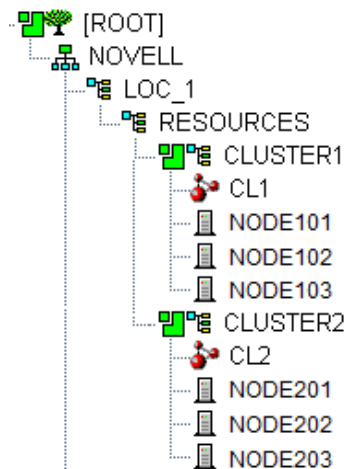
If the servers in the cluster are in separate eDirectory containers, the user that administers the cluster must have rights to the cluster server containers and to the containers where any cluster-enabled pool objects are stored. You can do this by adding trustee assignments for the cluster administrator to a parent container of the containers where the cluster server objects reside. See "eDirectory Rights" (http://www.novell.com/documentation/edir88/edir88/data/fbachifb.html) in the *eDirectory 8.8 Administration Guide* for more information.

Cluster Objects Stored in eDirectory

After you create a cluster, the following cluster objects are automatically created and stored in eDirectory under the Cluster object ( ):

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

The following objects are added to eDirectory when you add nodes or create cluster resources:

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

LDAP Server List

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

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

For information about configuring a list of LDAP servers for the cluster, see .

4.5.3 SLP

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

SLP is not automatically installed when you select Novell Cluster Services. SLP is installed as part of the Novell eDirectory configuration during the OES 2 Linux install. You can enable and configure SLP on the eDirectory Configuration - NTP & SLP page. For information, see "Specifying SLP Configuration Options" in the *OES 2 SP3: Installation Guide*.

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

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

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

You can use the `ndsconfig set` command to set the `n44.nds.advertise-life-time` parameter. To reset the parameter in a cluster, perform the following tasks on each node of the cluster:

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

2 Take offline all of the cluster resources on the node, or cluster migrate them to a different server. At a command prompt, enter

```
cluster offline <resource_name>

or

cluster migrate <resource_name> <target_node_name>
```

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

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

rcndsd restart

rcslpd restart
```

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

```
cluster online <resource_name>

or

cluster migrate <resource_name> <node_name>
```

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

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

For more information about configuring and managing SLP, see "Configuring OpenSLP for eDirectory" in the *Novell eDirectory 8.8 SP7 Administration Guide*.

4.5.4 Novell Cluster Services

Novell Cluster Services is required for creating and managing clusters and shared resources. It must be installed on the server to allow devices to be set to a shared state, such as the device used for the SBD partition and for shared storage resources.

4.5.5   Novell iManager 2.7.4

Novell iManager 2.7.4 (or later) is required for configuring and managing clusters on OES 2 SP3 Linux. iManager must be installed on at least one computer in the same tree as the cluster. It can be installed in the cluster or not in the cluster. For information about using iManager, see the iManager 2.7*x* documentation Web site (http://www.novell.com/documentation/imanager27/index.html).

For OpenWBEM and CIMOM requirements, see Section 4.5.9, "OpenWBEM and CIMOM," on page 65.

For browser configuration requirements, see "Web Browser" on page 66.

4.5.6   Clusters Plug-in for iManager

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

The following components must be installed in iManager:

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

    For information, see "Storage-Related Plug-Ins for iManager" on page 64.

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

The Clusters plug-in also provides an integrated management interface for Novell Business Continuity Clustering (BCC). The additional interface is present only if BCC is installed on the server. See the following table for information about the versions of BCC that are supported. BCC is sold separately from OES 2 Linux. For purchasing information, see the BCC product page (http://www.novell.com/products/businesscontinuity/).

BCC Release	OES Support	iManager and Clusters Plug-In
BCC 2.0	OES 11 SP1	Novell iManager 2.7.6 or later
		Requires the Clusters plug-in for OES 11 SP1 with the latest patches applied.
		See the *BCC 2.0 Administration Guide for OES 11 SP1* (http://www.novell.com/documentation/bcc/bcc20_admin_lx/data/bookinfo.html).
BCC 1.2.2	OES 2 SP3	Novell iManager 2.7.4 or later
		Requires the Clusters plug-in for OES 2 SP3 and the OES 2 SP3 April 2011 Scheduled Maintenance patch.
		See the *BCC 1.2.2: Administration Guide for OES 2 SP3* (http://www.novell.com/documentation/bcc/bcc122_admin_lx/data/bookinfo.html).

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

4.5.7 Storage-Related Plug-Ins for iManager

In OES 2 Linux, the following storage-related plug-ins for iManager share code in common in the `storagemgmt.rpm` file:

Product	Plug-In	NPM File
Novell Apple Filing Protocol (AFP) (for OES 2 SP1 Linux and later)	File Protocols > AFP	`afpmgmt.rpm`
Novell Archive and Version Services	Archive Versioning	`arkmgmt.rpm`
Novell CIFS (for OES 2 SP1 Linux and later)	File Protocols > CIFS	`cifsmgmt.rpm`
Novell Cluster Services	Clusters	`ncsmgmt.rpm`
Novell Distributed File Services	Distributed File Services	`dfsmgmt.rpm`
Novell Storage Services	Storage	`nssmgmt.rpm`

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

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

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

The plug-in files are included on the installation disk. The latest Novell storage-related plug-ins can be downloaded as a single zipped download file from the Novell Downloads Web site (http://download.novell.com). For information about installing plug-ins in iManager, see "Downloading and Installing Plug-in Modules" in the *Novell iManager 2.7.6 Administration Guide*.

For information about working with storage-related plug-ins for iManager, see "Understanding Storage-Related Plug-Ins" in the *OES 2 SP3: NSS File System Administration Guide for Linux*.

4.5.8　EVMS

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

Updates to EVMS are received through the update channel for SUSE Linux Enterprise Server 10 SP4 or later. Ensure that you install the latest patches for EVMS before you create any cluster resources for this server.

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

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

4.5.9　OpenWBEM and CIMOM

OpenWBEM is a PAM-enabled Linux utility that must be enabled and running on the OES 2 Linux server when managing services with Novell Remote Manager for Linux and Novell iManager. OpenWBEM must be configured to start with `chkconfig`, and be running when you manage the cluster with Novell iManager. For information on setup and configuration, see the *OES 2 SP3: OpenWBEM Services Administration Guide*.

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

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

For OES 2 and later, the Clusters plug-in (and all other storage-related plug-ins) for iManager require CIMOM connections for tasks that transmit sensitive information (such as a user name and password) between iManager and the `_admin` volume on the OES 2 Linux that server you are managing. Typically, CIMOM is running, so this should be the normal condition when using the server. CIMOM connections use Secure HTTP (HTTPS) for transferring data, and this ensures that sensitive data is not exposed.

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

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

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

To perform this task	At a terminal console prompt, enter as the `root` user
To check openWBEM status	`rcowcimomd status`
To restart openWBEM	`rcowcimomd restart`

4.5.10 CASA

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

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

To perform this task	At a terminal console prompt, enter as the `root` user
To check CASA status	`rcmicasad status`
To check that CASA is running correctly	`CASAcli -l`
To restart CASA	`rcmicasad restart`

4.5.11 Web Browser

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

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

Supported language codes are Unicode (UTF-8) compliant. To avoid display problems, ensure that the Character Encoding setting for the browser is set to Unicode (UTF-8) or ISO 8859-1 (Western, Western European, West European).

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

For a information about supported browsers, see "Using a Supported Web Browser" in the *Novell iManager 2.7.6 Administration Guide*.

4.6 Software Requirements for Cluster Resources

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

4.6.1 NCP Server for Linux

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

NCP Server for Linux is required to be installed and running before you can cluster-enable the following storage resources:

- Novell Storage Services (NSS) pools and volumes
- NCP volumes on Linux POSIX file systems
- Dynamic Storage Technology shadow volume composed of a pair of NSS volumes
- Shared Linux POSIX file system if you want to give it an `NCS:NCP Server` object that can be seen in the tree view.

This is necessary in order to provide authenticated access using the Novell Trustee model. NCP Server must be installed and running in order for the *NCP* option to be available as an advertising protocol for the NSS pool cluster resource.

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

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

For information, see "Configuring Cross-Protocol File Locks for NCP Server" in the *OES 2 SP3: NCP Server for Linux Administration Guide*.

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

For information about configuring and managing NCP Server for Linux, see the *OES 2 SP3: NCP Server for Linux Administration Guide*.

For information about creating and cluster-enabling NCP volumes on Linux POSIX file systems, see "Configuring NCP Volumes with Novell Cluster Services" in the *OES 2 SP3: NCP Server for Linux Administration Guide*.

4.6.2   Novell Storage Services File System for Linux

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

◆ Intializing and sharing devices used for the SBD (split-brain detector) and for shared pools. For information, see Section 4.7.2, "SBD Partitions," on page 70.

◆ Creating and cluster-enabling a shared pool. For information, see Chapter 12, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 173.

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

◆ NetWare Core Protocol (NCP), which is selected by default and is mandatory for NSS. For information, see "NCP Server for Linux" on page 67.

◆ Novell Apple Filing Protocol (AFP). For information, see "Novell AFP" on page 69.

◆ Novell CIFS. For information, see "Novell CIFS" on page 69.

4.6.3   Linux POSIX File Systems

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

After the cluster is configured, you can create Linux POSIX cluster resources as described in Chapter 13, "Configuring and Managing Cluster Resources for Shared Linux POSIX Volumes," on page 213. Linux POSIX file systems must be created in EVMS in order to first lay down the Cluster Segment Manager on the shared device (disk or LUN).

NCP Server is required if you want to create a virtual server name (NCS:NCP Server object) for the Linux POSIX volume cluster resource. However, this object does not give the users NCP access to data unless an NCP volume is also added. For information, see Section 4.6.1, "NCP Server for Linux," on page 67.

4.6.4   NCP Volumes on Linux POSIX File Systems

Novell Cluster Services supports clustering for NCP volumes on Linux POSIX file systems. NCP Server is required. For information, see Section 4.6.1, "NCP Server for Linux," on page 67.

For information about creating and cluster-enabling NCP volumes on Linux POSIX file systems, see "Configuring NCP Volumes with Novell Cluster Services" in the *OES 2 SP3: NCP Server for Linux Administration Guide*.

4.6.5   Dynamic Storage Technology Shadow Volume Pairs

Novell Cluster Services supports clustering for Novell Dynamic Storage Technology (DST) shadow volume pairs on OES 2 Linux. DST is installed automatically when you install NCP Server for Linux. To use cluster-enabled DST volume pairs, select the *NCP Server and Dynamic Storage Technology* option during the install.

For information about creating and cluster-enabling Dynamic Storage Technology volumes on Linux, see "Configuring DST Shadow Volume Pairs with Novell Cluster Services" in the *OES 2 SP3: Dynamic Storage Technology Administration Guide*.

4.6.6 NCP File Access

Novell Cluster Services requires NCP file access to be enabled for cluster-enabled NSS volumes, NCP volumes, and DST volumes, even if users do not access files via NCP. This is required to support access control via the Novell Trustee model. For information, see Section 4.6.1, "NCP Server for Linux," on page 67.

4.6.7 Novell AFP

Novell Cluster Services supports using Novell AFP as an advertising protocol for cluster-enabled NSS pools and volumes on OES 2 SP1 Linux and later.

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

To install Novell AFP, select the *Novell AFP* option during the install. For information about configuring and managing the Novell AFP service, see the *OES 2 SP3: Novell AFP For Linux Administration Guide*.

4.6.8 Novell CIFS

Novell Cluster Services supports using Novell CIFS as an advertising protocol for cluster-enabled NSS pools and volumes on OES 2 SP1 Linux and later.

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

To install Novell CIFS, select the *Novell CIFS* option during the install. For information about configuring and managing the Novell CIFS service, see the *OES 2 SP3: Novell CIFS for Linux Administration Guide*.

4.6.9 Novell Samba

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

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

For information about configuring and managing Novell Samba, see the *OES2 SP3: Samba Administration Guide*.

4.6.10 Novell Domain Services for Windows

Novell Cluster Services supports using clusters in Domain Services for Windows (DSfW) contexts for OES 2 SP1 Linux and later. If Domain Services for Windows is installed in the eDirectory tree, the nodes in a given cluster can be in the same or different DSfW subdomains. Port 1636 is used for DSfW communications. This port must be opened in the firewall.

For information using Domain Services for Windows, see the *OES 2 SP3: Domain Services for Windows Administration Guide*.

4.7 Shared Disk Configuration Requirements

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

Ensure that your shared storage devices meet the following requirements:

4.7.1 Shared Devices

Novell Cluster Services supports the following shared disks:

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

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

Prior to installation, verify that all the drives in your shared disk system are recognized by Linux by viewing a list of the devices on each server that you intend to add to your cluster. If any of the drives in the shared disk system do not show up in the list, consult the OES 2 documentation or the shared disk system documentation for troubleshooting information.

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

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

4.7.2 SBD Partitions

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

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

For information about how SBD partitions work and how to create an SBD partition for an existing cluster, see Section 9.14, "Creating or Deleting Cluster SBD Partitions," on page 136.

- ◆ "Preparing SAN Devices for the SBD" on page 71
- ◆ "Initializing and Sharing a Device for the SBD" on page 71
- ◆ "Determining the SBD Partition Size" on page 71
- ◆ "Creating the SBD After the Initial Cluster Setup" on page 72

Preparing SAN Devices for the SBD

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

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

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

NOTE: For an iSCSI SAN, a very small iSCSI device might rarely have uncommon CHS (cylinder-head-sector) geometry values. NSS prefers at least 32 sectors per track. If there are fewer than 32 sectors per track, NSS tools using EVMS can fail to create the partition or to mark the device as shareable. You can use fdisk to check for valid CHS values before you initialize and share the device. If necessary, you can use fdisk to set the sectors for the device (such as /dev/sdx) to 32:

```
fdisk -H 64 -S 32 /dev/sdx
```

Initializing and Sharing a Device for the SBD

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

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

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

Determining the SBD Partition Size

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

Because an SBD partition ends on a cylinder boundary, the partition size might be slightly smaller than the size you specify.

Creating the SBD After the Initial Cluster Setup

You can also create the SBD by using the SBD Utility (`sbdutil`) before you configure a second node for the cluster and before you create file system cluster resources. The utility also supports mirroring the SBD partition. For information, see Section 9.14.3, "Creating a Non-Mirrored Cluster SBD Partition," on page 138.

4.7.3 Shared iSCSI Devices

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

4.7.4 Shared RAID Devices

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

Consider the following when using software RAIDs:

- NSS software RAID is supported for shared disks.
- Linux software RAID can be used in shared disk configurations that do not require the RAID to be concurrently active on multiple nodes. Linux software RAID cannot be used underneath clustered file systems (such as OCFS2, GFS, and CXFS) because Novell Cluster Services does not support concurrent activation.

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

4.8 SAN Rules for LUN Masking

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

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

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

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

4.9 Multipath I/O Configuration Requirements

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

4.9.1 Path Failover Settings for Device Mapper Multipath

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

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

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

Change the Retry setting in the `/etc/modprobe.conf.local` and `/etc/multipath.conf` files so that Novell Cluster Services works correctly with DM-MP. For information, see Section 4.9.2, "Modifying the Retry Setting for the modprobe.conf.local File," on page 73 and Section 4.9.3, "Modifying the Polling Interval, No Path Retry, and Failback Settings in the multipath.conf File," on page 74.

Also consider changes as needed for the retry settings in the HBA BIOS. For information, see Section 4.9.4, "Modifying the Port Down Retry and Link Down Retry Settings for an HBA BIOS," on page 75.

4.9.2 Modifying the Retry Setting for the modprobe.conf.local File

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

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

```
options qla2xxx qlport_down_retry=1
```

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

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

```
options qla2xxx qlport_down_retry=2
```

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

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

Polling Interval

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

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

No Path Retry

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

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

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

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

Failback

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

```
failback "manual"
```

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

Example of Multipath I/O Settings

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

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

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

For information about configuring the multipath.conf file, see "Managing Multipath I/O for Devices" (http://www.suse.com/documentation/sles10/stor_admin/data/multipathing.html) in the *SLES 10 SP4 Storage Administration Guide* (http://www.suse.com/documentation/sles10/stor_admin/data/bookinfo.html).

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

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

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

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

5 Installing and Configuring Novell Cluster Services on OES 2 Linux

This section describes how to install the Novell Cluster Services software on Novell Open Enterprise Server 2 Linux servers, how to configure the cluster on the first node, and how to configure other nodes for the cluster.

IMPORTANT: Before you install or configure Novell Cluster Services, ensure that you understand the requirements for it and have configured the environment as described in Chapter 4, "Planning for Novell Cluster Services," on page 55.

See the following resources for information about rolling cluster upgrades or conversions to the latest version of Novell Cluster Services for OES 2 Linux:

Upgrading from	For information, see:
OES 2 Linux	Chapter 6, "Upgrading OES 2 Linux Clusters," on page 97
OES 1 Linux	Chapter 7, "Upgrading OES 1 Linux Clusters to OES 2 Linux," on page 101
NetWare 6.5 SP8	*OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*

- Section 5.1, "Novell Cluster Services Licensing," on page 77
- Section 5.2, "Extending the eDirectory Schema to Add Cluster Objects," on page 78
- Section 5.3, "Assigning Install Rights for Container Administrators (or Non-Administrator Users)," on page 80
- Section 5.4, "Installing Novell Cluster Services," on page 81
- Section 5.5, "Configuring Novell Cluster Services," on page 84
- Section 5.6, "Configuring Additional Administrators," on page 92
- Section 5.7, "Installing or Updating the Clusters Plug-in for iManager," on page 92
- Section 5.8, "Patching Novell Cluster Services," on page 94
- Section 5.9, "What's Next," on page 95

5.1 Novell Cluster Services Licensing

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

5.2   Extending the eDirectory Schema to Add Cluster Objects

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

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

In OES 2 SP1 Linux and later, a tree administrator user with the eDirectory credentials to do so can extend the eDirectory schema before a cluster is installed anywhere in a tree. This allows container administrators (or non-administrator users) to install a cluster in a container in that same tree without needing full administrator rights for the tree. After the schema has been extended, you must assign some eDirectory rights to the container administrators (or non-administrator users) who will install Novell Cluster Services clusters.

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

IMPORTANT: For information about the eDirectory rights needed to install Novell Cluster Services in a tree after the schema has been extended, see Section 5.3, "Assigning Install Rights for Container Administrators (or Non-Administrator Users)," on page 80.

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

- Section 5.2.1, "Prerequisites for Extending the Schema," on page 78
- Section 5.2.2, "Extending the Schema," on page 79

5.2.1   Prerequisites for Extending the Schema

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

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

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

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

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

Parameter	Description	Example
admin_password	The password of the administrator user.	`pas5W0rd`
server_ip_address	The IP address of the eDirectory server that contains the schema files.	`10.10.10.102`

5.2.2 Extending the Schema

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

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

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

1 On an OES 2 SP1 Linux (or later) server, open a terminal console, then log in as the `root` user to the tree.

2 In a text editor, create a text file, specify the configuration information for the Novell Cluster Services cluster in it, then save the file.

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

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

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

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

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

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

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

Replace `configuration_filename` with the actual name of the file that you created in Step 2.

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

This file contains a password in clear text. Ensure that you delete the file for security reasons.

6 Continue with Section 5.3, "Assigning Install Rights for Container Administrators (or Non-Administrator Users)," on page 80.

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

If the eDirectory schema has been extended in the tree where you want to create clusters, the container administrator (or non-administrator user) needs the following eDirectory rights to install and configure Novell Cluster Services. These rights are also required if a different user configures Novell Cluster Services after the install as described in Section 5.5.3, "Using Different LDAP Credentials for the Cluster Configuration," on page 87.

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

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

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

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

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

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

◆ Object Create rights where the cluster container will be.

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

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

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

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

5.4 Installing Novell Cluster Services

Novell Cluster Services for Linux is included on the add-on media for OES 2 Linux (or later). It is necessary to install OES 2 Linux on every server that you want to add to a cluster. You can install up to 32 nodes in each cluster. For information, see Section 5.1, "Novell Cluster Services Licensing," on page 77.

Installing Novell Cluster Services does the following:

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

IMPORTANT: Ensure that your system meets the requirements and guidelines in Chapter 4, "Planning for Novell Cluster Services," on page 55.

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

If the eDirectory schema has not been extended in the tree as described in Section 5.2, "Extending the eDirectory Schema to Add Cluster Objects," on page 78, the administrator user that installs Novell Cluster Services must have the rights to extend the schema, such as the tree administrator.

You can install Novell Cluster Services when you install OES 2 Linux, or on an existing OES 2 Linux server. The software must already be installed in order to be able to mark devices as shareable, such as the device (or devices) you want to use for the SBD (split-brain detector) when you configure the cluster or run the sbdutil utility.

- Section 5.4.1, "Installing Novell Cluster Services during a OES 2 Linux Installation," on page 81
- Section 5.4.2, "Installing Novell Cluster Services on an Existing OES 2 Linux Server," on page 83

5.4.1 Installing Novell Cluster Services during a OES 2 Linux Installation

This section describes only those steps in the install that are directly related to installing Novell Cluster Services. For detailed instructions on installing OES 2 Linux, see the *OES 2 SP3: Installation Guide*.

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

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

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

A one-node cluster that has shared disk storage can be configured without an SBD, but the SBD must be created before you add another node.

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

Typically, you create the SBD when you configure the cluster on the first node. For information about manually creating an SBD, see:

- Section 9.14.3, "Creating a Non-Mirrored Cluster SBD Partition," on page 138
- Section 9.14.4, "Creating a Mirrored Cluster SBD Partition," on page 140

2 Start the YaST install for SUSE Linux Enterprise Server 10 and continue to the Installation Mode page.

3 Select *New Installation*, select *Include Add-On Products from Separate Media*, click *Next*, then continue through the OES 2 add-on part of the install until you get to the Installation Settings page.

4 On the Installation Settings page, click *Software* to open the Software Selection and System Tasks page.

5 Under *OES Services*, select *Novell Cluster Services* and any other OES components that you want to install, then click *Accept*.
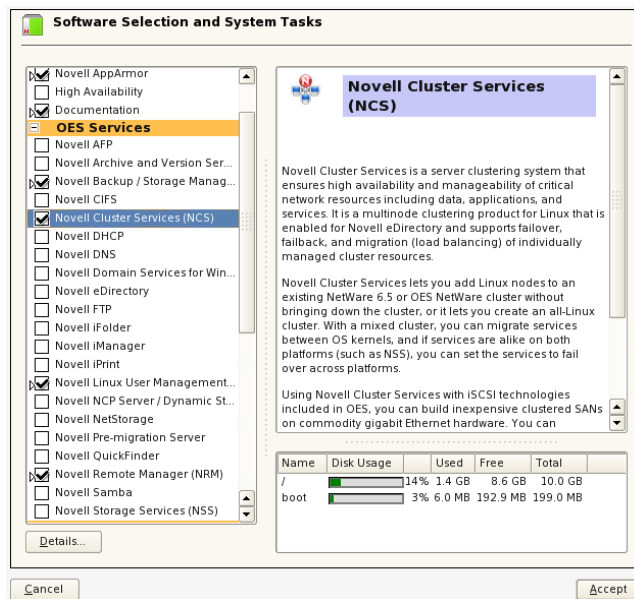
When you select Novell Cluster Services, the following basic services for managing OES 2 are automatically selected:

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

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

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

Select other protocols and services according to your planned setup. For information, see Section 4.5, "Software Requirements for Cluster Services," on page 59.

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

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

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

6 Continue through the installation process, but do not configure Novell Cluster Services when you reach the Novell Open Enterprise Server Configuration page. You configure it later.

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

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

9 Continue with Section 5.5, "Configuring Novell Cluster Services," on page 84.

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

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

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

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

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

A one-node cluster that has shared disk storage can be configured without an SBD, but the SBD must be created before you add another node.

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

Typically, you create the SBD when you configure the cluster on the first node. For information about creating the SBD manually, see:

   ◆ Section 9.14.3, "Creating a Non-Mirrored Cluster SBD Partition," on page 138
   ◆ Section 9.14.4, "Creating a Mirrored Cluster SBD Partition," on page 140

2 Log in to the server as the `root` user.

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

4 On the Software Selection page under *OES Services*, click *Novell Cluster Services* and any other compatible OES components that you want to install.
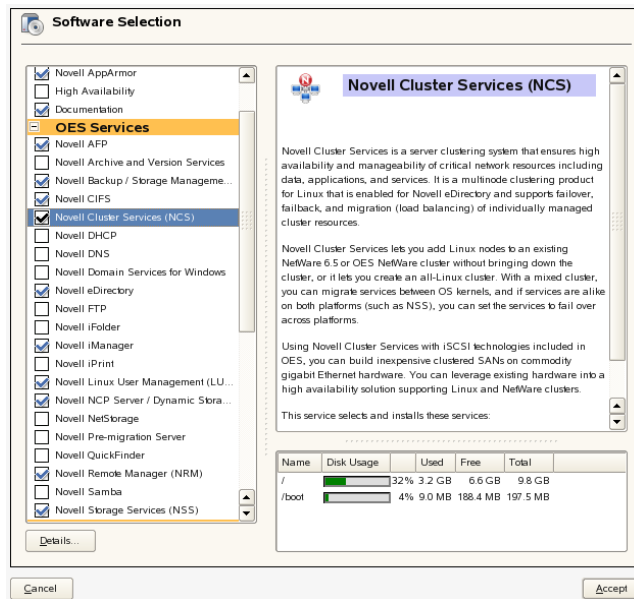
Services that you have already installed are indicated by a blue check mark in the status check box next to the service.

For information about the options, see Step 5 in Section 5.4.1, "Installing Novell Cluster Services during a OES 2 Linux Installation," on page 81.

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

6 Continue through the installation process, but do not configure Novell Cluster Services when you reach the Novell Open Enterprise Server Configuration page. You configure it later.

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

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

9 Continue with Section 5.5, "Configuring Novell Cluster Services," on page 84.

5.5 Configuring Novell Cluster Services

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

IMPORTANT: The YaST-based configuration is not used to modify the settings for an existing cluster. For information about modifying the settings for an existing cluster, see Section 8.9, "Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster," on page 116.

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

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

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

- ◆ Creates a new Cluster Node object in eDirectory.

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

5.5.1 Sharing a Device to Use for the SBD Partition

If you plan to use shared storage in the cluster, the cluster needs a split-brain detector (SBD). You must create the SBD before you add a second node to the cluster.

You can create the SBD when you set up the cluster on the first node. In preparation, you must initialize the partition, and mark its device as shareable for clustering. If you plan to mirror the SBD, you must also initialize the partition that you want to use as the mirrored SBD, and mark its device as shareable for clustering. You need at least 20 MB for the SBD.

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

2 At the console prompt, enter

```
nssmu
```

3 In the NSSMU, select *Devices* and press Enter.

4 In the *Devices* list, select the device that you want to use for the SBD.

5 If the device has not been initialized or if you want to delete the current partitioning structures on the device, press F3 to initialize the selected device.

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

Wait for the page to refresh before continuing.

6 Press F6 to mark the device as shareable for clustering.

The *Shareable for Clustering* value changes from *No* to *Yes*.

7 If you plan to mirror the SBD, repeat Step 4 through Step 6 for the second device.

8 Exit NSSMU.

9 Continue with Section 5.5.2, "Opening the Novell Open Enterprise Server Configuration Page," on page 86.

5.5.2 Opening the Novell Open Enterprise Server Configuration Page

1 Log in to the server as the `root` user.

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

3 On the Software Selection page under *OES Services*, verify that the *Novell Cluster Services* option is already installed as indicated by a blue check mark.
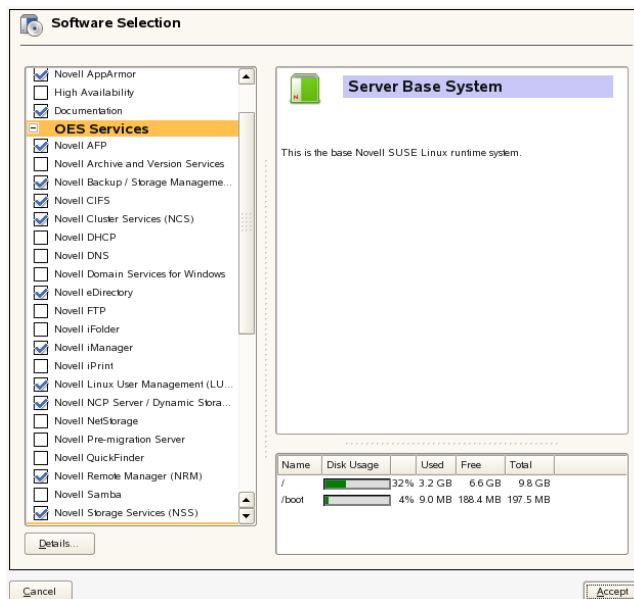
4 Click *Accept* to proceed to the Novell Open Enterprise Server Configuration page.

5 Do one of the following:

- ◆ **Same Administrator:** To use the same administrator credentials that were used to install Novell Cluster Services, continue with Section 5.5.4, "Accessing the Novell Cluster Services Configuration Page in YaST," on page 87.
- ◆ **Different Administrator:** To use different administrator credentials than those used to install Novell Cluster Services, continue with Section 5.5.3, "Using Different LDAP Credentials for the Cluster Configuration," on page 87.

5.5.3 Using Different LDAP Credentials for the Cluster Configuration

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

For information about what rights are needed, see Section 5.3, "Assigning Install Rights for Container Administrators (or Non-Administrator Users)," on page 80.

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

The sentence changes to *Reconfiguration is enabled*.

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

3 Specify the following values:

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

- ◆ **Admin password:** The password of the container administrator (or non-administrator user).

4 Click *Next*.

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

5 Continue with Section 5.5.4, "Accessing the Novell Cluster Services Configuration Page in YaST," on page 87.

5.5.4 Accessing the Novell Cluster Services Configuration Page in YaST

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

The configuration is currently disabled.

Novell Cluster Services (NCS)

Configure is disabled

2 Click the *disabled* link to enable configuration.

The sentence changes to *Configure is enabled*.

Novell Cluster Services (NCS)

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

Configure is enabled

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

4 If you are prompted for credentials, specify the password of the specified Administrator user, then click OK.

If you did not specify a different administrator user in Section 5.5.3, "Using Different LDAP Credentials for the Cluster Configuration," on page 87, this is the Administrator user whose credentials you specified for eDirectory when the OES Services were installed on the server.

5 On the Novell Cluster Services Configuration page, continue with one of the following:

- ◆
- ◆

5.5.5 Configuring a New Cluster

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

1 Go to the Novell Cluster Services Configuration page as described in .

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

New or Existing Cluster

Select *New Cluster* to define a new cluster.

Directory Server Address

Select the check box next to one or multiple IP addresses of the LDAP server to use for LDAP communications.

The IP addresses shown are the LDAP servers available for this service to use. The LDAP servers must have a master replica or a Read/Write replica of eDirectory.

Cluster FDN

Specify the FDN (fully distinguished name) you will give the new cluster and the eDirectory context where the new Cluster object will reside. The name is case sensitive.

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

For example:

```
cn=cluster1,ou=clusters,o=mycompany
```

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

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

Cluster IP address

The cluster IP address represents the cluster on the network. It is different than the server IP address. For example:

```
10.10.10.44
```

The cluster IP address provides a single point for cluster access, configuration, and management. A Master IP Address resource is created automatically during the Cluster Services installation. The cluster IP address is bound to the master node and remains with the master node regardless of which server is the master node.

The cluster IP address is required to be on the same IP subnet as the other servers in the same cluster, which is necessary for certain external network management programs to get cluster status alerts.

SBD Settings

A split-brain detector (SBD) is required if you plan to use shared disks in the cluster. For information about SBD requirements, see "SBD Partitions" on page 70.

If you need the SBD, you can create the SBD partition now, or you can create it manually at any time before you add a second node to the cluster.

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

To create the SBD partition now:

1. **Select the storage device with shared media:** Select the shared device that you want to use for the SBD from the drop-down list.

   The drop-down menu shows only devices that have been initialized and shared.

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

3. **Desired partition size of shared media:** Specify a size of 20 MB or greater to use for the SBD partition.

   If you selected a device for the SBD mirror, the specified size is also used for the mirror segment.

To manually create the SBD partition later, you can use either of the following procedures:

- Section 9.14.3, "Creating a Non-Mirrored Cluster SBD Partition," on page 138
- Section 9.14.4, "Creating a Mirrored Cluster SBD Partition," on page 140

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

- **OES Common Proxy User:** If the OES Common Proxy User is enabled in eDirectory, the *Use OES Common Proxy User* check box is automatically selected and the *NCS Proxy User Name* and *Specify NCS Proxy User Password* fields are populated with the credentials of the OES Common Proxy User.

- **LDAP Admin User:** If the OES Common Proxy User is disabled in eDirectory, the *Use OES Common Proxy User* check box is automatically deselected and the *NCS Proxy User Name* and *Specify NCS Proxy User Password* fields are populated with the credentials of the LDAP Admin user. The fields are also automatically populated with the LDAP Admin credentials if you deselect the *Use OES Common Proxy User* check box.

- **Another Administrator User:** Deselect the *Use OES Common Proxy User* check box, then specify the credentials of an existing administrator user.

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

4 On the Configuration page from the *IP address of this node* drop-down box, select the IP address that Novell Cluster Services will use for this node.

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

5 On the Configuration page, select the *Start Cluster Services now* check box to start Novell Cluster Services software on this node after configuring it, then click *Finish*.

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

7 Continue through the setup process.

8 After the configuration is completed, click *Finish* to exit the OES Configuration, then start Novell Cluster Services using one of these methods:

Setup Condition	Instructions
Start Cluster Services now was enabled in Step 5	Novell Cluster Services starts automatically after the configuration completes.
Start Cluster Services now was disabled in Step 5	Start Novell Cluster Services manually by using one of these methods:
	◆ Reboot the cluster server.
	◆ At a terminal console prompt, go to the `/etc/init.d` directory, then enter the following as the `root` user: `./novell-ncs start`
	◆ At a terminal console prompt, enter the following as the `root` user: `rcnovell-ncs start`

9 Continue with Section 5.6, "Configuring Additional Administrators," on page 92.

5.5.6 Adding a Node to an Existing Cluster

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

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

- ◆ Section 9.14.3, "Creating a Non-Mirrored Cluster SBD Partition," on page 138
- ◆ Section 9.14.4, "Creating a Mirrored Cluster SBD Partition," on page 140

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

IMPORTANT: An SBD partition requires least 20 MB of free space on a device that has been previously initialized and marked as shareable for clustering.

2 Go to the Novell Cluster Services Configuration page as described in Section 5.5.4, "Accessing the Novell Cluster Services Configuration Page in YaST," on page 87.

3 When you are prompted, enter the credentials of the LDAP administrator that is configured for the server.

This is either the LDAP administrator identity that was used when the OES Services were installed, or the identity configured after the install in Section 5.5.3, "Using Different LDAP Credentials for the Cluster Configuration," on page 87.

4 On the Novell Cluster Services Configuration page, select *Existing Cluster*.

5 In the *Cluster FDN* field, click *Browse* to open the *Browse for Container* window, locate and select the existing Cluster object where you want to add this server as a node, then click *OK*.

The comma-delimited fully distinguished name (FDN) of the cluster is automatically populated in the *Cluster FDN* field.

The *Browse* option is available in OES 2 SP3 or later.

You can also specify the fully distinguished name (FDN) (the name and eDirectory context) of the cluster where you want to add this server.

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

6 Click *Next*.

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

- ◆ **OES Common Proxy User:** If the OES Common Proxy User is enabled in eDirectory, the *Use OES Common Proxy User* check box is automatically selected and the *NCS Proxy User Name* and *Specify NCS Proxy User Password* fields are populated with the credentials of the OES Common Proxy User.

- ◆ **LDAP Admin User:** If the OES Common Proxy User is disabled in eDirectory, the *Use OES Common Proxy User* check box is automatically deselected and the *NCS Proxy User Name* and *Specify NCS Proxy User Password* fields are populated with the credentials of the LDAP Admin user. The fields are also automatically populated with the LDAP Admin credentials if you deselect the *Use OES Common Proxy User* check box.

- ◆ **Another Administrator User:** Deselect the *Use OES Common Proxy User* check box, then specify the credentials of an administrator user.

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

8 On the Configuration page from the *IP address of this node* drop-down box, select the IP address that Novell Cluster Services will use for this node.

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

9 On the Configuration page, select the *Start Cluster Services now* check box to start Novell Cluster Services software on this node after configuring it, then click *Finish*.

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

11 After the configuration is completed, click *Finish* to exit the OES Configuration page, then start Novell Cluster Services using one of these methods:

Setup Condition	Instructions
Start Cluster Services now was enabled in Step 9	Novell Cluster Services starts automatically after the configuration completes.

Start Cluster Services now was disabled in Step 9	Start Novell Cluster Services manually by using one of these methods:
	
	◆ Reboot the cluster server.
	
	◆ At a terminal console prompt, go to the `/etc/init.d` directory, then enter the following as the `root` user:
	
	`./novell-ncs start`
	
	◆ At a terminal console prompt, enter the following as the `root` user:
	
	`rcnovell-ncs start`

12 Continue with Section 5.6, "Configuring Additional Administrators," on page 92.

5.6 Configuring Additional Administrators

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

◆ The user is a trustee and has the Supervisor right to the Server object of each server node in the cluster.

◆ The user is enabled for Linux with Linux User Management. This gives the user a Linux UID in addition to the users eDirectory GUID.

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

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

You can modify the default administrator user name or password after the install by following the procedure in Section 8.9, "Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster," on page 116.

5.7 Installing or Updating the Clusters Plug-in for iManager

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

◆ Section 5.7.1, "Prerequisites for Installing or Updating the Clusters Plug-In," on page 93

◆ Section 5.7.2, "Installing the Clusters Plug-In," on page 93

5.7.1 Prerequisites for Installing or Updating the Clusters Plug-In

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

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

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

   For information, see "Storage-Related Plug-Ins for iManager" on page 64.

5.7.2 Installing the Clusters Plug-In

If you are installing a new instance of Novell iManager 2.7.4 or later and the Clusters plug-in on a computer, follow the instructions in see "Downloading and Installing Plug-in Modules" in the *Novell iManager 2.7.6 Administration Guide*. Ensure that you install the Clusters plug-in (ncsmgmt.rpm) as well as the common code for storage-related plug-ins (storagemgmt.rpm).

For the Clusters plug-in for iManager 2.7.5 or later, verify that the /var/opt/novell/iManager/nps/WEB-INF/lib directory contains the class loader file commons-lang-2.6.jar (or later version), then remove all earlier versions of the file. You must remove the earlier versions (such as commons-lang.jar or commons-lang-2.4.jar) to ensure that iManager loads the 2.6 or later version of the class loader file. Restart Tomcat when you are done.

To configure role based services for the Clusters plug-in, see "RBS Configuration" in the *Novell iManager 2.7.6 Administration Guide*.

5.7.3 Uninstalling and Reinstalling the Clusters Plug-In

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

1 In iManager, uninstall the currently installed storage-related plug-ins.

2 Copy the new version of the plug-in files to the iManager plug-ins location, manually overwriting the older version of the plug-in in the packages folder with the newer version of the plug-in.

3 In iManager, install all of the storage-related plug-ins, or install the plug-ins you need, plus the common code.

4 If you are updating the Clusters plug-in, delete the contents of the /var/opt/novell/tomcat6/work/Catalina/localhost/nps work folder.

   You must manually delete the old compiled JSPs that are cached in the nps work folder. This ensures that only the latest files are cached.

5 If you are updating with the Cluster plug-in for Novell iManager 2.7.5 or later, verify that the /var/opt/novell/iManager/nps/WEB-INF/lib directory contains the class loader file commons-lang-2.6.jar (or later version), then remove all earlier versions of the file from the directory.

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

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

```
rcnovell-tomcat6 restart
```

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

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

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

1 In iManager, click *Configuration*, then select *Role Based Services > RBS Configuration*.

2 Select the *iManager 2.x Collections* tab.

3 On the iManager 2.x Collections page, view the entry for *Role Based Services 2.novell*.

There are numbers in the columns for *Modules*, *Installed*, *Out-of-Date*, and *Not-Installed*.

4 Click the number link in the *Out-of-Date* column for *Role Based Services 2.novell*.

5 Select the check box next to *Clusters*, then click *Update*.

This updates the RBS configuration information for the Clusters plug-in.

6 Click *OK* to confirm the reinstallation.

7 After the reinstallation is complete, click *OK* to confirm the success message.

8 Log out of iManager.

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

```
rcnovell-tomcat6 restart
```

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

5.8 Patching Novell Cluster Services

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

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

2 Use the cluster leave command to remove the node from the cluster.

3 Stop Novell Cluster Services by entering

```
rcnovell-ncs stop
```

4 Apply the OES 2 Linux or Novell Cluster Services patches.

5 Start Novell Cluster Services by entering

```
rcnovell-ncs start
```

6 Use the `cluster join` command to rejoin the node to the cluster.

5.9 What's Next

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

- Chapter 8, "Configuring Cluster Policies and Priorities," on page 105
- Chapter 9, "Managing Clusters," on page 125.
- Chapter 10, "Configuring and Managing Cluster Resources," on page 147
- Chapter 12, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 173
- Chapter 13, "Configuring and Managing Cluster Resources for Shared Linux POSIX Volumes," on page 213

If you install Novell Cluster Services at the host level of an OES 2 Linux (Xen) server, you can create cluster resources for the virtual machines. For information, see Section 14.2, "Virtual Machines as Cluster Resources," on page 238.

6 Upgrading OES 2 Linux Clusters

You can upgrade a Novell Open Enterprise (OES) 2 Linux (or later) cluster to OES 2 SP3 Linux. This section describes the upgrade process and how to manage the temporarily mixed cluster during the upgrade.

6.1 Requirements for Upgrading Clusters

Ensure that your environment meets the requirements for installing Novell Cluster Services that are described in Chapter 4, "Planning for Novell Cluster Services," on page 55.

Novell Cluster Services added support for OES 2 SP3 services and file systems on the SUSE Linux Enterprise Server (SLES) 10 SP4 operating system. You can upgrade to SLES 10 SP4 by using the `move-to-sles10-sp4 patch` in the SLES patch channel. After a cluster upgrade, all nodes in the cluster must be running on the same SLES 10 release with the latest patches applied.

6.2 Upgrading OES 2 Clusters (Rolling Cluster Upgrade)

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

During a rolling cluster upgrade, one server is upgraded to the most recent version of OES 2 Linux while the other servers in the cluster continue running an older version of OES 2 Linux. Then another server can be upgraded in turn, and then another, until all servers in the cluster have been upgraded.

You can also add new OES 2 Linux servers running the most recent version of OES 2 Linux to the existing OES 2 Linux (or later) cluster, and remove the old servers from the cluster.

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

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

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

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

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

2 Stop Novell Cluster Services (`rcnovell-ncs stop`).

3 Upgrade the server by using the *Update* option on the OES 2 Linux installation.

See "Upgrading to OES 2 SP3" in the *OES 2 SP3: Installation Guide*.

4 Start Novell Cluster Services (`rcnovell-ncs start`).

5 If necessary, manually cluster migrate the resources that were previously loaded on this Server back to the upgraded server.

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

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

6 Repeat Step 2 through Step 7 for each OES 2 Linux cluster server until your entire cluster has been upgraded.

6.3 Upgrade Issues for OES 2 SP1 or Later

Ensure that any services that are available only in OES 2 SP1 or later (such as Novell CIFS or Novell AFP) are set up with preferred nodes for failover that are running OES 2 SP1 or later.

Consider the following issues when working with a mixed-release OES 2 cluster:

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

6.4 Upgrade Issues for OES 2 SP2

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

- Section 6.4.1, "Updating the iFolder Resource Template," on page 99

6.4.1 Updating the iFolder Resource Template

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

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

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

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

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

6.5 Upgrade Issues for OES 2 SP3

This section contains known issues for upgrading clusters from OES 2 SP2 Linux or earlier to OES 2 SP3.

- Section 6.5.1, "Change in Reboot Behavior for the Tolerance Setting in ldncs," on page 99
- Section 6.5.2, "Support for SLES 10 SP4," on page 99

6.5.1 Change in Reboot Behavior for the Tolerance Setting in ldncs

In OES 2 SP3, the Novell Cluster Services reboot behavior conforms to the kernel panic setting for the Linux operating system. By default the kernel panic setting is set for no reboot after a node shutdown.

During the upgrade, the $TOLERANCE setting is removed from the /opt/novell/ncs/bin/ldncs file. Any change you have made to the $TOLERANCE setting is lost after you update. You must manually set your preferred behavior by modifying the /etc/sysctl.conf file. For information, see Section 9.11, "Preventing a Cluster Node Reboot after a Node Shutdown," on page 134.

6.5.2 Support for SLES 10 SP4

Novell Cluster Services added support for OES 2 SP3 services and file systems on the SUSE Linux Enterprise Server (SLES) 10 SP4 operating system. You can upgrade to SLES 10 SP4 by using the move-to-sles10-sp4 patch in the SLES patch channel. After a cluster upgrade, all nodes in the cluster must be running on the same SLES 10 release with the latest patches applied.

7 Upgrading OES 1 Linux Clusters to OES 2 Linux

This section provides information to help you upgrade a Novell Open Enterprise (OES) 1 SP2 Linux cluster to OES 2 SP3 on SUSE Linux Enterprise Server (SLES) 10 SP4, and how to manage the temporarily mixed cluster during the upgrade.

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

* Section 7.1, "Supported Upgrade Paths from OES 1 to OES 2," on page 101
* Section 7.2, "Requirements and Guidelines for Upgrading Clusters from OES 1 Linux and OES 2 Linux," on page 101
* Section 7.3, "Upgrading Existing OES 1 Linux Cluster Nodes to OES 2 (Rolling Cluster Upgrade)," on page 102
* Section 7.4, "Adding New OES 2 Linux Cluster Nodes to Your OES 1 Linux Cluster," on page 103
* Section 7.5, "Modifying Cluster Resource Scripts for Mixed OES 1 Linux and OES 2 Linux Clusters," on page 103

7.1 Supported Upgrade Paths from OES 1 to OES 2

This section describes how to upgrade from OES 1 SP2 (32-bit) to OES 2 SP3 (32-bit). For information about supported upgrade paths for OES 2, see "Supported Upgrade Paths" in the *OES 2 SP3: Installation Guide*.

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

In addition to the Chapter 4, "Planning for Novell Cluster Services," on page 55, consider the following rules and recommendations for mixed OES 1 Linux and OES 2 Linux clusters:

* Mixed OES 1 Linux and OES 2 Linux clusters should be considered a temporary configuration that exists only during an upgrade.
* Adding an OES 1 Linux cluster node to a mixed-node OES 1 Linux and OES 2 Linux cluster is not supported
* If you have configured resource monitoring for a resource running on an OES 2 Linux cluster node, resource monitoring does not function if the resource fails over or is migrated to an OES 1 Linux cluster node.
* The use of EVMS is recommended for upgrading file system resources.

- You should ensure that all resource policies are configured to the settings that existed before the upgrade.

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

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

Performing a rolling cluster upgrade from OES 1 Linux to OES 2 Linux lets you keep your cluster up and running and lets your users continue to access cluster resources while the upgrade is being performed.

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

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

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

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

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

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

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

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

3 Stop Novell Cluster Services (`rcnovell-ncs stop`).

4 Upgrade the server by using the *Update* option on the OES 2 Linux installation, making sure to install the components that are currently installed on the server that you noted in Step 1.

See "Upgrading to OES 2 SP3" in the *OES 2 SP3: Installation Guide*.

5 Repeat Step 1 through Step 4 for each OES 1 Linux cluster server until your entire cluster has been upgraded to OES 2.

6 Start Novell Cluster Services (`rcnovell-ncs start`).

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

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

- The failback mode for the resources was set to Auto.

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

8 After the last OES 1 cluster node has been upgraded to OES 2 and only OES 2 nodes remain, the conversion is complete.

9 (Optional) Enable Monitoring for the cluster resources and configure the Monitor scripts as described in Section 10.6, "Enabling Monitoring and Configuring the Monitor Script," on page 155.

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

The process for adding a new OES 2 Linux cluster node to an existing OES 1 Linux cluster is the same as for adding an OES 1 cluster node to an OES 1 Linux cluster and adding a new OES 2 cluster node to an existing OES 2 cluster. See Section 5.4, "Installing Novell Cluster Services," on page 81.

However, you should be aware of the rules that apply to mixed OES 1 and OES 2 clusters. See Section 7.2, "Requirements and Guidelines for Upgrading Clusters from OES 1 Linux and OES 2 Linux," on page 101.

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

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

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

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

8 Configuring Cluster Policies and Priorities

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

8.1 Understanding Cluster Settings

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

8.1.1 Cluster Policies

Table 8-1 describes the configurable cluster policies. You can manage cluster policies in iManager by going to the *Clusters > Cluster Options > Policies* page.

Table 8-1 *Cluster Policies*

Property	Description
Cluster IP address	Specifies the IP address for the cluster. You specify the IP address when you install Novell Cluster Services on the first node of the cluster. Rarely, you might need to modify this value. For information, see Section 8.8, "Modifying the Cluster IP Address and Port Properties," on page 115.
Port	Specifies the port used for cluster communication. The default cluster port number is 7023, and is automatically assigned when the cluster is created. You might need to modify this value if there is a port conflict. For information, see Section 8.8, "Modifying the Cluster IP Address and Port Properties," on page 115.
Quorum membership	Specifies number of nodes that must be up and running in the cluster in order for cluster resources to begin loading. Specify a value between 1 and the number of nodes. For instructions, see Section 8.2, "Configuring Quorum Membership and Timeout Properties," on page 107.
Quorum timeout	Specifies the maximum amount of time in seconds to wait for the specified quorum to be met before cluster resources begin loading on whatever number of nodes are actually up and running. For instructions, see Section 8.2, "Configuring Quorum Membership and Timeout Properties," on page 107.
Email notification	Enables or disables email notification for the cluster. If it is enabled, you can specify up to eight administrator email addresses for cluster events notification. For instructions, see Section 8.4, "Configuring Cluster Event Email Notification," on page 109.

8.1.2 Cluster Protocols Properties

Table 8-2 describes the configurable cluster protocols properties that govern inter-node communication transmission and tolerances. You can manage cluster protocols policies in iManager by going to the *Clusters > Cluster Options > Protocols* page. For instructions, see Section 8.3, "Configuring Cluster Protocol Properties," on page 108.

Table 8-2 *Cluster Policies*

Property	Description
Heartbeat	Specifies the interval of time in seconds between signals sent by each of the non-master nodes in the cluster to the master node to indicate that it is alive.
Tolerance	Specifies the maximum amount of time in seconds that a master node waits to get an alive signal from a non-master node before considering that node to have failed and removing it from the cluster. The default is 8 seconds.
Master watchdog	Specifies the interval of time in seconds between alive signals sent from the master node to non-master nodes to indicate that it is alive.

Property	Description
Slave watchdog	Specifies the maximum amount of time in seconds that the non-master nodes wait to get an alive signal from the master node before considering that the master node has failed, assigning another node to become the master node, and removing the old master node from the cluster.
Maximum retransmits	This value is set by default and should not be changed.

8.2 Configuring Quorum Membership and Timeout Properties

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

1 In iManager, select *Clusters*, then select *Cluster Options*.

2 Specify the cluster name, or browse and select the Cluster object.

3 Click the *Properties* button under the cluster name.

4 Click the *Policies* tab.

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

For information, see Section 8.2.1, "Quorum Triggers (Number of Nodes)," on page 107.

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

For information, see Section 8.2.2, "Quorum Triggers (Timeout)," on page 107.

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

8 Restart Cluster Services by entering the following at a terminal console prompt:

```
rcnovell-ncs restart
```

8.2.1 Quorum Triggers (Number of Nodes)

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

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

8.2.2 Quorum Triggers (Timeout)

Timeout specifies the amount of time to wait for the number of servers defined in the *Number of Nodes* field to be up and running. If the timeout period elapses before the quorum membership reaches its specified number, resources automatically start loading on the servers that are currently up and

running in the cluster. For example, if you specify a *Number of Nodes* value of 4 and a timeout value equal to 30 seconds, and after 30 seconds only two servers are up and running in the cluster, resources begin to load on the two servers that are up and running in the cluster.

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

8.3 Configuring Cluster Protocol Properties

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

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

1 In iManager, select *Clusters*, then select *Cluster Options*.

2 Specify the cluster name, or browse and select the Cluster object.

3 Click the *Properties* button under the cluster name.

4 Click the *Protocols* tab.

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

5 Specify values for the cluster protocols properties.

For information, see the following:

- ◆ Heartbeat
- ◆ Tolerance
- ◆ Master Watchdog
- ◆ Slave Watchdog
- ◆ Maximum Retransmits

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

7 Restart all nodes in the cluster to make the changes take effect.

8.3.1 Heartbeat

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

8.3.2 Tolerance

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

8.3.3   Master Watchdog

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

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

8.3.4   Slave Watchdog

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

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

8.3.5   Maximum Retransmits

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

8.4   Configuring Cluster Event Email Notification

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

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

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

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

To bind a service to its resource's secondary IP address, refer to the cluster documentation for the service. For example, the GroupWise Internet Agent (GWIA) uses SMTP. Because both Postfix and the GWIA default to using port 25, you must configure the GWIA to bind exclusively to its resource's secondary IP address in order to avoid a port conflict between Postfix and the GWIA. For information

about how to set up the secondary IP address and to bind the GWIA to it, see the following sections in the *GroupWise 2012 Interoperability Guide* (http://www.novell.com/documentation/groupwise2012/gw2012_guide_interop/data/a20gkue.html):

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

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

1 In iManager, select *Clusters*, then select *Cluster Options*.

2 Specify the cluster name, or browse and select the Cluster object.

3 Click the *Properties* button under the cluster name.

4 Click the *Policies* tab.

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

6 If you enable email notification, add up to eight email addresses in the field provided.

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

7 If you enable email notification, specify the type of cluster events you want administrators to receive messages for.

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

8 If you enable email notification, specify whether you want to receive notification of all cluster state changes in XML format by selecting the *XML Messages* option.

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

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

8.5 Configuring Cascade Failover Prevention

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

In OES 2 SP3 and later, Novell Cluster Services provides a Cascade Failover Prevention function. It detects if a node has failed because of a bad cluster resource, and prevents that bad resource from failing over to other servers in the cluster.

- Section 8.5.1, "Understanding the Cascade Failover Prevention Quarantine," on page 111
- Section 8.5.2, "Releasing a Resource from Quarantine," on page 112
- Section 8.5.3, "Enabling or Disabling Cascade Failover Prevention," on page 112

8.5.1 Understanding the Cascade Failover Prevention Quarantine

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

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

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

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

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

1. Traces the history of node failures for the suspected bad resource. This includes:
   - What node the resource was running on or loading on
   - If the node failed
   - The state the resource was in when the node failed
   - If there were other resources trying to load when the node failed

2. Repeats the above process until one of the following happens:
   - The end of the cluster log file is reached
   - Enough consecutive node failures are found
   - Found that the node did not fail
   - Found that the whole cluster was down
   - The entries in the log file are more than 365 days old

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

For a particular resource, there is no fixed number of node failures that triggers a quarantine. Generally, three consecutive node failures triggers a resource quarantine. However, the actual number of failures considered can vary based on other factors like how many nodes are in the cluster and what other resources are doing at the time of the node failures. For example, if no other resources are in a running or loading state when a resource loads but never reaches a running state, then two consecutive node failures may trigger resource quarantine.

Factors that might contribute to a resource being quarantined include:

- A large number of consecutive node failures (generally three or more)
- No other resources are causing node failures
- The resource never reaches a running state

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

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

The resource quarantine is disabled if:

◆ Cascade Failover Prevention is turned off.

For information, see

◆ There is no shared storage (SAN) or SBD partition.

◆ There are enough nodes in the cluster to form a quorum.

8.5.2 Releasing a Resource from Quarantine

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

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

1 Log in to the node as the `root` user.

2 Disable Cascade Failover Prevention. See .

3 Re-enable Cascade Failover Prevention. See .

8.5.3 Enabling or Disabling Cascade Failover Prevention

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

◆ A value of 0 disables the function.

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

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

◆

◆

Disabling Cascade Failover Prevention

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

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

◆ To stop using the Cascade Failover Prevention function in the cluster, you can disable the function on every node in the cluster.

To disable Cascade Failover Prevention:

1 Log in to the node as the `root` user.

2 Navigate to the `/etc/modprobe.d` folder.

3 In a text editor, create a configuration file (such as `novell-ncs.conf`) under the `/etc/ modprobe.d` folder.

If the file already exists, open it.

4 Add the following content to the file. If the line already exists, change the flag value to 0.

```
options crm ncs_cascade_failover_detection_flag=0
```

The flag value 0 disables Cascade Failover Prevention for the node.

5 Save the file, then close the text editor.

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

```
rcnovell-ncs restart
```

7 Repeat this procedure on other nodes where you want to disable Cascade Failover Prevention.

8 (Optional) Re-enable Cascade Failover Prevention as described in .

Manually Enabling Cascade Failover Prevention

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

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

To enable Cascade Failover Prevention:

1 Log in to the node as the `root` user.

2 Navigate to the `/etc/modprobe.d` folder.

3 Use one of the following methods to enable Cascade Failover Prevention:

- **Delete the file:** Delete the `/etc/modprobe.d/novell-ncs.conf` file that you created in .

- **Change the value to 2:** In a text editor, open the `/etc/modprobe.d/novell-ncs.conf` file that you created in , change the flag value from 0 to 2, then save the file.

  ```
  options crm ncs_cascade_failover_detection_flag=2
  ```

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

- **Delete the line:** In a text editor, open the `/etc/modprobe.d/novell-ncs.conf` file that you created in , remove the `ncs_cascade_failover_detection_flag` line, then save the file.

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

```
rcnovell-ncs restart
```

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

8.6 Viewing the Cluster Node Properties

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

1 In iManager, select *Clusters*, then select *Cluster Options*.

2 Specify the cluster name, or browse and select the Cluster object.

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

4 View the desired information, then click *OK*.

8.7 Changing the NCS Proxy User Assignment in the NCS_Management_Group

When you configure a cluster, the NCS_Management group is created in the Cluster object container. The group contains the user that you specify for the NCS Proxy User on the Proxy User Configuration page as described in Step 3 in Section 5.5.5, "Configuring a New Cluster," on page 88:

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

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

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

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

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

2 At the console prompt, enter

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

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

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

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

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

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

8.8   Modifying the Cluster IP Address and Port Properties

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

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

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

2 Specify the cluster name, or browse and select the Cluster object.

3 Click the *Properties* button under the cluster name.

4 Click the *Policies* tab.

5 Specify the new value for the IP address or port.

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

The action also automatically changes the cluster IP address in the master IP resource scripts.

Because you are changing the IP address, the resulting response is a communications error. This is expected.

7 Exit iManager, then close the web browser.

8 Restart NCS.

```
rcnovell-ncs restart
```

A cluster restart is required before iManager can continue to manage the cluster with its new master IP address.

```
avalon:~ # rcnovell-ncs restart
Stopping Novell Cluster Services
Leaving...
No longer a member of cluster

#------------------------------------------------------------+
#                                                            |
#   Please wait for the Novell Cluster Services KOs to unload  |
#                                                            |
#------------------------------------------------------------+

Novell Cluster Services Unloaded

                                                         done

Starting Novell Cluster Services
Mounting adminfs at /admin ... already mounted          done
                                                         done
avalon:~ #
avalon:~ # Joining...
Now a member of

        Cluster cl242
        This node avalon [ epoch 0 master node avalon ]
        Cluster nodes [ avalon ]
```

9 In iManager, visually verify that the IP address change was made by looking at the *Properties > Policies* tab and in the scripts for the Master IP resource.

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

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

- Section 8.9.1, "Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster," on page 116
- Section 8.9.2, "Moving a Cluster or Changing IP Addresses of Cluster Nodes and Resources," on page 118

8.9.1 Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster

You can modify the administrator credentials or LDAP server settings that you assigned when you created the cluster. You must modify this cluster configuration information in the following cases:

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

You should list the LDAP servers in the following order:

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

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

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

1 Ensure that the IP addresses and administrator user name that you plan to use meet the requirements specified in Section 4.2, "IP Address Requirements," on page 58.

2 Log in as the root user to the master node of the cluster.

3 In a text editor, create a text file, specify the configuration information for the Novell Cluster Services cluster in it, then save the file.

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

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

The following lines are the content of a sample configuration file for a Novell Cluster Services cluster when you have a single LDAP server.

```
CONFIG_NCS_CLUSTER_DN="cn=svr1_oes2_cluster.o=context"
CONFIG_NCS_LDAP_IP="10.1.1.102"
CONFIG_NCS_LDAP_PORT="636"
CONFIG_NCS_ADMIN_DN="cn=admin.o=context"
CONFIG_NCS_ADMIN_PASSWORD="password"
```

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

```
CONFIG_NCS_CLUSTER_DN="cn=svr1_oes2_cluster.o=context"
CONFIG_NCS_LDAP_INFO="ldaps://10.1.1.102:636,ldaps://10.1.1.101:636"
CONFIG_NCS_ADMIN_DN="cn=admin.o=context"
CONFIG_NCS_ADMIN_PASSWORD="password"
```

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

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

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

5 Delete the configuration file that you created.

6 For each of the other nodes in the cluster, log in as the root user for the node, then repeat Step 3 to Step 5.

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

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

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

8.9.2 Moving a Cluster or Changing IP Addresses of Cluster Nodes and Resources

If you move a cluster to a different subnet, you must change the IP addresses of the cluster nodes and the cluster resources, information about the LDAP servers used by the cluster, and possibly the administrator credentials for the cluster.

When you move the cluster to a new IP subnet, you must replace the existing unique static IP addresses with ones that are valid in that subnet. You can make the IP address changes in the old location or the new location. If you start the servers in the different IP subnet with the old IP addresses, the cluster does not come up until you make the changes described in this section.

To modify the IP addresses of servers being used in a Novell Cluster Services cluster, perform the following tasks in the order given:

- "Prerequisites" on page 118
- "Changing the IP Addresses of Cluster Resources" on page 118
- "Changing the IP Addresses of Servers in a Cluster" on page 120
- "Changing the IP Address of the Master IP Resource" on page 120
- "Modifying the Cluster Configuration Information" on page 121

Prerequisites

Before you begin, ensure that the IP addresses that you plan to use meet the requirements specified in Section 4.2, "IP Address Requirements," on page 58. Ensure that the administrator user name that you will use in the new location has sufficient rights as described in Section 4.1.3, "Cluster Administrator or Administrator-Equivalent User," on page 57.

Changing the IP Addresses of Cluster Resources

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

1 Offline the cluster resources whose IP addresses are changing.

    1a In iManager, click *Cluster*s, then click *Cluster Manager*.

    1b Browse to locate and select the Cluster object of the cluster you want to manage.

    1c Select the check boxes next to the resources you want to take offline, then click *Offline*.

2 For each NSS pool cluster resource, modify the IP address on the resource's Protocols page.

    2a On the Cluster Manager or Cluster Options page, click the pool cluster resource to open its Properties page.

    2b Click the *Protocols* tab.

    2c Modify the IP address for the NSS pool cluster resource.

    2d Click *OK* to save the changes.

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

       Do not online the resource at this time.

3 For each non-NSS cluster resource, modify the IP address information for the RESOURCE_IP variable as needed in the resource's load, unload, and monitor scripts.

  3a On the Cluster Manager page or Cluster Options page, click the non-NSS cluster resource to open its Properties page.

  3b Click the *Scripts* tab, then click the *Load Script* link.

  3c Edit the script by replacing the RESOURCE_IP variable's value with new IP address. You might also need to edit the values used in the command lines if you did not use the variable there.

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

  3d Click *Apply* to save the changed script.

  3e Make similar changes to the non-NSS resource's *Unload Script* and *Monitor Script*.

  3f Repeat this process for each non-NSS cluster resource.

  Do not online the resources at this time.

4 For each Linux cluster resource that has an NCP virtual server object, delete and re-create the NCS:NCP Server Object with the new IP address.

  4a In iManager, select the *Directory Administration > Delete Object*, select the NCS:NCP Server object for the Linux resource, then click *OK*.

  4b Re-create the NCS:NCP Server object with the new IP address:

    4b1 On the master cluster node, open a terminal console, then log in as the `root` user.

    4b2 In the console, use the `cd` command to go to the `/opt/novell/ncs/bin` directory.

    4b3 At the terminal console prompt, enter

    ```
    ./ncs_ncpserv.py -c lx_volumename -i resource_ip_address
    ```

    Replace the *lx_volumename* and *resource_ip_address* with the information for your particular solution.

    Do not use periods in cluster resource names. Novell clients interpret periods as delimiters. If you use a space in a cluster resource name, that space is converted to an underscore.

    For example, to create the NCS:NCP Server object for the `lxvol44` cluster resource where the IP address is 10.10.10.44 and the cluster context is `ou=clusters,ou=city,o=mycompany`, enter

    ```
    ./ncs_ncpserv.py -c lxvol44 -i 10.10.10.44
    ```

    The confirmation message is displayed:

    ```
    NCP Server
    'cn=cluster1_lxvol44_server,ou=clusters,ou=city,o=mycompany' created.
    ```

5 Stop Novell Cluster Services for every node in the cluster by entering the following at the terminal console prompt as the `root` user:

```
rcnovell-ncs stop
```

6 Continue with .

Changing the IP Addresses of Servers in a Cluster

After preparing the cluster resources for the IP address change and stopping Novell Cluster Services (see "Changing the IP Addresses of Cluster Resources" on page 118), you are ready to change the IP addresses of the servers in the cluster.

1 For each server in the cluster, change the server's IP address by following the instructions "Changing an OES 2 SP3 Server's IP Address" in the *OES 2 SP3: Planning and Implementation Guide*.

2 The server IP address changes are not complete until you make those changes known to Novell Cluster Services and eDirectory. Continue with "Changing the IP Address of the Master IP Resource" on page 120.

Changing the IP Address of the Master IP Resource

1 In iManager, click *Directory Administration*, then click *Modify Object*.

2 Browse to locate and select the Master IP Address Resource object of the cluster you want to manage.
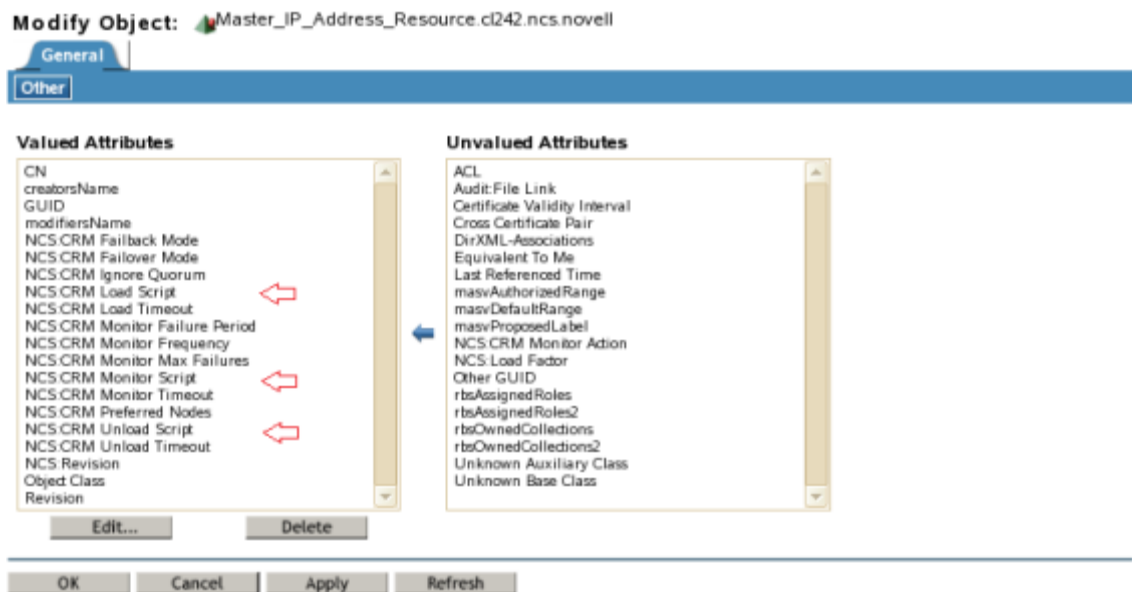
3 In the *Valued Attributes* list, select the attribute *NCS:CRM Load Script*, click *Edit*, modify the IP address in the script, then click *OK*.

Edit Attribute

NCS:CRM Load Script:

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

OK   Cancel

4 In the *Valued Attributes* list, select the attribute *NCS:CRM Unload Script*, click *Edit*, modify the IP address in the script, then click *OK*.

5 In the *Valued Attributes* list, select the attribute *NCS:CRM Monitor Script*, click *Edit*, modify the IP address in the script, then click *OK*.

6 At the bottom of the page, click *OK* to save the changes.

7 Continue with .

Modifying the Cluster Configuration Information

Before restarting Novell Cluster Services, you must update the cluster configuration information in Novell Cluster Services and eDirectory with the new IP addresses. You might also need to update the IP address information for the LDAP server and administrator credentials that the cluster uses in the new subnet.

1 If the cluster is using a different LDAP server or administrator in the new IP subnet, change the LDAP server IP address and administrator credentials for the cluster in the Novell Cluster Services configuration settings.

Follow the procedure in .

2 For each node in the cluster, modify the *NCS: Network Address* attribute of its Cluster Node object.

  2a In iManager, click *Directory Administration*, then click *Modify Object*.

  2b Browse to locate and select the Cluster Node object of the cluster node you want to manage.

  2c In the *Valued Attributes* list, select the attribute *NCS: Network Address*, click *Edit*, modify the IP address, then click *OK*.

Edit Attribute

NCS:Network Address:

IP: 10.10.10.242

OK   Cancel

  2d Repeat this process for each node in the cluster.

3 For the cluster container of the cluster you want to manage, modify the *NCS: Network Address* and *Network Address* attributes of its Cluster object to specify the new IP address information. Both *TCP* and *UDP* addresses need to be replaced.

    3a In iManager, click *Directory Administration*, then click *Modify Object*.

    3b Browse to locate and select the Cluster object of the cluster you want to manage.
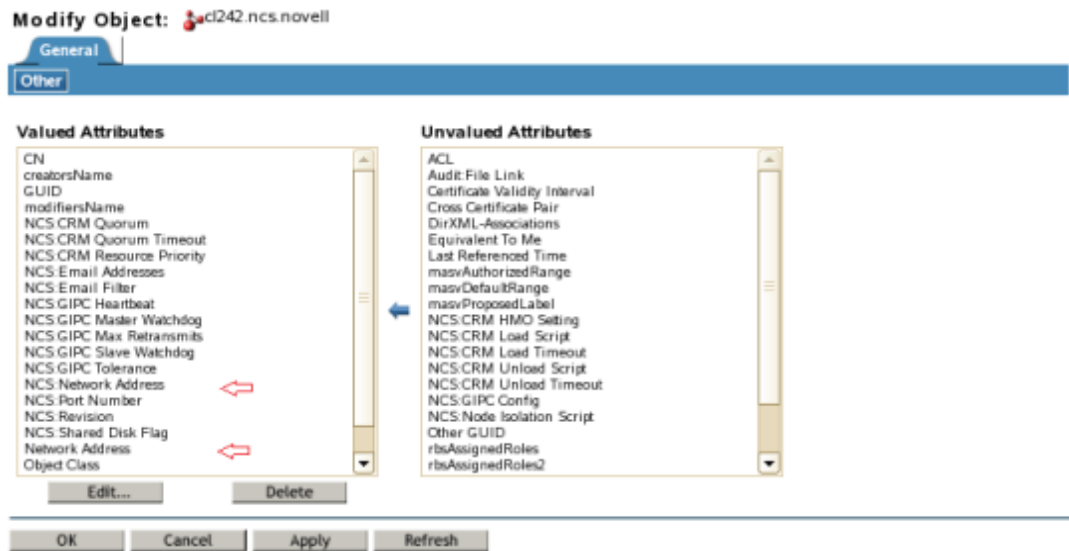
    3c In the *Valued Attributes* list, select the attribute *NCS: Network Address* (the attribute for the TCP address), click *Edit*, modify the IP address, then click *OK*.

    3d In the *Valued Attributes* list, select the attribute *Network Address* (the attribute for the UDP address), click *Edit*, modify the IP address, then click *OK*.

4 Ensure that LDAP server is running before restarting Novell Cluster Services.

> **IMPORTANT:** Novell Cluster Services requires LDAP.

5 Ensure that NSS is running if there are NSS cluster resources that you will bring online.

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

```
rcnovell-ncs start
```

7 Online the cluster resources:

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

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

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

8.10 What's Next

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

- Chapter 10, "Configuring and Managing Cluster Resources," on page 147
- Chapter 11, "Creating and Managing Service Cluster Resources," on page 171

- Chapter 12, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 173
- Chapter 13, "Configuring and Managing Cluster Resources for Shared Linux POSIX Volumes," on page 213

For information about managing the cluster, see Chapter 9, "Managing Clusters," on page 125.

9 Managing Clusters

After you have installed, set up, and configured Novell Cluster Services for your specific needs and configured cluster resources, use the information in this section to help you effectively manage your cluster. This section provides instructions for migrating resources, identifying cluster and resource states, and customizing cluster management.

IMPORTANT: For information about using console commands to manage a cluster, see Appendix A, "Console Commands for Novell Cluster Services," on page 263.

9.1 Starting and Stopping Novell Cluster Services

Novell Cluster Services automatically starts after it is installed. Novell Cluster Services also automatically starts when you reboot your Novell Open Enterprise Server (OES) 2 Linux server.

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

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

```
chkconfig open-iscsi on
```

9.1.1 Starting Novell Cluster Services

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

1 Open a terminal console, then log in as the `root` user.

2 Use one of the following methods to start Novell Cluster Services:

  ◆ At the terminal console prompt, go to the `/etc/init.d` directory and enter

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

  ◆ At the terminal console prompt, enter

    ```
    rcnovell-ncs start
    ```

9.1.2 Stopping Novell Cluster Services

1 Open a terminal console, then log in as the `root` user.

2 Use one of the following methods to stop Novell Cluster Services:

  ◆ Go to the `/etc/init.d` directory and enter

    ```
    ./novell-ncs stop
    ```

  ◆ At the terminal prompt, enter

    ```
    rcnovell-ncs stop
    ```

9.1.3 Enabling and Disabling the Automatic Start of Novell Cluster Services

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

To cause Novell Cluster Services to not start automatically after a server reboot:

1 Open a terminal console, then log in as the `root` user.

2 Enter the following at a Linux terminal console:

  ```
  chkconfig novell-ncs off
  ```

3 Reboot the server.

4 After rebooting, you must manually start Novell Cluster Services by entering

  ```
  rcnovell-ncs start
  ```

To cause Novell Cluster Services to resume starting automatically after a server reboot:

1 Open a terminal console, then log in as the `root` user.

2 Enter the following at a Linux terminal console:

  ```
  chkconfig novell-ncs on
  ```

3 Reboot the server.

9.2 Monitoring Cluster and Resource States

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

1 In iManager, click *Cluster*, then click *Cluster Manager*.

2 Type the name of the desired cluster, or browse to locate and select the Cluster object.

A list of resources and resource states displays.

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

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

Table 9-1 *Cluster Operating States*

State	Icon	Description
Normal		A green ball indicates that the server or resource is online or running.
Stopped		A red ball with a horizontal white line indicates that the node is stopped.
Offline		A white ball with a horizontal red line indicates that the node is offline.
Critical		A white ball with a red X indicates that the node has failed or is comatose.
Warning		A white ball with a yellow diamond indicates that an alert condition has occurred, and the resource needs administrator attention.
Unknown		A yellow ball with a question mark indicates that the state of the node is unknown.

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

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

Table 9-2 identifies the different resource states and gives descriptions and possible actions for each state.

Table 9-2 *Cluster Resource States*

Resource State	Description	Possible Actions
Alert Alert types are: Failback Alert Failover Alert Start Alert	Alerts can result if the Start mode, Failover mode, or Failback mode for the resource has been set to Manual. The resource is waiting for the administrator to intervene to start, fail over, or fail back the resource on the specified server.	On the Cluster Manager page in Novell iManager, click the *Respond to Alert* option to resolve the alert. Depending on the resource state, you are prompted to start, fail over, or fail back the resource. If you attempt to offline a resource that is in the *Start Alert* state, nothing happens and the following message is displayed: `This operation cannot be completed. It is only available when the resource is in an offline state.` You must clear the *Start Alert* before you can offline the resource. After the resource is offline, you can online the resource.
Comatose	The resource is not running properly and requires administrator intervention.	Click the *Comatose* 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 *Offline* status indicator and, if desired, click the *Online* button to load the resource on the best node possible, given the current state of the cluster and the resource's preferred nodes list.
Quorum Wait	The resource is waiting for the quorum to be established so it can begin loading.	None.
Running	The resource is in a normal running state.	Click the *Running* status indicator and choose to either migrate the resource to a different server in your cluster or unload (bring offline) the resource.
Unassigned	There isn't an assigned node available that the resource can be loaded on.	Click the *Unassigned* status indicator and, if desired, offline the resource. Offlining the resource prevents it from running on any of its preferred nodes if any of them join the cluster.
Unloading	The resource is in the process of unloading from the server it was running on.	None.

9.3 Generating a Cluster Configuration Report

A cluster configuration report can help you diagnose problems with the cluster nodes and resources. The report contains the following information for the selected cluster:

Field	Description
Cluster name	The dot delimited distinguished name of the selected cluster (*clustername.context*).
Date	The date and time the report was generated.
Cluster status	Current Nodes Quorum Triggers List of resources (Name, State, Location, Lives, Up Since)
Cluster options	Quorum Triggers (Number of nodes, Timeout) Protocol Setting (Heartbeat, Tolerance, Master Watchdog, Slave Watchdog, Maximum retransmits) List of nodes (Name, State, Number, IP Address, and Distinguished Name) Resource Mutual Exclusion Groups
Resource details	Information listed (as it applies) by resource for the master IP address resource and each of the cluster resources: Policies (Follows masters, Ignore quorum, Start mode, Failover mode, Failback mode, Assigned nodes) Load script and timeout Unload script and timeout Monitoring script Protocols Virtual server name CIFS server name IP address Advertising protocols

To generate a cluster configuration report:

1 In iManager, select *Clusters > Cluster Manager*.

2 Browse to select the cluster you want to manage.

The *Run Report* button is dimmed until the page is refreshed with the cluster information.

3 Click *Run Report*.

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

4 Use the scroll bar to view the report in the browser.

5 (Optional) Use the browser options to print the report.

6 (Optional) Save the report to a text file.

    6a Click the cursor in the browser window where the report results are displayed, press Ctrl+A to select all, press Ctrl+C to copy the report to the clipboard.

    6b In a text editor, open a new document, then press Ctrl+V to paste the report in the document.

    6c Save the report document.

9.4 Cluster Migrating Resources to Different Nodes

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

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

Using iManager

1 In iManager, click *Clusters*, then click *Cluster Manager*.

2 Browse to locate and select the Cluster object of the cluster you want to manage.

3 Select the check box next to the resource you want to migrate and click *Migrate*.

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

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

Using Console Commands

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

```
cluster migrate resource_name node_name
```

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

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

- Section 9.5.1, "Guidelines for Onlining and Offlining Resources," on page 130
- Section 9.5.2, "Using iManager to Online and Offline Resources," on page 131
- Section 9.5.3, "Using the Cluster Online Command," on page 131
- Section 9.5.4, "Using the Cluster Offline Command," on page 131
- Section 9.5.5, "Using the Cluster Migrate Command," on page 131

9.5.1 Guidelines for Onlining and Offlining Resources

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

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

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

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

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

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

9.5.2 Using iManager to Online and Offline Resources

1 In iManager, select *Cluster*s, then select *Cluster Manager*.

2 Browse to locate and select the Cluster object of the cluster you want to manage.

A list of resources and resource states displays.

3 Select the check box next to the resource you want to manage, then click *Online* or click *Offline*.

9.5.3 Using the Cluster Online Command

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

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

```
cluster online resource_name
```

To bring a resource online on a specific active node:

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

```
cluster online resource_name node_name
```

9.5.4 Using the Cluster Offline Command

To take a resource offline:

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

```
cluster offline resource_name
```

9.5.5 Using the Cluster Migrate Command

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

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

```
cluster migrate resource_name node_name
```

9.6 Removing (Leaving) a Node from the Cluster

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

- Section 9.6.1, "Removing One Node at a Time," on page 132
- Section 9.6.2, "Removing All Nodes at the Same Time," on page 132

9.6.1 Removing One Node at a Time

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

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

```
cluster leave
```

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

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

9.6.2 Removing All Nodes at the Same Time

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

Beginning in OES 2 SP3, you are prompted to confirm the cluster down.

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

```
cluster down
```

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

9.7 Joining a Node to the Cluster

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

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

```
cluster join
```

9.8 Configuring the EVMS Remote Request Timeout

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

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

To modify the EVMS timeout:

1 Open the `/etc/evms.conf` file in a text editor.

2 Edit the following section to increase the `remote_request_timeout` value:

```
engine {

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

3 Save the file.

4 At the terminal console prompt, enter

```
evms_activate
```

9.9 Shutting Down Linux Cluster Servers When Servicing Shared Storage

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

9.10 Enabling or Disabling Cluster Maintenance Mode

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

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

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

```
cluster maintenance on
```

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

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

```
cluster maintenance off
```

9.11 Preventing a Cluster Node Reboot after a Node Shutdown

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

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

* Section 9.11.1, "OES 2 SP2 with Patches and Later," on page 134
* Section 9.11.2, "OES 2 SP2 Release Version and Earlier," on page 135

9.11.1 OES 2 SP2 with Patches and Later

Beginning in the OES 2 SP2 Maintenance Patch for May 2010, the Novell Cluster Services reboot behavior conforms to the kernel panic setting for the Linux operating system. By default the kernel panic setting is set for no reboot after a node shutdown.

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

1 As the `root` user, open the `/etc/sysctl.conf` file in a text editor.

2 If the `kernel.panic` token is not present, add it.

```
kernel.panic = 0
```

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

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

```
kernel.panic = 0
```

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

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

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

```
kernel.panic = 60
```

4 Save your changes.

9.11.2 OES 2 SP2 Release Version and Earlier

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

1 Open the `opt/novell/ncs/bin/ldncs` file in a text editor.

2 Find the following line:

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

3 Replace $TOLERANCE with a value of 0 to cause the server to not automatically reboot after a shutdown.

4 After editing the `ldncs` file, you must reboot the server to cause the change to take effect.

9.12 Adding a Node That Was Previously in the Cluster

1 If your storage array and devices are not configured, configure them before you install Novell Cluster Services.

2 If necessary, install OES 2 Linux and Novell Cluster Services, including the latest Service Pack on the server using the same node name and IP address of the node that was previously in the cluster.

3 If the Cluster object for the server is still present, delete the object.

For information, see Section 9.13, "Deleting a Cluster Node from a Cluster," on page 135.

4 Run the Novell Cluster Services installation by following the procedure outlined in Section 5.4.2, "Installing Novell Cluster Services on an Existing OES 2 Linux Server," on page 83.

The node assumes its former identity.

9.13 Deleting a Cluster Node from a Cluster

To permanently remove a cluster node from a cluster:

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

```
cluster leave
```

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

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

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

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

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

    3a Select *eDirectory Administration > Delete Objects*.

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

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

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

5 If the deleted cluster node persists in iManager, do one of the following:

- Run the following command as the `root` user on the master cluster node. It is safe to run the command on an active cluster:

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

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

```
rcnovell-ncs restart
```

9.14 Creating or Deleting Cluster SBD Partitions

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

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

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

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

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

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

IMPORTANT: For instructions about setting up the SBD partition during the Novell Cluster Services configuration, see Section 5.5.5, "Configuring a New Cluster," on page 88.

- Section 9.14.1, "Prerequisites for Creating an SBD Partition," on page 137
- Section 9.14.2, "Before Creating an Cluster SBD Partition," on page 137
- Section 9.14.3, "Creating a Non-Mirrored Cluster SBD Partition," on page 138
- Section 9.14.4, "Creating a Mirrored Cluster SBD Partition," on page 140
- Section 9.14.5, "Deleting a Non-Mirrored Cluster SBD Partition," on page 142
- Section 9.14.6, "Deleting a Mirrored Cluster SBD Partition," on page 142
- Section 9.14.7, "Removing a Segment from a Mirrored Cluster SBD Partition," on page 143

9.14.1 Prerequisites for Creating an SBD Partition

You must have a shared disk system (such as a Fibre Channel SAN or an iSCSI SAN) connected to your cluster nodes before attempting to create a cluster partition. For information, see Section 4.7, "Shared Disk Configuration Requirements," on page 70 for more information.

You need one small LUN in your storage array to use exclusively for the SBD partition. You can create another LUN of the same size to use as its mirror. The device should have at least 20 MB of free available space. About 4 MB of space that is used to store its shareable state information. For information about the SBD partition requirements, see "SBD Partitions" on page 70.

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

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

NOTE: For an iSCSI SAN, a very small iSCSI device might rarely have uncommon CHS (cylinder-head-sector) geometry values. NSS prefers at least 32 sectors per track. If there are fewer than 32 sectors per track, NSS tools using EVMS can fail to create the partition or to mark the device as shareable. You can use fdisk to check for valid CHS values before you initialize and share the device. If necessary, you can use fdisk to set the sectors for the device (such as `/dev/sdx`) to 32:

```
fdisk -H 64 -S 32 /dev/sdx
```

9.14.2 Before Creating an Cluster SBD Partition

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

IMPORTANT: If a cluster SBD partition already exists, do not create another one. Before you create a new SBD partition, delete the existing one as described in Section 9.14.2, "Before Creating an Cluster SBD Partition," on page 137.

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

```
sbdutil -f
```

This tells you whether an SBD partition exists and identifies the device on the SAN where the SBD partition is located.

2 If the SBD partition already exists, use one of the following methods to delete the existing partition before attempting to create another one:

- Section 9.14.5, "Deleting a Non-Mirrored Cluster SBD Partition," on page 142
- Section 9.14.7, "Removing a Segment from a Mirrored Cluster SBD Partition," on page 143

9.14.3 Creating a Non-Mirrored Cluster SBD Partition

If you did not create a cluster partition during the Novell Cluster Services installation on the first node of the cluster, you can create one on it later by using the SBDUTIL utility (`/opt/novell/ncs/bin/sbdutil`). You might also need to delete and re-create an SBD partition if the SBD becomes corrupted or its device fails. See the man page for `sbdutil` for more information on how to use it.

If a cluster partition does not exist, create one by doing the following:

1 Ensure that the device you want to use has been initialized and marked as *Shareable for Clustering* as described in Section 9.14.1, "Prerequisites for Creating an SBD Partition," on page 137.

2 Enter `cluster down` at the terminal console of one cluster server.

This causes all cluster servers to leave the cluster.

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

`rcnovell-ncs stop`

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

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

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

For the -s option, replace *size* with the size (in MB) to use for the SBD partition. If the -s option is not used, the default size is 8 MB. The -s option is available in OES 2 SP1 and later. For information about size requirements, see "SBD Partitions" on page 70.

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

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

If the creation is not successful, the response is `No Shared Disk`. If the creation is successful, you can use the `sbdutil -f` command to view the path and name of the SBD partition, such as

`/dev/evms/.nodes/mycluster1.sbd`

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

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

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

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

If the SBD command is not successful, the response is `No Shared Disk`. If the SBD is created, there is no response. You can use the `sbdutil -f -s` command to view the path and name of the SBD partition.

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

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

5a In a Web browser, open iManager, then log in to the Novell eDirectory tree that contains the cluster you want to manage.

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

5b Select *Directory Administration*, then select *Modify Object*.

5c Browse to locate and select the Cluster object of the cluster you want to manage, then click *OK*.

5d Under *Valued Attributes*, select the *NCS: Shared Disk Flag*, then click *Edit*.
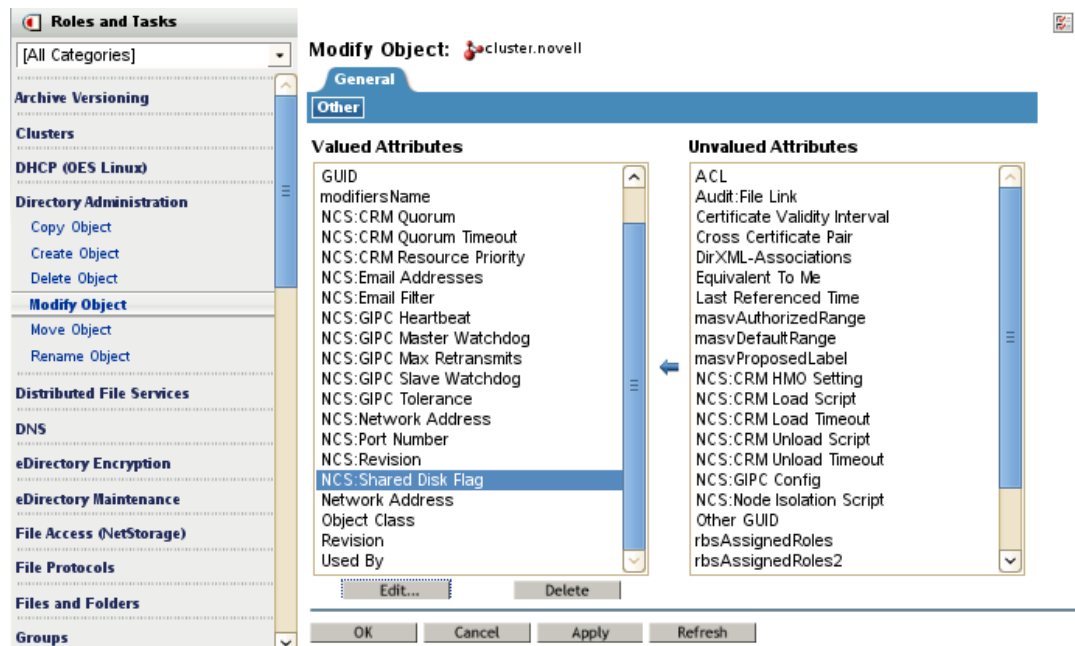
5e Select (enable) the *NCS: Shared Disk Flag* check box, then click *OK*.

5f Click *Apply* to save changes.

6 Restart Novell Cluster Services on all cluster nodes.

7 Rejoin each node to the cluster.

9.14.4 Creating a Mirrored Cluster SBD Partition

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

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

- "Using SBDUTIL to Create a Mirrored Cluster SBD Partition" on page 140
- "Using EVMSGUI to Create a Mirrored Cluster SBD Partition" on page 141

Using SBDUTIL to Create a Mirrored Cluster SBD Partition

1 Ensure that both of the devices you want to use have been initialized and marked as *Shareable for Clustering* as described in Section 9.14.1, "Prerequisites for Creating an SBD Partition," on page 137.

2 Enter `cluster down` at the terminal console of one cluster server.

This causes all cluster servers to leave the cluster.

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

`rcnovell-ncs stop`

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

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

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

For the -s option, replace *size* with the size (in MB) to use for the SBD partition. If the -s option is not used, the default size is 8 MB. The -s option is available in OES 2 SP1 and later. For information about size requirements, see "SBD Partitions" on page 70.

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

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

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

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

If the SBD command is not successful, the response is `No Shared Disk`. If the SBD is created, there is no response. You can use the `sbdutil -f -s` command to view the path and name of the SBD RAID device. You can use NSSMU or the Storage plug-in to iManager to view the partitions used by the RAID device.

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

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

5a In a Web browser, open iManager, then log in to the Novell eDirectory tree that contains the cluster you want to manage.

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

5b Click *Directory Administration*, then click *Modify Object*.

5c Browse to locate and select the Cluster object of the cluster you want to manage, then click *OK*.

5d Under *Valued Attributes*, select the *NCS: Shared Disk Flag*, then click *Edit*.

5e Select (enable) the *NCS: Shared Disk Flag* check box, then click *OK*.
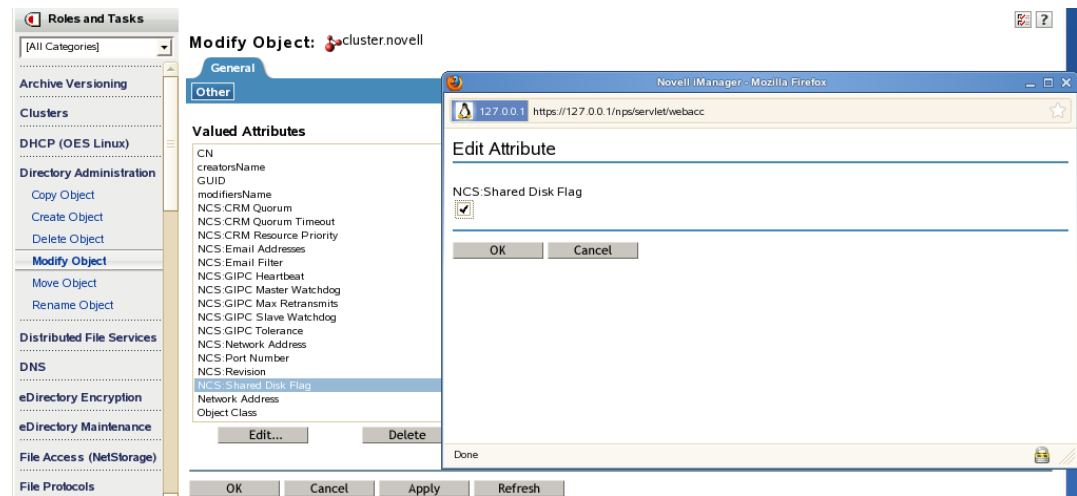
5f Click *Apply* to save changes.

6 Restart Novell Cluster Services on all cluster nodes.

7 Rejoin each node to the cluster.

Using EVMSGUI to Create a Mirrored Cluster SBD Partition

1 Ensure that both of the devices you want to use have been initialized and marked as *Shareable for Clustering* as described in Section 9.14.1, "Prerequisites for Creating an SBD Partition," on page 137.

2 At the Linux terminal console of a cluster server, log in as the `root` user, then enter `evmsgui` to start the EVMS GUI utility.

3 In evmsgui, create an SBD partition:

3a Click *Action*, then click *Create*.

3b Click *Segment*, choose the *NetWare Segment Manager*, then click *Next*.

3c Select *Free Space Storage Object*, then click *Next*.

3d Specify 20 MB or larger as the size of the cluster partition, then choose *SBD* as the partition type.

For information about size requirements, see "SBD Partitions" on page 70.

3e Specify the name of your cluster as the *Label*, then click *Create*.

3f Click *Save* to save your changes.

4 In evmsgui, mirror the newly created SBD partition:

4a Click *Segments*.

4b Locate the SBD partition and right-click it.

4c Select *Mirror Segment*, then click *OK*.

This creates a RAID 1 device with the name `/dev/evms/.nodes/cluster.sbd`.

4d Click *Save* to save your changes.

5 Exit evmsgui.

6 Reboot all cluster nodes.

9.14.5 Deleting a Non-Mirrored Cluster SBD Partition

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

1 At a Linux terminal console of a cluster server, log in as the `root` user.

2 Enter `cluster down` at the terminal console of one cluster server.

This causes all cluster servers to leave the cluster.

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

`rcnovell-ncs stop`

4 Delete the SBD partition.

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

4b Select the SBD partition you want to delete.

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

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

  ◆
  ◆

Ensure that the SBD partition exists before continuing.

6 Start Novell Cluster Services by entering

`rcnovell-ncs start`

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

9.14.6 Deleting a Mirrored Cluster SBD Partition

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

1 At a Linux terminal console of a cluster server, log in as the `root` user.

2 Enter `cluster down` at the terminal console of one cluster server.

This causes all cluster servers to leave the cluster.

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

`rcnovell-ncs stop`

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

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

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

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

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

    ◆ Section 9.14.3, "Creating a Non-Mirrored Cluster SBD Partition," on page 138

    ◆ Section 9.14.4, "Creating a Mirrored Cluster SBD Partition," on page 140

Ensure that the SBD partition exists before continuing.

6 Start Novell Cluster Services by entering

```
rcnovell-ncs start
```

7 Join the nodes to the cluster by entering `cluster join` at the terminal console on one of the nodes in the cluster:

9.14.7 Removing a Segment from a Mirrored Cluster SBD Partition

You can remove a segment from a mirrored cluster SBD partition and keep the remaining SBD partition. The software RAID definition remains, so if you delete the remaining partition later, you must delete the software RAID instead of simply deleting the partition as with a standalone SBD partition.

IMPORTANT: To get rid of the software RAID definition, you must delete the mirrored SBD partition as described in Section 9.14.6, "Deleting a Mirrored Cluster SBD Partition," on page 142, then re-create a non-mirrored SBD partition, as described in Section 9.14.3, "Creating a Non-Mirrored Cluster SBD Partition," on page 138.

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

1 At a Linux terminal console of a cluster server, log in as the `root` user.

2 At the terminal console of one cluster server, enter

```
cluster maintenance on
```

This causes all cluster servers to enter maintenance mode.

3 Enter `nssmu` to open the NSS Utility, then select *Software RAIDS*.

4 Select the software RAID1 device for the cluster SBD partition that you want to manage.

5 Press Enter to show its member segments.

6 Select the member segment you want to delete, then select *Delete* to remove the segment.

The RAID definition will still exist for the remaining segment of the mirrored SBD partition, but it reports that it is not mirrored.

7 At the terminal console of one cluster server, enter

```
cluster maintenance off
```

This causes all cluster servers to return to normal mode.

9.15 Customizing Cluster Services Management

Some portions of Novell Cluster Services management can be performed and customized by using virtual XML files that exist in the /admin/Novell/Cluster directory on Linux.

The cluster-related virtual XML files (management access points) are created on each server's /admin/Novell/Cluster directory. These files let you manage the cluster from any node in the cluster. This means that as long as the cluster is running, you can always access the cluster-related XML virtual files in the /admin/Novell/Cluster directory.

There are two types of virtual files in the /admin/Novell/Cluster directory: XML files and CMD files. The XML files are read-only and contain cluster configuration or cluster state information. The CMD files are write-then-read command files that are used to issue commands to the cluster and retrieve resulting status.

Table 9-3 lists the cluster-related virtual XML files and gives a brief description of each.

Table 9-3 *Cluster-Related Virtual XML Files*

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

Table 9-4 lists the cluster-related CMD files and gives a brief description of each.

Table 9-4 *Cluster-Related CMD Files*

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

10 Configuring and Managing Cluster Resources

After you create and configure a Novell Cluster Services cluster, you are ready to create and configure cluster resources for the cluster. This section provides general instructions for creating cluster resources and configuring their behavior in the cluster.

For information about viewing and managing resource status on the cluster, see Chapter 9, "Managing Clusters," on page 125.

- Section 10.1, "Planning for Cluster Resources," on page 147
- Section 10.2, "Using Cluster Resource Templates," on page 149
- Section 10.3, "Creating Cluster Resources," on page 152
- Section 10.4, "Configuring a Load Script for a Cluster Resource," on page 153
- Section 10.5, "Configuring an Unload Script for a Cluster Resource," on page 154
- Section 10.6, "Enabling Monitoring and Configuring the Monitor Script," on page 155
- Section 10.7, "Configuring the Start, Failover, and Failback Modes for Cluster Resources," on page 159
- Section 10.8, "Configuring Preferred Nodes for a Resource," on page 160
- Section 10.9, "Configuring Resource Priorities for Load Order," on page 161
- Section 10.10, "Configuring Resource Mutual Exclusion Groups," on page 162
- Section 10.11, "Controlling Resource Monitoring," on page 166
- Section 10.12, "Changing the IP Address of a Cluster Resource," on page 166
- Section 10.13, "Renaming a Cluster Resource," on page 166
- Section 10.14, "Deleting Cluster Resources," on page 167
- Section 10.15, "Additional Information for Creating Cluster Resources," on page 169

10.1 Planning for Cluster Resources

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

- Section 10.1.1, "Naming Conventions for Cluster Resources," on page 148
- Section 10.1.2, "Using Parameter Values with Spaces in a Cluster Script," on page 148
- Section 10.1.3, "Using Double Quotation Marks in a Cluster Script," on page 148
- Section 10.1.4, "Script Length Limits," on page 148
- Section 10.1.5, "Planning Cluster Maintenance," on page 149
- Section 10.1.6, "Number of Resources," on page 149

10.1.1 Naming Conventions for Cluster Resources

Cluster resource names can be up to 63 characters. Novell Cluster Services supports only alphanumeric characters and the underscore character in cluster resource names:

```
ABCDEFGHIJKLMNOPQRSTUVWXYZabcdefghijklmnopqrstuvwxyz0123456789_
```

Special characters, such as the following, are not supported in cluster resource names:

```
!@#$%&()
```

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

10.1.2 Using Parameter Values with Spaces in a Cluster Script

In cluster scripts, if a parameter value contains spaces, you should enclose the value with both double-quotes (") and single-quotes ('), such as "'quoted text'". The key is to use both sets of quotation marks with the single-quotes on the inside. For example:

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

You can alternatively use two sets of double-quotes and escape the inside set of double-quotes with a backslash (\). For example:

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

Another alternative is to escape each of the spaces in the value. For example:

```
exit_on_error echo Errors\ will\ be\ reported\ here.
```

10.1.3 Using Double Quotation Marks in a Cluster Script

In cluster scripts, if a command requires double-quotation marks, you should enclose the quoted part of the command with both double-quotes (") and single-quotes ('), such as "'quoted text'". The key is to use both sets of quotation marks with the single-quotes on the inside. This works for spaces and allows you to pass variables to the command. For example:

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

You can alternatively use two sets of double-quotes and escape the inside set of double-quotes with a backslash (\). For example:

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

10.1.4 Script Length Limits

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

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

10.1.5 Planning Cluster Maintenance

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

◆ Perform maintenance tasks during non-peak hours so that users are minimally affected.

◆ Before performing maintenance on a node, cluster migrate its cluster resources to another node if you want the related users to be undisturbed.

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

10.1.6 Number of Resources

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

10.2 Using Cluster Resource Templates

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

◆ Section 10.2.1, "Default Resource Templates," on page 149

◆ Section 10.2.2, "Creating a Resource Template," on page 150

◆ Section 10.2.3, "Synchronizing Locally Modified Resourse Templates with eDirectory," on page 152

10.2.1 Default Resource Templates

Table 10-1 identifies the cluster resource templates that Novell Cluster Services provides for use on physical servers and Xen virtual machine (VM) guest servers (DomU). You can also create your own templates or personalize the default templates by using iManager. For information, see Section 10.2.2, "Creating a Resource Template," on page 150. Third-party templates might also available for third-party applications; see the vendor documentation.

Table 10-1 *Cluster Resource Templates for Physical Servers and Xen VM Guest Servers*

Cluster Resource Template	OES 2 Linux Product
AV	Novell Archive and Version Services
DHCP	Novell Dynamic Host Configuration Protocol using an NSS pool
	Novell Dynamic Host Configuration Protocol using a Linux POSIX File System
DNS	Novell Domain Name System
Generic File System (Generic_FS)	Linux POSIX file systems

Cluster Resource Template	OES 2 Linux Product
Generic IP Service	This template can be modified to create cluster resources for certain server applications that run on your cluster.
iFolder	Novell iFolder 3.8*x*
iPrint	Novell iPrint
MySQL	Novell MySQL
Samba	Novell Samba
Third-party templates	See your vendor documentation.

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

Table 10-2   Cluster Resource Templates for Xen-Based Virtualization Host Environments

Cluster Resource Template	Use
Xen	Automatically configure the cluster resource for the virtual machine.
XenLive	Automatically configure the cluster resource for the virtual machine. Provides an additional function to allow a virtual machine resource migration (manual) without the need to boot or bring up the virtual machine on the cluster node where the virtual machine has been migrated.

To view the default templates in iManager:

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

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

2 Enter your Administrator user name and password.

3 In *Roles and Tasks*, click *Clusters*, then click *Cluster Options*.

4 Browse to locate and select the Cluster object of the cluster you want to manage.

A list of available resources and resource templates are displayed.

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

10.2.2   Creating a Resource Template

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

1 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 server in the cluster that has iManager installed.

2 Enter your Administrator user name and password.

3 In *Roles and Tasks*, click *Clusters*, then click *Cluster Options*.

4 Browse to locate and select the Cluster object of the cluster you want to manage.

5 Click the *New* link.

6 Specify *Template* as the resource type you want to create by clicking the *Template* radio button, then click *Next*.
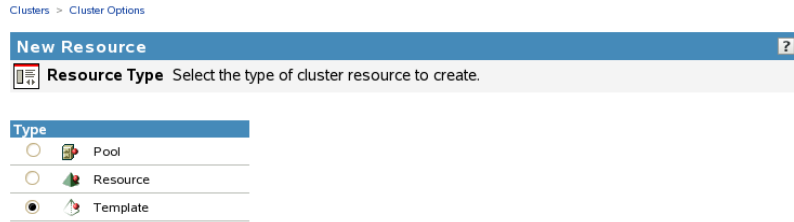
7 In *Cluster Resource Name*, specify the name of the template you want to create.

8 If desired, in *Inherit from Template*, browse to the Cluster object and select the existing resource template in the Cluster container that you want to personalize for the new template.

9 Ensure that the *Define Additional Properties* check box is selected, then click *Next* to continue to the Load Script page.

10 On the Load Script page, configure the load script for the cluster resource template.

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

   10b Edit or add any lines to the load script that are required to load dependent services such as Web servers or file access protocols.

   10c Edit or add the necessary commands to the script to load the resource on the server.

       For example, this might include bind command for the NCP service and the mount commands for the shared disks and file systems.

   10d Specify the default *Load Script Timeout* value, then click *Next* to continue to the Unload Script page.

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

11 On the Unload Script page, configure the unload script for the cluster resource template.

   11a Edit or add variables with example values for your template configuration, such as the mount point, IP address, container name, file system type, and device.

   11b Edit or add the necessary commands to the script to unload the resource from the server.

       For example, this might include unbind command for the NCP service and the dismount commands for the shared disks and file systems.

   11c Edit or add any lines to the unload script that are required to unload the dependent services that are loaded by this cluster resource.

11d Specify the default *Unload Script Timeout* value, then click *Next* to continue to the Monitor Script page.

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

12 On the Monitor Script page, configure the monitor script for the cluster resource template.

12a Edit or add the variables with example values for your template configuration, such as the mount point, IP address, container name, file system type, and device.

12b Edit or add the necessary commands to the script to monitor the resource on the server.

You can use the same commands that are used at the Linux terminal console.

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

12c Specify the default *Monitor Script Timeout* value, then click *Next*.

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

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

14 On the Resource Preferred Nodes page, specify the node assignments for the resource template, then click *Finish*.

The template you created is saved to the Cluster container of the cluster you selected. If you personalized an existing template, both the old template and the new template are in the container.

10.2.3   Synchronizing Locally Modified Resourse Templates with eDirectory

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

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

1 Log in to the master node of the cluster, then open a terminal console.

2 Enter the following commands:

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

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

10.3   Creating Cluster Resources

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

1 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 user name and password.

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

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

4 Browse to locate and select the Cluster object of the cluster you want to manage, then click the *New* link.

5 Specify Resource as the resource type you want to create by clicking the *Resource* radio button, then click *Next*.

6 Specify the name of the resource you want to create.

NOTE: Do not use periods in cluster resource names. Novell clients interpret periods as delimiters. If you use a space in a cluster resource name, that space is converted to an underscore.

---

7 In the *Inherit From Template* field, specify one of the available templates, such as the *Generic_FS_Template*.

For information about cluster resource templates, see Section 10.2, "Using Cluster Resource Templates," on page 149.

8 Select the *Define Additional Properties* check box, then click *Next*.

9 If you are creating a new cluster resource, continue with "Configuring a Load Script for a Cluster Resource" on page 153.

10.4 Configuring a Load Script for a Cluster Resource

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

Example load scripts are available in the following sections:

 ◆ Section 12.5, "Configuring a Load Script for the Shared NSS Pool," on page 189
 ◆ Section 13.7.1, "Sample Load Script for the Linux POSIX Volume Cluster Resource," on page 232

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

1 In iManager, click *Clusters,* then click *Cluster Options*.

2 Browse to locate and select the Cluster object of the cluster you want to manage.

3 Select the check box next to the resource whose load script you want to edit, then click the *Details* link.

4 Click the *Scripts* tab, then click the *Load Script* link.

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

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

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

IMPORTANT: Do not comment out commands that are automatically generated for parameters that define the cluster resource, such as the mount point, IP address, container name, file system type, and device. If you need to modify the IP address, administrator credentials, or other attributes of an existing resource, follow the procedure in Section 8.9, "Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster," on page 116.

6 Specify the *Load Script Timeout* value, then click *Apply* to save the script or, if you are creating a new cluster resource, click *Next*.

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

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

7 Do one of the following:

- ◆ If you are configuring a new resource, click *Next*, then continue with Section 10.5, "Configuring an Unload Script for a Cluster Resource," on page 154.
- ◆ Click *Apply* to save your changes.

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

10.5 Configuring an Unload Script for a Cluster Resource

Depending on your cluster application or resource, you can add an unload script to specify how the application or resource should terminate. An unload script is not required by all resources, but is required for cluster-enabled Linux partitions. Consult your application vendor or documentation to determine if you should add commands to unload the resource.

Example unload scripts are available in the following sections:

- ◆ Section 12.6, "Configuring an Unload Script for the Shared NSS Pool," on page 190
- ◆ Section 13.7.2, "Sample Unload Script for the Linux POSIX Volume Cluster Resource," on page 233

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

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

2 Browse to locate and select the Cluster object of the cluster you want to manage.

3 Select the check box next to the resource whose unload script you want to edit, then click the *Details* link.

4 Click the *Scripts* tab, then click the *Unload Script* link.

5 Edit or add the necessary commands to the script to unload the resource on the server.

You can add any lines to the unload script that are required to unload services that are loaded by this cluster resource.

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

6 Specify the *Unload Script Timeout* value, then click *Apply* to save the script or, if you are creating a new cluster resource, click *Next*.

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

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

7 Do one of the following:

   - If you are configuring a new resource, click *Next*, then continue with Section 10.6.2, "Configuring Resource Monitoring," on page 157.

   - Click *Apply* to save your changes.

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

10.6 Enabling Monitoring and Configuring the Monitor Script

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

   - Section 10.6.1, "Understanding Resource Monitoring," on page 155
   - Section 10.6.2, "Configuring Resource Monitoring," on page 157
   - Section 10.6.3, "Example Monitoring Scripts," on page 158
   - Section 10.6.4, "Monitoring Services That Are Critical to Clustering," on page 158

10.6.1 Understanding Resource Monitoring

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

   - "Polling Interval" on page 155
   - "Failure Rate" on page 156
   - "Failure Action" on page 156
   - "Timeout Value" on page 156
   - "How Resource Monitoring Works" on page 156

Polling Interval

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

Failure Rate

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

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

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

Failure Action

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

If the failure action initiates and you chose the option to migrate the resource to another server, the resource migrates to the next server in its *Assigned Nodes* list, which you previously ordered according to your preferences. The resource remains on the server it has migrated to unless you migrate it to another server or the failure action initiates again, in which case it again migrates to the next server in its *Assigned Nodes* list.

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

With resource monitoring, the *Start*, *Failover*, and *Failback* Modes have no effect on where the resource migrates. This means that a resource that has been migrated by the resource monitoring failure action does not migrate back (fail back) to the node it migrated from unless you manually migrate it back.

Timeout Value

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

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

How Resource Monitoring Works

1 The monitoring script runs at the frequency you specify as the polling interval.

2 There are two conditions that trigger a response by Novell Cluster Services:

   ◆ An error is returned. Go to Step 3.

   ◆ The script times out, and the process fails. Go to Step 4.

3 Novell Cluster Services tallies the error occurrence, compares it to the configured failure rate, then does one of the following:

  ◆ **Total errors in the interval are less than or equal to the Maximum Local Failures:** Novell Cluster Services tries to resolve the error by offlining the resource, then onlining the resource.

    If this problem resolution effort fails, Novell Cluster Services goes to Step 4 immediately regardless of the failure rate condition at that time.

  ◆ **Total errors in the interval are more than the Maximum Local Failures:** Go to Step 4.

4 Novell Cluster Services initiates the configured failure action. Possible actions are:

  ◆ Puts the resource in a comatose state

  ◆ Migrates the resource to another server

  ◆ Reboots the hosting node (without synchronizing or unmounting the disks)

10.6.2 Configuring Resource Monitoring

The resource monitoring function allows you to monitor the health of a specified resource by using a script that you create or customize. If you want Novell Cluster Services to check the health status of a resource, you must enable and configure resource monitoring for that resource. Enabling resource monitoring requires you to specify a polling interval, a failure rate, a failure action, and a timeout value.

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

1 In iManager, click *Clusters*, then click *Cluster Options*.

2 Browse to locate and select the Cluster object of the cluster you want to manage.

3 Select the check box next to the resource that you want to configure monitoring for, then click the *Details* link.

4 Click the *Monitoring* tab.

5 Select the *Enable Resource Monitoring* check box to enable resource monitoring for the selected resource.

Resource monitoring is disabled by default.

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

You can specify the value in minutes or seconds.

7 Specify the number of failures (*Maximum Local Failures*) for the specified amount of time (*Time Interval*).

For information, see "Failure Rate" on page 156.

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

For information, see "Failure Action" on page 156.

9 Click the *Scripts* tab, then click the *Monitor Script* link.

10 Edit or add the necessary commands to the script to monitor the resource on the server.

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

You can use the same commands that would be used at the Linux terminal console. For example, see .

11 Specify the *Monitor Script Timeout* value, then click *Apply* to save the script.

The timeout value determines how much time the script is given to complete. If the script does not complete within the specified time, the failure action you chose in initiates.

12 Do one of the following:

- ◆ If you are configuring a new resource, click *Next*, then continue with .
- ◆ Click *Apply* to save your changes.

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

10.6.3 Example Monitoring Scripts

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

Example monitor scripts are available in the following sections:

- ◆
- ◆

10.6.4 Monitoring Services That Are Critical to Clustering

Monitoring scripts can also be used for monitoring critical services needed by the resources, such as Linux User Management (`namcd`) and Novell eDirectory (`ndsd`). However, the monitoring is in effect only where the cluster resource is running.

IMPORTANT: The monitor script runs only on the cluster server where the cluster resource is currently online. The script does not monitor the critical services on its assigned cluster server when the resource is offline. The monitor script does not monitor critical services for any other cluster node.

For example, to monitor whether the `namcd` and `ndsd` services are running, add the following commands to the Monitor script:

```
# (optional) status of the eDirectory service
exit_on_error rcndsd status

# (optional) status of the Linux User Management service
exit_on_error rcnamcd status
```

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

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

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

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

- Section 10.7.1, "Understanding Cluster Resource Modes," on page 159
- Section 10.7.2, "Setting the Start, Failover, and Failback Modes for a Resource," on page 159

10.7.1   Understanding Cluster Resource Modes

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

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

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

The preferred node is the first server in the *Assigned Nodes* list for the resource.

IMPORTANT: Resources fail back only to the first node in their *Assigned Nodes* list. For example, if a resource has failed over to three servers since it originally ran on its preferred node, and the second server the resource was running on comes back up, the resource does not fail back to that second server.

Resources do not automatically move from node to node just because a node higher in the *Assigned Nodes* list rejoins the cluster, unless the *Failback* mode is set to AUTO and the first node in the *Assigned Nodes* list rejoins the cluster.

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

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

1. In iManager, click *Clusters* and then click *Cluster Options*.
2. Browse to locate and select the Cluster object of the cluster you want to manage.
3. Select the box next to the resource whose *Start*, *Failover*, or *Failback* modes you want to view or edit, then click the *Details* link.

4 Click the *Policies* tab.

5 (Optional) Select the *Resource Follows Master* check box if you want to ensure that the resource runs only on the master node in the cluster.

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

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

The quorum default values were set when you installed Novell Cluster Services. You can change the quorum default values by accessing the properties page for the Cluster object.

Selecting this box ensures that the resource is launched immediately on any server in the Assigned Nodes list as soon as any server in the list is brought online.

7 Specify the *Start*, *Failover*, and *Failback* modes for this resource.

The default for both *Start* and *Failover* modes is AUTO, and the default for *Failback* mode is DISABLE.

8 Do one of the following:

◆ If you are configuring a new resource, click *Next*, then continue with Section 10.8, "Configuring Preferred Nodes for a Resource," on page 160.

◆ Click *Apply* to save your changes.

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

10.8 Configuring Preferred Nodes for a Resource

The Preferred Nodes page allows you to control the order in which a resource attempts to fail over to one of its assigned nodes when the resource is brought online or during a failover. For example, if a node fails, the resource attempts to fail over to the one of the assigned nodes in its Preferred Nodes list. The preferred failover order determines which node is tried first. This is useful for ensuring that the resource runs on its most preferred node of the nodes available when it fails over or is brought online. Changes are not allowed for the Preferred Nodes list for the Master_IP_Address_Resource.

By default, a new resource is configured to run on all nodes, including future nodes. If its preferred node list is ever changed (directly or indirectly) by administrators, future nodes are not automatically added to the resource's preferred nodes list.

IMPORTANT: Ensure that you prepare the node for the services in the resource before you migrate or fail over the resource to it.

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

1 In iManager, click *Clusters*, then click *Cluster Options*.

2 Browse to locate and select the Cluster object of the cluster you want to manage.

3 Select the box next to the resource whose preferred node list you want to view or edit, then click the *Details* link.

4 Click the *Preferred Nodes* tab.

5 From the *Unassigned Nodes* list, select the server you want the resource assigned to, then click the right-arrow button to move the selected server to the *Assigned Nodes* list.

Repeat this step for all servers you want assigned to the resource.

6 From the *Assigned Nodes* list, select the servers you want to unassign from the resource, then click the left-arrow button to move the selected servers to the *Unassigned Nodes* list.

7 Use either of the following methods to change the preferred failover order of the nodes assigned to a resource:

   ◆ **Arrows:** Select one of the assigned nodes by clicking on it, then click the up-arrow and down-arrow buttons to move the node up or down in the list. The page refreshes between each click on an arrow.

   ◆ **Edit:** Open the Preferred Nodes Edit function by clicking the *Edit* pen icon, list the assigned nodes in the preferred failover order with one node per line, then click *OK* to accept the revised order.

     Be careful to not alter the names of the nodes while editing the list.
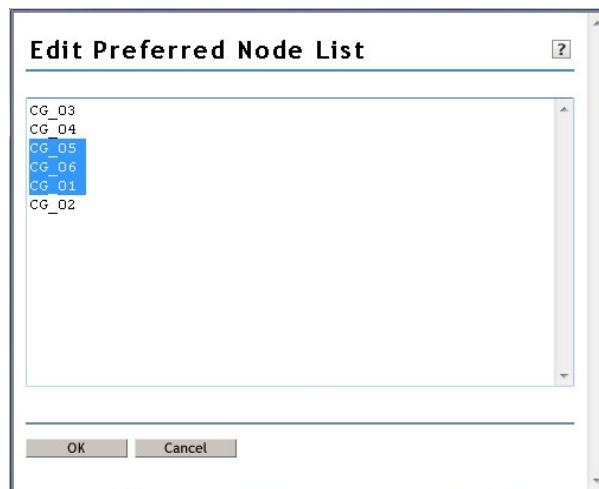
8 Click *Apply* to save node assignment changes.

9 Offline the resource, then online it again to apply the changes to the resource.

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

10.9 Configuring Resource Priorities for Load Order

Cluster resource priorities control the load order of a resource relative to other cluster resources on the same cluster node when bringing up a cluster, or during a failover or failback. This is useful for ensuring that the most critical resources load first and are available to users before less critical resources.

The Resource Priority setting controls the order in which multiple resources start on a given node when the cluster is brought up or during a failover or failback. For example, if a node fails and two resources fail over to another node, the resource priority determines which resource loads first.

1 In iManager, select *Clusters*, then select *Cluster Options*.

2 Browse to locate and select the Cluster object of the cluster you want to manage.

3 Click the *Properties* button under the cluster name.

4 Click the *Priorities* tab.

5 Use either of the following methods to change the load order (from highest priority to lowest priority) of a resource relative to other cluster resources on the same node:

- ◆ **Arrows:** Select a resource in the list by clicking on it, then click the up-arrow or down-arrow buttons to move the resource up or down in the list. The page refreshes between each click on an arrow.

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

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

```
Edit Cluster Resource Priorities              [?]

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

   OK          Cancel
```
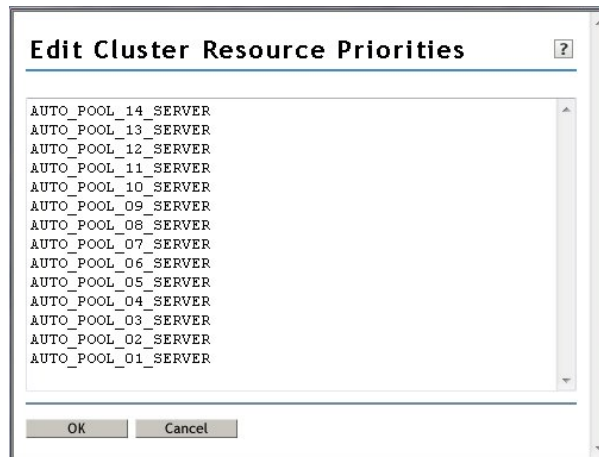
6 Click *Apply* or *OK* to save changes.

10.10 Configuring Resource Mutual Exclusion Groups

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

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

- ◆ Section 10.10.1, "Understanding Resource Mutual Exclusion," on page 163
- ◆ Section 10.10.2, "Setting Up RME Groups," on page 165
- ◆ Section 10.10.3, "Viewing RME Group Settings," on page 165

10.10.1 Understanding Resource Mutual Exclusion

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

How RME Groups Work

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

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

Rules for Node Assignment

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

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

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

Resources that are not members of any of the RME Groups are not restricted. They are allowed to run at any time and in any combination on any node in their Preferred Nodes lists.

Examples

If you try to cluster migrate a resource to a node where another member of its RME group is running, the resource is not migrated and remains in the Running state on the node where it is currently loaded. Only one resource in the group is allowed to run on the node at a time.

If you try to cluster migrate multiple resources in an RME Group to the same node, the resources are brought online on different nodes. Only one of the resources (randomly selected) is brought online on the specified node. The other resources are migrated to different nodes in their Preferred Nodes lists.

Planning for RME Groups

An RME Group can contain any combination of the resources that are available to the cluster. Resources that are members of the same group are not allowed to run concurrently on a node. A resource can be a member of more than one group.

You can define up to four groups (Group A, Group B, Group C, and Group D). The group names are fixed; they cannot be customized.

For example, the following image shows an RME Groups page where three groups (A, B, and D) have been defined. The selected check boxes in a column indicate which resources belong to that group. Group A has four member resources, and Group B and Group D have three members each. The AUTO_POOL_12_SERVER resource belongs to Group A and Group B. It has mutually exclusive relationships with two different sets of resources, and each exclusion is managed separately.

Figure 10-1   *Sample RME Groups for an OES 2 SP3 Cluster*

RME Groups are supported when all cluster nodes are running OES 2 SP3 or later. If you are managing an older version of Novell Cluster Services (without the RME code) from a new Clusters plug-in in iManager, the tables on the RME Groups page are shown as *No items*. You cannot set up RME Groups for clusters running on older platforms.

Figure 10-2   *Sample RME Groups Page for an Unsupported Platform*

If you are upgrading from an OES 2 SP2 Linux or earlier version of Novell Cluster Services where RMS Groups are not available, you should wait until the rolling cluster upgrade to OES 2 SP3 Linux or later has been completed on all nodes before you define RME Groups. If you define RME Groups while the cluster is in a mixed-node condition, the RME Groups are honored only if the resource fails over to a cluster node where OES 2 SP3 Linux or later is running.

10.10.2   Setting Up RME Groups

To define sets of resources that must not be assigned to the same node at the same time:

1 In iManager, click *Clusters*, then click *Cluster Options*.

2 Browse to locate and select the Cluster object of the cluster you want to manage.

3 Click *Properties*, then click the *RME Groups* tab.

The RME Groups page lists the resources in the cluster and columns for four possible RME Groups for the cluster. In each Group's column, the check boxes correspond to the resources listed at the left.

If you are managing an older version of Novell Cluster Services (without the RME code) from a new Clusters plug-in in iManager, the tables are shown as *No Items*. You cannot set up RME Groups for clusters running on older platforms.

4 To set up the members in a group, select the check boxes in the same column for two or more resources that must not be assigned on the same node at the same time.

The selected resources in a column are not allowed to run concurrently on a node. A resource can be a member of more than one group. You can define up to four groups (A, B, C, and D).

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

10.10.3   Viewing RME Group Settings

The RME Groups settings are displayed in the following locations:

Cluster Report

The Cluster Report has an RME Groups section that lists the member resources of each group.

1 In iManager, click *Clusters*, then click *Cluster Manager*.

2 Browse to locate and select the Cluster object of the cluster you want to manage.

3 Click *Run Report*, then scroll down to view the *RME Groups*.

The RME Groups section lists only the resources in the cluster that are members in each of the RME Groups.

For example, the following report shows the members for Groups A, B, and D. Group C has no defined members. The AUTO_POOL_12_SERVER resource belongs to two RME Groups. It has mutually exclusive relationships with two different groups of resources. These exclusions are managed separately.

Group A	Group B	Group C	Group D
AUTO_POOL_12_SERVER	AUTO_POOL_14_SERVER		AUTO_POOL_10_SERVER
AUTO_POOL_09_SERVER	AUTO_POOL_12_SERVER		AUTO_POOL_07_SERVER
AUTO_POOL_06_SERVER	AUTO_POOL_05_SERVER		Master_IP_Address_Resource
AUTO_POOL_02_SERVER			

Resource Mutual Exclusion (RME) Groups

RME Groups Configuration Page

The RME Groups configuration page displays all possible resources, and indicates membership by the check boxes that are selected in the same column under the RME Group Name.

1 In iManager, click *Clusters*, then click *Cluster Options*.

2 Browse to locate and select the Cluster object of the cluster you want to manage.

3 Click *Properties*, then click the *RME Groups* tab.

The RME Groups configuration page lists the all of the resources in the cluster. The selected check boxes in the same group column indicate their memberships in the four possible RME Groups.

10.11 Controlling Resource Monitoring

You can use the `cluster monitor` command to start, stop, or view the status of monitoring for a specified cluster resource. Issue the command as the `root` user from the node where the resource is currently loaded:

```
cluster monitor <resourcename> {start | stop | status}
```

Monitoring must be enabled for the resource. For instructions, see Section 10.6, "Enabling Monitoring and Configuring the Monitor Script," on page 155.

10.12 Changing the IP Address of a Cluster Resource

You can change the IP address of a cluster resource by using the Cluster Properties page of the Clusters plug-in in iManager. Ensure that you offline the cluster resource before attempting to change its IP address. When you online the cluster resource, the load script is updated with the new IP address.

10.13 Renaming a Cluster Resource

Beginning in OES 2 SP3, a `cluster rename` option is available that allows you to rename a cluster resource. Renaming the resource does not modify the virtual server name (the NCS:NCP Server object).

Your setup must meet the following prerequisites:

- This command must be issued from the master node.
- The resource must be in the *Offline* state in order to be renamed.
- The new name must not exist prior to the renaming.
- Novell eDirectory must be running when you attempt to rename cluster resource.

To rename the cluster resource:

1 Log in to the master node in the cluster as the `root` user, and open a terminal console.

2 Offline the cluster resource that you want to rename by entering

```
cluster offline resource_name
```

For example:

```
cluster offline POOL1_SERVER
```

3 Rename the cluster resource by entering

```
cluster rename <resource_name> <new_resource_name>
```

For example:

```
cluster rename POOL1_SERVER customized_name22
```

4 Online the cluster resource by entering

```
cluster online new_resource_name
```

For example:

```
cluster online customized_name22
```

10.14 Deleting Cluster Resources

Ensure that you offline the cluster resource before attempting to delete either the cluster resource or the clustered pool.

WARNING: If you attempt to delete a cluster resource without first offlining it, deletion errors occur, and the data associated with the clustered pool is not recoverable.

To delete a resource and create a new one with the same name, you must wait to create the new one until eDirectory synchronizes all of the objects in the tree related to the deleted resource.

10.14.1 Deleting a Cluster Resource on a Master Node

We strongly recommend that when you need to delete a cluster resource, that you do so only from the master node in the cluster. If the resource cannot be migrated to the master node, follow the procedure in Section 10.14.2, "Deleting a Cluster Resource on a Non-Master Node," on page 168.

You might want to delete the shared storage area if you no longer need the data.

WARNING: Deleting a pool or Linux POSIX volume destroys all data on it.

1 If the resource is on a non-master node in the cluster, migrate it to the master node.

2 If the cluster resource is online, offline it before continuing by using one of the following methods:

- Enter the following at the terminal console prompt as the root user:

```
cluster offline resource
```

- In iManager, go to *Clusters > Cluster Manager*, specify the cluster you want to manage, select the cluster resource, then click *Offline*.

3 Delete the resource on the master node by using the appropriate storage management tool:

- For shared NSS pools and volumes, use NSSMU or the Storage plug-in to iManager. Deleting the pool automatically deletes the volumes on it.
- For shared Linux POSIX volumes, use evmsgui.

4 In eDirectory, look at the Cluster Resource objects in the Cluster container to verify that the resource has been deleted from the Cluster container.

   If necessary, you can delete the Cluster Resource object and its virtual server object manually. In iManager, go to *Directory Administration > Delete Objects*, select the objects, then click *OK*.

10.14.2 Deleting a Cluster Resource on a Non-Master Node

We strongly recommend that when you need to delete a cluster resource, that you do so only from the master node in the cluster. If the resource can be migrated, migrate it to the master node and follow the procedure in Section 10.14.1, "Deleting a Cluster Resource on a Master Node," on page 167.

You might want to delete a cluster resource on a non-master node when deleting NSS pool and volume resources (by using NSSMU).

If you must delete a cluster resource while it resides a non-master node, use the following procedure:

1 Log in as the `root` user to the non-master node where the cluster resource currently resides, then open a terminal console.

2 If the cluster resource is online, offline it by entering

   `cluster offline resource`

3 At the terminal console prompt on the non-master node, enter

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

4 Look at the file `/var/opt/novell/ncs/resource-priority.conf` to verify that it has the same information (REVISION and NUMRESOURCES) as the file on the master node.

5 Delete the resource on the non-master node by using the appropriate storage management tool:

   ◆ For shared NSS pools and volumes, use NSSMU or the Storage plug-in to iManager.

   ◆ For shared Linux POSIX volumes, use evmsgui.

6 In eDirectory, look at the objects in the Cluster container to verify that the resource has been deleted from the Cluster container.

7 On the master node, log in as the `root` user, then open a terminal console.

8 At the terminal console prompt on the master node, enter

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

9 Look at the file `/var/opt/novell/ncs/resource-priority.conf` to verify that it has the same information (REVISION and NUMRESOURCES) as that of the non-master node where you deleted the cluster resource.

10 In iManager, select *Clusters > Cluster Options*, then browse to select the Cluster object.

11 Click *Properties*, select the *Priorities* tab, then click *Apply* on the Priorities page.

12 At the terminal console, enter

   `cluster view`

   The cluster view should be consistent.

13 Look at the file `/var/opt/novell/ncs/resource-priority.conf` on the master node to verify that the revision number increased.

   If the revision number increased, you are done. Do not continue with Step 14.

If the deleted resource is the only one in the cluster, the priority won't force the update. A phantom resource might appear in the interface. You need to restart Cluster Services to force the update, which also removes the phantom resource.

14 If the revision number did not automatically update in the previous steps, restart Novell Cluster Services by entering the following on one node in the cluster:

```
cluster restart [seconds]
```

For *seconds*, specify a value of 60 seconds or more. The cluster leave process begins immediately. Specify a value of 60 or more seconds as the time to wait before the cluster join begins. For example:

```
cluster restart 120
```

10.15 Additional Information for Creating Cluster Resources

Cluster resources can be configured for storage, services, or virtual machines.

- Section 10.15.1, "Creating Storage Cluster Resources," on page 169
- Section 10.15.2, "Creating Service Cluster Resources," on page 169
- Section 10.15.3, "Creating Virtual Machine Cluster Resources," on page 170

10.15.1 Creating Storage Cluster Resources

For information about creating cluster resources for shared storage on Linux, see the following:

Table 10-3 *Cluster-Enabling Shared Storage*

Shared Storage	Refer to
NSS pools and volumes	Chapter 12, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 173
Linux POSIX volumes	Chapter 13, "Configuring and Managing Cluster Resources for Shared Linux POSIX Volumes," on page 213
NCP volumes	"Configuring NCP Volumes with Novell Cluster Services" in the *OES 2 SP3: NCP Server for Linux Administration Guide*
Dynamic Storage Technology shadow volume pairs (shared NSS pools and volumes configured as DST pairs)	"Configuring DST Shadow Volume Pairs with Novell Cluster Services" in the *OES 2 SP3: Dynamic Storage Technology Administration Guide*

10.15.2 Creating Service Cluster Resources

For information about creating cluster resources for various services on your OES 2 Linux server, see Chapter 11, "Creating and Managing Service Cluster Resources," on page 171.

10.15.3 Creating Virtual Machine Cluster Resources

If you install Novell Cluster Services at the host level of an OES 2 Linux (Xen) server, you can create cluster resources for the virtual machines. For information, see Section 14.2, "Virtual Machines as Cluster Resources," on page 238.

11 Creating and Managing Service Cluster Resources

The individual service guides contain information about creating and managing service cluster resources with Novell Cluster Services on your Novell Open Enterprise Server (OES) 2 Linux server. See Table 11-1 for links to those resources. You can also find this list by going to "Clustering Linux Services" on the Clustering (High Availability) Documentation Web site (http://www.novell.com/documentation/oes2/cluster-services.html#clust-config-resources).

Table 11-1   *Clustering OES 2 Linux Services with Novell Cluster Services*

OES 2 Linux Service	For information, see
Apache Web Server	"Apache HTTP Server" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*
Archive and Version Services	"Configuring Archive and Version Service for Novell Cluster Services" in the *OES 2 SP3: Novell Archive and Version Services 2.1 Administration Guide for Linux*
Certificate Server	The eDirectory Certificate Server is not cluster-enabled. The Certificate Server service issues Server Certificate objects that might need to reside on each node in a cluster, depending on the service that is clustered. See "eDirectory Server Certificates" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*.
DFS VLDB (Distributed File Services volume location database)	"Clustering Novell Distributed File Services" in the *OES 2 SP3: Novell Distributed File Services Administration Guide for Linux*
DHCP Server	In the *OES 2 SP3: Novell DNS/DHCP Administration Guide*, see: ◆ "Configuring DHCP with Novell Cluster Services for the NSS File System" ◆ "Configuring DHCP with Novell Cluster Services for the Linux File System"
DNS Server	"Configuring DNS with Novell Cluster Services" in the *OES 2 SP3: Novell DNS/DHCP Administration Guide*
eDirectory	Novell eDirectory is not clustered because it has its own replica system.
File, AFP (Apple Filing Protocol)	"Configuring AFP with Novell Cluster Services for an NSS File System" in the *OES 2 SP3: Novell AFP For Linux Administration Guide*
File, CIFS (Windows File Services)	"Configuring CIFS with Novell Cluster Services for an NSS File System" in the *OES 2 SP3: Novell CIFS for Linux Administration Guide*

OES 2 Linux Service	For information, see
File, FTP	"Cluster Enabling Pure-FTPd in an OES 2 Environment" in the *OES 2 SP3: Planning and Implementation Guide*.
File, NetStorage	"Configuring NetStorage with Novell Cluster Services" in the *OES 2 SP3: NetStorage Administration Guide*.
File, Samba	"Configuring Samba for Novell Cluster Services" in the *OES2 SP3: Samba Administration Guide*
iFolder 3.8.4	"Clustering iFolder Servers with Novell Cluster Services for Linux" (http://www.novell.com/documentation/ifolder3/ifolder38_admin/data/bx6m7nn.html)in the *Novell iFolder 3.8 Administration Guide*.
MySQL	A MySQL template is available that uses a shared Linux Posix file system that you have already created.
iPrint	"Configuring iPrint with Novell Cluster Services" in the *OES 2 SP3: iPrint for Linux Administration Guide*
QuickFinder (Server Synchronization Feature)	"Configuring QuickFinder Server for Novell Cluster Services" in the *OES 2 SP3: Novell QuickFinder Server 5.0 Administration Guide*
Storage, DST shadow volumes built with NSS volumes	In the *OES 2 SP3: Dynamic Storage Technology Administration Guide*, see: ◆ "Configuring DST Shadow Volume Pairs with Novell Cluster Services" ◆ "Technology Preview: Using DST Shadow Volumes with Remote Secondary NSS Volumes and Novell Cluster Services"
Storage, Linux POSIX volumes	Chapter 13, "Configuring and Managing Cluster Resources for Shared Linux POSIX Volumes," on page 213
Storage, NCP volumes	"Configuring NCP Volumes with Novell Cluster Services" in the *OES 2 SP3: NCP Server for Linux Administration Guide*
Storage, NSS pools and volumes	Chapter 12, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 173
Xen virtual machines	Section 14.2, "Virtual Machines as Cluster Resources," on page 238

12 Configuring Cluster Resources for Shared NSS Pools and Volumes

After you have installed and configured Novell Cluster Services on a server, you can configure pool cluster resources. This section describes how to create and cluster-enable shared Novell Storage Services pools volumes as pool cluster resources with Novell Cluster Services.

12.1 Requirements for Creating Pool Cluster Resources

Your system must meet the requirements in this section in addition to the cluster requirements described in Chapter 4, "Planning for Novell Cluster Services," on page 55.

12.1.1 Novell Cluster Services

Novell Cluster Services must be installed and running when you create and manage the shared NSS pools and volumes. The cluster must be active.

12.1.2 Resource IP Address

Each cluster-enabled NSS pool requires a unique static IP address. The IP address is used to provide access and failover capability to the cluster-enabled pool (virtual server). Users access the pool by using the resource IP address instead of the server IP address where the pool is active. The IP address you assign to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.

IMPORTANT: The IP address for the virtual server must be in the same IP subnet as the server nodes in the cluster where you plan to use it.

12.1.3 Novell eDirectory

Novell eDirectory must be running and working properly when you cluster-enable a pool and when you manage the pool.

12.1.4 Novell Storage Services

NSS must be installed and running on each server in the cluster. For information about installing NSS and managing NSS pools and volumes, see the *OES 2 SP3: NSS File System Administration Guide for Linux*.

In addition, the following requirements must be met:

  ◆ "Pool" on page 174
  ◆ "Volume" on page 174
  ◆ "Naming Conventions" on page 175
  ◆ "Pool Cluster Resource Name" on page 175
  ◆ "Location of Cluster, Server, and Storage Objects" on page 175

Pool

We recommend that your pool cluster resource be one device, one pool, one volume. You can create and cluster-enable a shared NSS pool by using the Storage plug-in for Novell iManager or the server-based NSS Management Utility (nssmu). You can also use these tools to create NSS volumes on the shared pool.

Volume

You must create at least one volume in the shared pool. Typically, you create all volumes for a shared pool when you set up the pool cluster resource and before you need to cluster migrate or fail over the resource to a different node in the cluster.

Naming Conventions

Pool names and volume names must comply with the same naming conventions. Names must be 2 to 15 characters in length. Novell Cluster Services supports pool names with characters A to Z, 0 to 9, and underscores (_). The name cannot begin or end with an underscore; it cannot contain multiple adjacent underscores (__). It does not allow the following special characters in a shared pool name:

```
!@#$%&()
```

IMPORTANT: Before you cluster-enable an existing NSS pool, you must rename the pool and its volumes to use names that do not contain special characters.

Pool Cluster Resource Name

The pool name is used to form the default name of the pool cluster resource and the virtual server (the NCS:NCP Server object). If you modify the resource name, ensure that the name conforms to the resource naming conventions as described in Section 10.1.1, "Naming Conventions for Cluster Resources," on page 148.

Storage Object	Default Name
Resource name	`<servername>_<poolname>`
Clustered Volume object name	`<cluster_name>_<volume_name>`
	`\\<cluster_name>-<poolname>-SERVER\<volume_name>`
Virtual server name	`<servername>_<poolname>_SERVER`

Location of Cluster, Server, and Storage Objects

The Server, Pool, Volume, Cluster Resource, and Cluster objects are recommended to be in the same context (such as ou=ncs,o=novell). It is supported for the objects to be in the same context or different contexts. If the objects are in different contexts, you might need to cluster migrate the pool cluster resource back to the node where the pool was created in order to modify the pool or volume, or to perform other tasks like setting up Distributed File Services junctions or home directories. You receive an eDirectory error if the operation cannot find the information that it needs in the same context.

12.1.5 Shared Storage

You should carefully plan how you want to configure your shared storage prior to installing Novell Cluster Services. Consider the guidelines and requirements in the following sections when planning your NSS storage solution.

- Section 4.7, "Shared Disk Configuration Requirements," on page 70
- Section 4.8, "SAN Rules for LUN Masking," on page 72
- Section 4.9, "Multipath I/O Configuration Requirements," on page 73

12.1.6 NCP Server for Linux

NetWare Core Protocol (NCP) is the Novell networking protocol used by the Novell Client. NCP is automatically selected as an advertising protocol when you cluster-enable an NSS pool. This is necessary to provide authenticated access to data using the Novell Trustee model.

Novell Storage Services requires that the NCP Server for Linux service be installed and running on each node in the server. NCP Server must be running even if users access volumes on the shared NSS pool only via other protocols.

WARNING: Cross-protocol file locking is required when using multiple protocols for data access on the same volume. This helps prevent possible data corruption that might occur from cross-protocol access to files.

Beginning in OES 2 SP2, the NCP Cross-Protocol File Lock parameter is enabled by default when you install NCP Server. If you modify the Cross-Protocol File Lock parameter, you must modify the setting on all nodes in the cluster.

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

For information, see "Configuring Cross-Protocol File Locks for NCP Server" in the *OES 2 SP3: NCP Server for Linux Administration Guide*.

NCP Server for Linux is installed by selecting *NCP Server and Dynamic Storage Technology* from the OES Services menu in the YaST install interface. For information about NCP Server for Linux, see the *OES 2 SP3: NCP Server for Linux Administration Guide*.

12.1.7 Novell CIFS for Linux

Common Internet File System (CIFS) is the Windows networking protocol. The Novell CIFS for Linux service is available beginning in the OES 2 SP1 Linux release. Novell CIFS allows you to give clients access via CIFS to volumes on the shared NSS pool.

WARNING: To prevent possible data corruption, enable the NCP Cross-Protocol File Locks parameter for NCP Server on all nodes in the cluster before you allow users to access the data. For information, see "Configuring Cross-Protocol File Locks for NCP Server" in the *OES 2 SP3: NCP Server for Linux Administration Guide*.

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

Novell CIFS must be installed, configured, and working properly before you can specify CIFS as an advertising protocol when cluster-enabling an NSS pool. Novell CIFS for Linux is installed by selecting *Novell CIFS* from the OES Services menu in the YaST install interface. For information about Novell CIFS for Linux, see the *OES 2 SP3: Novell CIFS for Linux Administration Guide*.

12.1.8 Novell AFP for Linux

Apple Filing Protocol (AFP) is the Macintosh networking protocol. The Novell AFP for Linux service is available beginning in the OES 2 SP1 Linux release. Novell AFP is required when you want to give Macintosh clients access via AFP to volumes on the shared NSS pool.

WARNING: To prevent possible data corruption, enable the NCP Cross-Protocol File Locks parameter for NCP Server on all nodes in the cluster before you allow users to access the data. For information, see "Configuring Cross-Protocol File Locks for NCP Server" in the *OES 2 SP3: NCP Server for Linux Administration Guide*.

Novell AFP must be installed, configured, and working properly before you can specify AFP as an advertising protocol when cluster-enabling an NSS pool. Novell AFP for Linux is installed by selecting *Novell AFP* from the OES Services menu in the YaST install interface. For information about Novell AFP for Linux, see the *OES 2 SP3: Novell AFP For Linux Administration Guide*.

12.1.9 Novell Samba

Novell Samba is supported for NSS pool cluster resources as an alternative to Novell CIFS. The setup is not integrated in the Advertising Protocols options when you cluster-enable a pool. If you do not use Novell CIFS, Novell Samba can be set up after you cluster-enable the pool. For information about setting up Novell Samba for storage resources, see "Configuring Samba for Novell Cluster Services" in the *OES2 SP3: Samba Administration Guide*.

Novell Samba requires that users are enabled for Linux with Linux User Management (LUM). For information, see *OES 2 SP3: Novell Linux User Management Administration Guide*.

WARNING: To prevent possible data corruption, enable the NCP Cross-Protocol File Locks parameter for NCP Server on all nodes in the cluster before you allow users to access the data. For information, see "Configuring Cross-Protocol File Locks for NCP Server" in the *OES 2 SP3: NCP Server for Linux Administration Guide*.

12.1.10 Domain Services for Windows

Cluster-enabled NSS volumes can be used in a Domain Services for Windows environment.

12.2 Guidelines for Using Pool Cluster Resources

Consider the guidelines in this section when working with shared NSS pools and volumes in the cluster:

12.2.1 Guidelines for Shared Pools

- When the pool cluster resource is brought online, the pool is automatically activated by the resource load script. You typically do not activate the pool at the terminal console.
- The unit of failover is the device (disk or LUN) that contains the pool. If you use multiple devices in the pool, all of those devices must be failed over together.

  IMPORTANT: As a best practice, we recommend a single LUN (or device) per shared pool, and one volume per pool. (one LUN=>one pool=>one volume)

- If you delete a pool cluster resource, Novell Cluster Services automatically removes the Pool Resource object and the virtual server object from Novell eDirectory. After you delete the resource, you can unshare the device and use the pool locally, or you can delete the pool and its volumes.

  Ensure that you offline the cluster resource before you attempt to delete the cluster resource or the pool. For information, see Section 10.14, "Deleting Cluster Resources," on page 167.

- Before you rename a cluster-enabled pool, ensure that you offline the pool resource, activate the pool by using iManager or NSSMU, rename the pool, deactivate the pool, then online the pool resource.

  Novell Cluster Services automatically updates the pool resource load and unload scripts to reflect the name change. Also, NSS automatically changes the Pool Resource object name in eDirectory.

12.2.2 Guidelines for Shared Volumes

You must create at least one volume on a shared pool before you migrate it. As a best practice, we recommend that you create only one volume per shared pool.

When you create a volume, commands are added to the pool resource load and unload scripts to automatically mount and dismount the volume when the scripts run. You can modify the load script to comment out the mount command so that you can manually mount the volume on a node in the cluster where the pool resource has been activated.

Volumes on cluster-enabled pools do not appear as independent cluster resources. A cluster-enabled volume does not have an associated load and unload script or an assigned IP address. If you want each volume to be in a separate cluster resource, each volume must have its own pool.

When a server fails, the cluster resources fail over to other servers in the cluster. Because the cluster-enabled pool fails over, all volumes in the pool also fail over, but only the volumes that have been added to the load script are automatically mounted. Any volumes in the pool that are not in the load script must be mounted manually. For this reason, volumes that you do not want to fail over should be in separate pools that are not cluster-enabled.

When you create an encrypted NSS volume in a shared pool, you must mount the volume manually by using NSSMU and enter the password. NSS uses the password to create a key. Instead of storing it in the server memory as it does for non-shared volumes, NSS asks Novell Cluster Services to store the key and to pass it to the other nodes. After all servers hold the key, the volume is available for access as long as any one of the servers is still participating actively in the cluster. If all of the servers in the cluster fail, you must repeat this manual mounting procedure when you recover the cluster and restart services.

If you delete a volume from a cluster-enabled pool, Novell Cluster Services automatically removes the volume mount command from the pool cluster resource load script.

12.3 Creating Cluster-Enabled Pools and Volumes

After you have installed and configured Novell Cluster Services, you can create a pool cluster resource by creating a cluster-enabled NSS pool and volume.

- Section 12.3.1, "Initializing and Sharing a Device," on page 179
- Section 12.3.2, "Creating a Cluster-Enabled Pool and Volume with iManager," on page 180
- Section 12.3.3, "Creating a Cluster-Enabled Pool and Volume with NSSMU," on page 182

12.3.1 Initializing and Sharing a Device

Devices (disks or LUNs) that you want to use for your cluster-enabled pool must be initialized and shared to be visible to the NSS management tools when you create a pool. Novell Cluster Services must be installed and running on the server in order to share the device.

You can perform the device management tasks by using the Novell Storage Services (NSS) Management Utility or the Storage plug-in to iManager.

- ◆ "Initializing and Sharing a Disk with iManager" on page 179
- ◆ "Initializing and Sharing a Disk with NSSMU" on page 179

Initializing and Sharing a Disk with iManager

1 Ensure that the SAN device is attached to all nodes in the cluster.

2 Log in to iManager as an administrator user, then select *Storage > Devices*.

3 Browse to select the Cluster object ( ) of the cluster.

Selecting the Cluster object automatically selects the server that is currently the master node in the cluster.

4 From the *Devices* list, select the device where you want to create the pool.

5 If the device has not been previously initialized or if you want to delete the current partitioning structures on the device, click *Initialize Disk.*

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

Wait for the page to refresh before continuing.

6 In the *Details* area, select the *Shareable for Clustering* check box, then click *OK* to apply the change.

7 Repeat Step 4 through Step 6 for each device that you want to use in the pool.

8 Exit iManager.

9 Continue with Section 12.3.2, "Creating a Cluster-Enabled Pool and Volume with iManager," on page 180.

Initializing and Sharing a Disk with NSSMU

1 Ensure that the SAN device is attached to all nodes in the cluster.

2 Log in as the root user to the master node in the cluster, then open a terminal console.

3 At the console prompt, enter

```
nssmu
```

4 In the NSSMU, select *Devices* and press Enter.

5 In the *Devices* list, select the device where you want to create the pool.

6 If the device has not been initialized or if you want to delete the current partitioning structures on the device, press F3 to initialize the selected device.

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

Wait for the page to refresh before continuing.

7 Press F6 to mark the device as shareable for clustering.

The *Shareable for Clustering* value changes from *No* to *Yes*.

8 Repeat Step 5 through Step 7 for each device that you want to use in the pool.

9 Exit NSSMU.

10 Continue with Section 12.3.3, "Creating a Cluster-Enabled Pool and Volume with NSSMU," on page 182.

12.3.2 Creating a Cluster-Enabled Pool and Volume with iManager

1 Ensure that the SAN device is attached to all of the nodes in the cluster.

2 Log in to iManager as an administrator user.

3 Create a clustered pool:

3a In *Roles and Tasks*, select *Storage > Pools*.

3b Browse to select the Cluster object ( ) of the cluster.

This automatically selects the server that is currently the master node in the cluster. The shared device will be assigned to this server and the pool will be created there. After you create the clustered pool and volume, you can modify the resource's preferred node order, then cluster migrate the volume to its most preferred node.

3c Click the *New* link to open the New Pool wizard.

3d Specify the new pool name, then click *Next*.

3e Select the check box next to the shared device where you want to create the pool, then specify the size of the pool.

3f Select or deselect the *Mount On Creation* option.

The Mount On Creation option determines if the pool you are creating is to be activated (the resource is brought online) as soon as it is created.

The option is enabled by default. This allows you to create volumes immediately after you create the clustered pool. Default settings apply to the pool cluster resource and scripts, but you can modify them later.

If you deselect the option, you can configure the resource and modify the scripts before you bring the resource online for the first time. You must bring the pool cluster resource online before you can create volumes on it.

3g Select the *Cluster Enable on Creation* check box.

This option is selected by default if the device is shared.

If you deselect this option, you are not shown the Cluster Pool Information page. After you create a volume on the non-clustered pool, you can cluster-enable the pool and volume as described in Section 12.4, "Cluster-Enabling an Existing NSS Pool and Its Volumes," on page 184.

3h On the Cluster Pool Information page, specify the following information:

Parameter	Action
Virtual Server Name	(Optional) Change the default name of the Virtual Server for the cluster resource.
	The default virtual server name for the resource is the cluster name plus the cluster resource name. For example, if the cluster name is `cluster1` and the pool cluster resource name is `POOL1_SERVER`, then the default virtual server name is `CLUSTER1_POOL1_SERVER`.

Parameter	Action
CIFS Server Name	If CIFS is enabled as an advertising protocol, specify the name of the CIFS virtual server that CIFS clients see when they browse the network. The name can be up to 15 characters.
	A default name is specified in the form of *clustername_poolname_W*. If this exceeds 15 characters, characters are dropped from the left.
	If Novell CIFS is not installed and running, this field value is `NOT_SUPPORTED`.
IP Address	Specify an IP address for the pool cluster resource. Tab between the address fields.
	Each pool cluster resource requires its own unique IP address. The IP address assigned to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.
Advertising Protocols	Select the check boxes of the advertising protocols (AFP, CIFS, NCP) that you want to enable for data requests to this shared pool. NCP is required to support authenticated access to data via the Novell Trustee model.
	Selecting a protocol causes commands to be added to the pool cluster resource's load and unload scripts to activate the protocol for the resource. This lets you ensure that the cluster-enabled pool is highly available to users via the specified protocol.
	Novell CIFS and Novell AFP are available for OES 2 SP1 Linux and later. If Novell CIFS or Novell AFP are not installed and running, selecting the CIFS or AFP check box has no effect.
	For OES 2 Linux and earlier, Novell CIFS and Novell AFP are not available and these options do no apply for pools on Linux. Selecting the CIFS and AFP check boxes has no effect.
Online Resource after Create	The check box is deselected by default and dimmed so that you cannot change the setting.
	The pool is currently active on the server. You must deactivate the pool from the server before attempting to bring the resource online. You should also configure the resource load, unload, and monitoring scripts before you bring the resource online.
Define Additional Properties	Select the *Define Additional Properties* check box.
	This allows you to configure the resource policies for the start, failover, and failback modes, and to configure the preferred nodes.

3i Click *Finish*.

Typically, the pool creation takes less than a minute. However, if you have a large tree or the server does not hold an eDirectory replica, the create time can take up to 3 minutes.

The pool cluster resource should be online and active on the master node if the Mount on Creation option was enabled.

4 Create a volume on the clustered pool:

Repeat the following procedure for each cluster volume that you want to create on the shared pool. We recommend using only one volume per shared pool.

4a In iManager, select *Storage*, then select the *Volumes*.

4b Browse to select the Cluster object (🔴) of the cluster.

Selecting the Cluster object automatically selects the server that is currently the master node in the cluster.

4c Click *New*.

4d Specify the new volume name, then click *Next*.

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

4e Select the check box next to the cluster pool where you want to create the volume, select *Allow the volume to grow to the size of the pool*, then click *Next*.

4f Review and change volume attributes by selecting or deselecting the check boxes next to the attributes.

The Backup and Salvage Files attributes are selected by default.

For information about volume attributes, see "Volume Attributes" in the *OES 2 SP3: NSS File System Administration Guide for Linux*.

4g Choose whether you want the volume activated and mounted when it is created, then click *Finish*.

Typically, the volume creation takes less than 10 seconds. However, if you have a large tree or the server does not hold an eDirectory replica, the create time can take up to 3 minutes.

5 Verify that the pool cluster resource was created and is online.

5a In iManager, select *Clusters > Cluster Options*.

5b Browse to select the Cluster.

5c In the Cluster Objects list, view the pool cluster resource, such as POOL1_SERVER.

5d In iManager, select *Clusters > Cluster Manager*.

5e If the resource is online, the state is *Running*.

6 Continue with Section 12.5, "Configuring a Load Script for the Shared NSS Pool," on page 189.

12.3.3 Creating a Cluster-Enabled Pool and Volume with NSSMU

1 Ensure that the SAN device is attached to all of the nodes in the cluster.

2 On the master node, log in as the root user, then open a terminal console.

3 Start NSSMU by entering nssmu at the terminal console prompt.

4 Create a clustered pool:

4a From the NSSMU main menu, select *Pools*.

4b On the Pools page, press Insert, type a name for the new pool you want to create, then press Enter.

4c From the list of available devices, select the shared device where you want the pool created (such as sdc), then press Enter.

4d Specify the amount of space (in MB) to use, then press Enter.

4e Press F3 to accept the device and size settings.

The Cluster Pool Information page opens automatically because the selected device is shared.

4f For the *Activate On Creation* option, specify whether you want the pool to be activated when it is created, then continue to the next field by pressing Enter.

The *Activate On Creation* option determines if the pool you are creating is to be activated (the resource is brought online) as soon as it is created.

The option is set to *Yes* (enabled) by default. This allows you to create volumes immediately after you create the clustered pool. Default settings apply to the pool cluster resource and scripts, but you can modify them later.

If you set the value to *No* (disabled), you can configure the resource and modify the scripts before you bring the resource online for the first time. You must bring the pool cluster resource online before you can create volumes on it.

4g Specify *Yes* for the *Cluster Enable on Creation* option.

This option is selected by default if the device is shared.

If you set the value to *No*, the values you set on the Cluster Pool Information page are ignored, and a pool cluster resource is not created. After you create a volume on the non-clustered pool, you can cluster-enable the pool and volume as described in Section 12.4, "Cluster-Enabling an Existing NSS Pool and Its Volumes," on page 184.

4h On the Cluster Pool Information page, specify the following information:

Parameter	Action
Virtual Server Name	(Optional) Change the default name of the Virtual Server for the cluster resource.
	The default virtual server name for the resource is the cluster name plus the cluster resource name. For example, if the cluster name is `cluster1` and the pool cluster resource name is `POOL1_SERVER`, then the default virtual server name is `CLUSTER1_POOL1_SERVER`.
CIFS Server Name	If CIFS is enabled as an advertising protocol, specify the name of the CIFS virtual server that CIFS clients see when they browse the network. The name can be up to 15 characters.
	A default name is specified in the form of `clustername_poolname_W`. If this exceeds 15 characters, characters are dropped from the left.
	If Novell CIFS is not installed and running, this field value is `NOT_SUPPORTED`.
IP Address	Specify an IP address for the pool cluster resource. Tab between the address fields.
	Each pool cluster resource requires its own unique IP address. The IP address assigned to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.
Advertising Protocols	Select the check boxes of the advertising protocols (AFP, CIFS, NCP) that you want to enable for data requests to this shared pool. NCP is required to support authenticated access to data via the Novell Trustee model.
	Selecting a protocol causes commands to be added to the pool cluster resource's load and unload scripts to activate the protocol for the resource. This lets you ensure that the cluster-enabled pool is highly available to users via the specified protocol.
	Novell CIFS and Novell AFP are available for OES 2 SP1 Linux and later. If Novell CIFS or Novell AFP are not installed and running, selecting the CIFS or AFP check box has no effect.
	For OES 2 Linux and earlier, Novell CIFS and Novell AFP are not available and these options do no apply for pools on Linux. Selecting the CIFS and AFP check boxes has no effect.

The default name of a pool cluster resource is the pool name plus the word "SERVER", such as `POOL1_SERVER`. You can modify the resource name after the resource has been created by using the `cluster rename` command. For information, see Section 10.13, "Renaming a Cluster Resource," on page 166. Changing the resource name does not modify the pool name or the virtual server name.

4i Select *Apply*, then press Enter to create and cluster-enable the pool.

Typically, the pool creation takes less than a minute. However, if you have a large tree or the server does not hold an eDirectory replica, the create time can take up to 3 minutes.

5 Press Esc to return to the NSSMU main menu.

6 Create a volume on the clustered pool.

Repeat the following procedure for each cluster volume that you want to create on the shared NSS pool. We recommend using only one volume per shared pool.

6a On the NSSMU main menu, select *Volumes*.

6b On the Volumes page, press Insert, type a name for the new volume you want to create, then press Enter.

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

6c Specify Y(es) to encrypt volume or N(o) to create a regular volume.

6d From the list of available pools, select the clustered pool where you want the volume to reside, then press Enter.

The new volume appears in the list of volumes.

Typically, the volume creation takes less than 10 seconds. However, if you have a large tree or the server does not hold an eDirectory replica, the create time can take up to 3 minutes.

6e (Optional) From the list of volumes, select the newly created volume, press F8 to view more options, press Enter to open the *Properties* page to review and change volume attributes, then select *Apply* and press Enter to save any changes.

The Backup and Salvage Files attributes are selected by default.

For information about volume attributes, see "Volume Attributes" in the *OES 2 SP3: NSS File System Administration Guide for Linux*.

6f Exit NSSMU.

7 Verify that the pool cluster resource was created.

7a In iManager, select *Clusters > Cluster Options*.

7b Browse to select the Cluster.

7c In the Cluster Objects list, view the pool cluster resource, such as `POOL1_SERVER`.

7d In iManager, select *Clusters > Cluster Manager*.

7e If the resource is online, the state is *Running*.

8 Continue with Section 12.5, "Configuring a Load Script for the Shared NSS Pool," on page 189.

12.4 Cluster-Enabling an Existing NSS Pool and Its Volumes

Cluster-enabling a pool allows it (and its volumes) to be moved or mounted on different servers in the cluster in a manner that supports transparent client reconnect. This makes the data highly available to users.

IMPORTANT: You can cluster-enable pools when you create them by following the procedures in Section 12.3, "Creating Cluster-Enabled Pools and Volumes," on page 178.

The procedure in this section describes how to enable clustering for an existing pool and its volumes. It assumes the following:

- The pool contains at least one volume.
- The pool and its volumes have Storage objects in the eDirectory tree where you are setting up the cluster.

  If the objects are missing, you can use the Update eDirectory options in the Storage plug-in for NSS to create them in eDirectory. For information, see "Updating eDirectory Pool Objects" and "Updating eDirectory Volume Objects" in the *OES 2 SP3: NSS File System Administration Guide for Linux*.

- The pool resides on a disk that has been enabled as *Sharable for Clustering*.
- Each pool resource requires its own static IP address that is in the same IP subnet as the cluster IP address.
- To add CIFS as an advertising protocol when you cluster-enable an existing pool, you must install and configure Novell CIFS on the server, and ensure that the CIFS daemon is running and functioning properly before you begin. The CIFS daemon should also be running before trying to online the resource.
- To add AFP as an advertising protocol when you cluster-enable an existing pool, you must install and configure Novell AFP on the server, and ensure that the AFP daemon is running and functioning properly before you begin. The AFP daemon should also be running before trying to online the resource.

Before you cluster-enable the pool, ensure that it is activated on the master node in the cluster. This allows the pool and volume information to be automatically added to the load, unload, and monitor scripts for the pool cluster resource. In addition, the volume entry is automatically removed from the /etc/fstab so that the pool cluster resource load script controls when the volume is mounted. If you use special settings for mounting the volume, you must manually add those settings to the load script.

The default name of a pool cluster resource is the pool name plus the word "SERVER", such as POOL1_SERVER. You can modify the resource name after the resource has been created by using the cluster rename command. For information, see Section 10.13, "Renaming a Cluster Resource," on page 166. Changing the resource name does not modify the pool name or the virtual server name.

After the pool is cluster-enabled, users can access the cluster-enabled pool by using the pool cluster resource's IP address or virtual server name. The default virtual server name of a pool cluster resource is the cluster name plus the default cluster resource name, such as CLUSTER1_POOL1_SERVER. You can specify the name of the virtual server name when you cluster-enable the pool.

To cluster-enable an existing NSS pool:

1 Ensure that the SAN device is attached to all of the nodes in the cluster.

2 Log in to iManager as an administrator user.

3 Deactivate the shared pool on the current server.

   3a In *Roles and Tasks*, select *Storage > Pools*.

   3b Browse to select the server where the pool is currently assigned.

   3c In the Pools list, select the pool, then click *Deactivate*.

      Wait until the page refreshes and confirms that the pool is deactive.

4 Verify that the device is marked as shareable for clustering. If the pool uses multiple devices, you must verify the shared setting of each device.

    4a In *Roles and Tasks*, select *Storage > Devices*.

    4b Browse to select the server where the pool is currently assigned.

    4c Select the device to view its Details page.

    4d If the *Shareable for Clustering* check box is deselected, select the check box, click *Apply*, read the warning, then click *Continue* to enable sharing.

       Marking the device as shareable automatically modifies the Shared attribute on the eDirectory objects for the pools and volumes on the device.

       Wait a couple of minutes before continuing to allow the attribute change to synchronize in eDirectory. You can verify the change in the shared state of the pool and its volumes by returning to the *Pools* page and *Volumes* page.

5 In *Roles and Tasks*, select *Clusters > Cluster Options*, then browse to select the Cluster object ( ) of the cluster.

6 On the Cluster Options page, click the *New* link in the Cluster Objects toolbar.

7 On the Resource Type page, specify Pool ( ) as the resource type you want to create by clicking the *Pool* radio button, then click *Next*.

8 On the Cluster Pool Information page, specify the following information:

Parameter	Action
Pool Name	Specify the *Pool Name* by browsing to select the Pool object ( ) of the pool you want to cluster-enable, such as *servername*_POOL1_POOL.
Virtual Server Name	(Optional) Change the default name of the Virtual Server for the cluster resource.
	The default virtual server name for the resource is the cluster name plus the cluster resource name. The default name of a pool cluster resource is the pool name plus the word "SERVER", such as POOL1_SERVER.
	For example, if the cluster name is cluster1 and the pool cluster resource name is POOL1_SERVER, then the default virtual server name is CLUSTER1_POOL1_SERVER.
CIFS Server Name	If CIFS is enabled as an advertising protocol, specify the name of the CIFS virtual server that CIFS clients see when they browse the network. The name can be up to 15 characters.
	A default name is specified in the form of *clustername_poolname*_W. If this exceeds 15 characters, characters are dropped from the left.
	If Novell CIFS is not installed and running, this field value is NOT_SUPPORTED.
IP Address	Specify an IP address for the pool cluster resource.
	Each pool cluster resource requires its own unique IP address. The IP address assigned to the pool remains assigned to the pool regardless of which server in the cluster is accessing the pool.

Parameter	Action
Advertising Protocols	Select the check boxes of the advertising protocols (AFP, CIFS, NCP) that you want to enable for data requests to this shared pool. NCP is required to support authenticated access to data via the Novell Trustee model.
	Selecting a protocol causes commands to be added to the pool cluster resource's load and unload scripts to activate the protocol for the resource. This lets you ensure that the cluster-enabled pool is highly available to users via the specified protocol.
	Novell CIFS and Novell AFP are available for OES 2 SP1 Linux and later. If Novell CIFS or Novell AFP are not installed and running, selecting the CIFS or AFP check box has no effect.
	For OES 2 Linux and earlier, Novell CIFS and Novell AFP are not available and these options do no apply for pools on Linux. Selecting the CIFS and AFP check boxes has no effect.
Deactivate the pool on the master node, and online resource after create	The check box is selected by default.
	If the pool is currently deactive, the resource is brought online automatically after you create the resource.
	If the pool is currently active as a local pool on the master node of the cluster, this option automatically deactivates the pool on the master node, and brings the resource online on a preferred node immediately after the pool cluster resource is created.
	If you deselect the option, the resource is in an offline state after it is created. You should configure the resource load, unload, and monitoring scripts before you bring the resource online. If the pool was in an active state when you started to cluster-enable it, the pool is still mounted locally on the master node. After you set up preferred nodes for the resource, you can deactivate the pool on the master node, and bring the resource online on a preferred node.
Define Additional Properties	Select the *Define Additional Properties* check box.
	This allows you to configure the resource policies for the start, failover, and failback modes, and to configure the preferred nodes.

9 Click *Next*.

10 On the Resource Policies page, configure the resource policies for the start, failover, and failback modes.

For information, see "Configuring the Start, Failover, and Failback Modes for Cluster Resources" on page 159.

11 Click *Next*.

12 On the Resource Preferred Nodes page, assign the preferred nodes to use for the resource.

For information, see "Configuring Preferred Nodes for a Resource" on page 160.

13 Click *Finish*.

The pool cluster resource appears in the *Cluster Objects* list on the Cluster Options page, such as POOL1_SERVER.

The load, unload, and monitoring scripts are automatically created with information about the pools and volumes. The volume entries are removed from the /etc/fstab file so that the resource load script can be used to control when the pools and volumes are mounted. The resource is offline.

14 Verify that the pool cluster resource was created and that it is offline.

14a Select *Clusters > Cluster Options*, then view the pool cluster resource in the Cluster Objects list, such as `POOL1_SERVER`.

14b Select *Clusters > Cluster Manager*, then view the pool cluster resource state as *Offline*.

15 Manually deactivate the pool on its currently assigned node.

You must deactivate the pool from the server before attempting to bring the pool cluster resource online.

15a In *Roles and Tasks*, select *Storage > Pools*.

15b In the Server field, browse to select the Cluster object (🔴) of the cluster.

The Cluster object is automatically associated with the node that is currently the master node in the cluster.

15c In the Pools list, select the pool, then click *Deactivate*.

Wait until the page refreshes and confirms that the pool is deactive.

16 Verify the cluster resource scripts for the pool cluster resource, and modify them if desired.

16a Select *Clusters > Cluster Options*, then select the cluster.

16b Click the pool cluster resource link to open its Properties page.

16c Configure the following resource settings:

Property	Action
Monitoring	(Optional) Select the *Monitoring* tab, select *Enable Resource Monitoring*, and configure monitoring as described in Section 10.6, "Enabling Monitoring and Configuring the Monitor Script," on page 155.
Load script	Select the Scripts tab, select Load Script, then modify the script as described in Section 12.5, "Configuring a Load Script for the Shared NSS Pool," on page 189.
Unload Script	Select the Scripts tab, select Unload Script, then modify the unload script as described in Section 12.6, "Configuring an Unload Script for the Shared NSS Pool," on page 190.
Monitoring Script	Select the Scripts tab, select Monitor Script, then modify the unload script as described in Section 10.6, "Enabling Monitoring and Configuring the Monitor Script," on page 155.

16d At the bottom of the page, click *Apply* or *Ok* to save your changes.

The changes are not applied until the next time the resource is brought online.

17 Select *Clusters > Cluster Manager*, select the check box next to the pool cluster resource, then click *Online* to bring the resource online.

The pool is activated and its volumes are mounted on the primary preferred node that is configured for the pool cluster resource.

12.5 Configuring a Load Script for the Shared NSS Pool

A cluster resource load script is automatically generated for the pool when you cluster-enable it. You can modify the script as needed to suit your needs by using iManager. For information about how to access scripts for the cluster resource, see Section 10.4, "Configuring a Load Script for a Cluster Resource," on page 153.

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

If you need to modify the IP address, administrator credentials, or other attributes of an existing resource, follow the procedure in Section 8.9, "Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster," on page 116.

If you specified the following values for the variables in the template, your load script would appear like the script below.

Variable	Your Value
Cluster resource's virtual server name	NCS1_SHPOOL43_SERVER
Resource IP address	10.10.10.43
Pool name	SHPOOL43
Volume name	SHVOL43
Volume ID	252 (valid values are 0 to 254)

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

exit_on_error add_secondary_ipaddress 10.10.10.43
exit_on_error nss /poolact=SHPOOL43
exit_on_error ncpcon mount SHVOL43=252

exit_on_error ncpcon bind --ncpservername=NCS1_SHPOOL43_SERVER --
ipaddress=10.10.10.43

exit 0
```

If you change the name space for an existing shared volume by using NSSMU or the NSS plug-in for iManager, you must modify the load script for the pool cluster resource to add the name space to the `ncpcon mount` command for the volume. Otherwise, the cluster assumes the default name space for mounting the volume. You can do this by using the `/opt=ns=<long|unix|dos|mac>` switch in the `ncpcon mount` command. Place the `/opt` switch before the volume information. The options are case-sensitive. For information about default name spaces for NSS volumes, see "Lookup Namespace" in the *OES 2 SP3: NSS File System Administration Guide for Linux*.

For example, to specify the Long name space, add the `/opt=ns=long` switch as follows:

```
ncpcon mount /opt=ns=long <VOLUMENAME>=<VOLUMEID>
```

For example, to specify the Unix name space, add the `/opt=ns=unix` switch as follows:

```
ncpcon mount /opt=ns=unix <VOLUMENAME>=<VOLUMEID>
```

12.6 Configuring an Unload Script for the Shared NSS Pool

A cluster resource unload script is automatically generated for the pool when you cluster-enable it. You can modify the script as needed to suit your needs by using iManager. For information about how to access the scripts for a cluster resource, see Section 10.5, "Configuring an Unload Script for a Cluster Resource," on page 154.

If you specified the following values for the variables in the template, your unload script would appear like the script below.

Variable	Your Value
Cluster resource's virtual server name	NCS1_SHPOOL43_SERVER
Resource IP address	10.10.10.43
Pool name	SHPOOL43
Volume name	SHVOL43
Volume ID	252 (valid values are 0 to 254)

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

ignore_error ncpcon unbind --ncpservername=NCS1_SHPOOL43_SERVER --
ipaddress=10.10.10.43

ignore_error nss /pooldeact=SHPOOL43
ignore_error del_secondary_ipaddress 10.10.10.43

exit 0
```

By default, `ncpcon dismount` commands are not added to the unload script. The pool deactivation automatically dismounts the volumes on the pool. However, you can add `ncpcon dismount` commands for each volume before the pool deactivation line. For example:

```
ignore_error ncpcon dismount SHVOL43
ignore_error nss /pooldeact=SHPOOL43
```

Adding the `ncpcon dismount` command provides predictable and clean dismount behavior. It allows the volume dismount to occur and be logged independently of the pool deactivation, which can help with troubleshooting dismount issues. However, the ncpcon dismount commands are not automatically maintained and updated if you change a volume name or delete a volume. You must modify the unload script yourself.

12.7 Configuring a Monitor Script for the Shared NSS Pool

A cluster resource monitor script is automatically generated for the pool when you cluster-enable it. It is disabled by default. To enable or disable monitoring, see Section 10.6, "Enabling Monitoring and Configuring the Monitor Script," on page 155.

The following is a sample monitoring script for an NSS pool cluster resource:

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

exit_on_error status_fs /dev/evms/POOL1 /opt/novell/nss/mnt/.pools/POOL1 nsspool

exit_on_error status_secondary_ipaddress 10.10.10.41
exit_on_error ncpcon volume VOL1

exit 0
```

You can add other monitoring commands to the script. For example, the CIFS monitor command helps to keep CIFS up and running. It checks to see if the `cifsd` daemon is running and takes action if it is not running. If CIFS is running, it returns its status. If CIFS is dead or not running, it starts the `cifsd` daemon and returns its status. If the restart fails, the resource's failover settings are invoked, such as failing over the resource to its preferred node, or going comatose. Add the following line above the exit command:

```
exit_on_error rcnovell-cifs monitor
```

The `monitor` command is the default line added by CIFS for newly created pool cluster resources on OES 2 SP3 servers with the September 2011 or later patches applied. Existing pool cluster resources are not automatically updated to use it. Instead, they use the CIFS `status` command in the monitor script. It checks the status and immediately invokes the resource's failover settings if the `cifsd` daemon is not running. To take advantage of the restart capability that is offered by the `monitor` command, you can modify the line to use "`monitor`" instead of "`status`".

You can also add lines to the monitor script that check the status of processes used by clustering. For information, see Section 10.6.4, "Monitoring Services That Are Critical to Clustering," on page 158.

12.8 Adding Advertising Protocols for NSS Pool Cluster Resources

Shared NSS pools support three advertising protocols: NCP (default, required), Novell AFP, and Novell CIFS. You can add Novell AFP or Novell CIFS as advertising protocols for an existing pool cluster resource.

You must install Novell CIFS or Novell AFP on all the servers where you want to fail over the pool resource, but you set up the advertising protocol only once on the pool cluster resource. Before you begin, ensure that the pool is cluster migrated back to the server where it was originally created.

1 Log in to iManager as an administrator user.

2 Offline the cluster resource that you want to modify.

For information, see Section 9.5, "Onlining and Offlining (Loading and Unloading) Cluster Resources from a Cluster Node," on page 130.

3 In *Roles and Tasks*, select *Storage*, then go the Volumes and Pools pages to verify that the volumes are not mounted and the pool is deactive.

4 In *Roles and Tasks*, select *Clusters*, click *Cluster Options*, then browse to locate and select the Cluster object of the cluster you want to manage.

5 On the Cluster Options page, use one of the following methods to open the Cluster Pool Properties page for the pool cluster resource that you want to manage:

- Select the check box next to the pool cluster resource, then click *Details*.
- Click the *Name* link for the pool cluster resource.
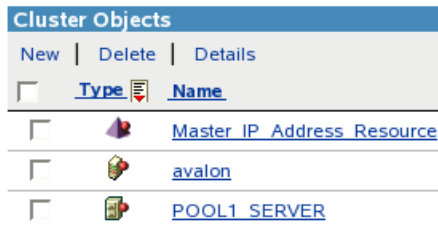
6 On the Cluster Pool Properties page, click the *Protocols* tab.

7 If you are enabling the resource for Novell AFP, select the *AFP* check box, then click *Apply*.

8 If you are enabling the resource for Novell CIFS, do the following:

8a On the Protocols page, select the *CIFS* check box and view the *CIFS Virtual Server Name*, then click *Apply*.

A default name (up to 15 characters) is suggested for the CIFS Virtual Server Name when you first add the advertising protocol.

After you have applied this setting, the CIFS share is created with the name you specified. If you later want to modify the CIFS Virtual Server Name, you must use the CIFS plug-in for iManager to manage the CIFS share name of an existing share. For information, see "Setting CIFS General Server Parameters" (http://www.novell.com/documentation/oes2/file_cifs_lx/data/bdkfk5j.html#af3ozzn) in the *OES 2 SP3: Novell CIFS for Linux Administration Guide*.

8b If the server is OES 2 SP1 or SP2 Linux, run the cifsPool.py script to modify the load and unload scripts. This step is not required in OES 2 SP3 and later.

IMPORTANT: In OES 2 SP1 and SP2 Linux, the Novell CIFS commands are not automatically added to the load and unload scripts for the NSS pool resource when you add Novell CIFS as an advertising protocol on an existing cluster-enabled pool. You must run the cifsPool.py script to add commands for Novell CIFS.

You can download the cifsPool.py script file as a patch Novell Downloads Web site (http://download.novell.com/Download?buildid=sA5Tdb_U2CE~).

This issue has been resolved in OES 2 SP3.

At a terminal console prompt, log in as the root user, then enter

```
python cifsPool.py pool_resource_DN cifs_servername ldaps://
ldap_server_ip:636 admin_DN admin_password
```

For example, enter

```
python cifsPool.py cn=CLUS1_POOL1_SERVER,cn=clus1,o=company CIFS_POOL
ldaps://10.10.10.1:636 cn=admin,o=company pa$$w0rD
```

where the values are:

Parameter	Sample Value
pool_resource_DN	cn=CLUS1_POOL1_SERVER,cn=clus1,o=company **NOTE:** The cluster container is stored in eDirectory as a CN and not an OU.
cifs_servername (up to 15 characters)	CIFS_POOL1
ldap_server_ip	10.10.10.1
admin_DN	cn=admin,o=company
admin_password	pa$$w0rD

9 Online the cluster resource.

For information, see Section 9.5, "Onlining and Offlining (Loading and Unloading) Cluster Resources from a Cluster Node," on page 130.

12.9 Adding NFS Export for a Clustered Pool Resource

You can add NFS support for a clustered NSS volume by adding an `exportfs(8)` command to the load script and unload script for a clustered pool resource. Basically, the order in load script should be NSS, NFS, IP, and NCP. The order is reversed in the unload script.

Before you use the `exportfs(8)` command in resource scripts, ensure that you install the `nfs-kernel-server` package on all of the cluster nodes in the resource's preferred nodes list. This contains the support utilities for the kernel `nfsd` daemon. It is not installed by default. You can use the *YaST > Software > Software Management* tool to find and install the package.

Figure 12-1 *Using YaST Software Management to Install the nfs-kernel-server Package*
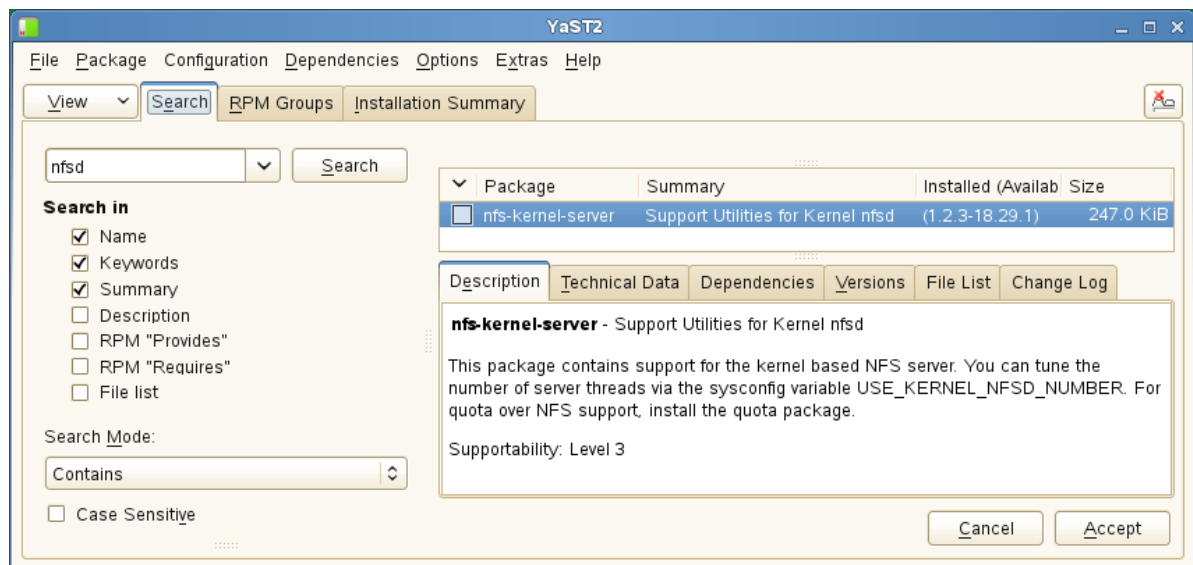
◆ Section 12.9.1, "Sample NFS Load Script," on page 194

◆ Section 12.9.2, "Sample NFS Unload Script," on page 194

12.9.1 Sample NFS Load Script

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

# Activate the NSS pool and mount its volume
exit_on_error nss /poolact=POOL_16
exit_on_error ncpcon mount VOL_161=224

# Start the NFS service
exit_on_error rcnfsserver start

# Export the volume for NFS
exit_on_error exportfs -o rw,sync,no_root_squash,fsid=216 *:/media/nss/VOL_161

exit_on_error add_secondary_ipaddress 192.168.188.16

exit_on_error ncpcon bind --ncpservername=CLUS1_POOL_16_SERVER --ipaddress=192.168.188.16

exit 0
```

12.9.2 Sample NFS Unload Script

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

ignore_error ncpcon unbind --ncpservername=CLUS1_POOL_16_SERVER --ipaddress=192.168.188.16

ignore_error del_secondary_ipaddress 192.168.188.16

#Unexport a volume for NFS
ignore_error exportfs -u *:/media/nss/VOL_161

# Deactivate the pool and volume
ignore_error nss /pooldeact=POOL_16

exit 0
```

12.10 Mirroring and Cluster-Enabling Shared NSS Pools and Volumes

12.10.1 Understanding NSS Mirroring

NSS mirroring is a checkpoint-based synchronous mirroring solution. Data blocks are written synchronously to multiple storage devices. It is an alternative to synchronous replication options that are provided by a SAN array.

Hardware configuration and placement for NSS mirroring can vary depending on geography and fault tolerance requirements. For example, Figure 12-2 depicts a simple hardware configuration in which one side of the mirrored NSS volume is located in a separate building from the rest of the cluster hardware. If a disaster occurs in one building, data is still safe on the mirrored NSS volume in the other building.

Figure 12-2 *Single Cluster with Mirrored NSS Volumes in Separate Buildings*
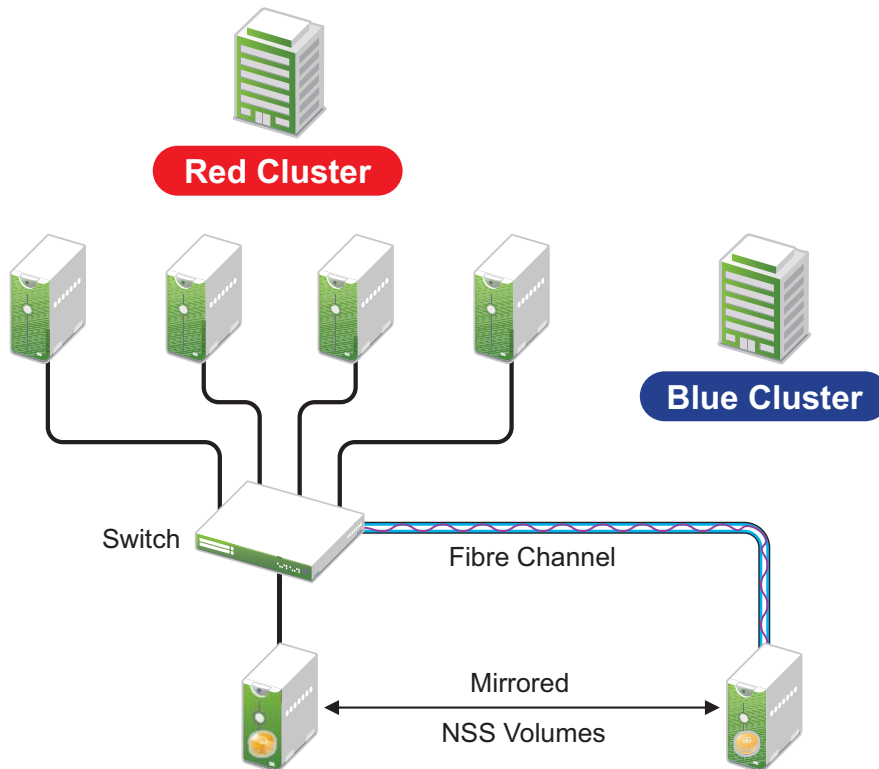
Figure 12-3 depicts a more complex hardware configuration in which two clusters are placed in separate buildings, but each has one side of a mirrored NSS volume in the building with the other cluster. If a disaster occur in one building, data is still safe and immediately accessible on the mirrored NSS volume in the other building.

Figure 12-3   *Two Clusters with Mirrored NSS Volumes in Separate Buildings*
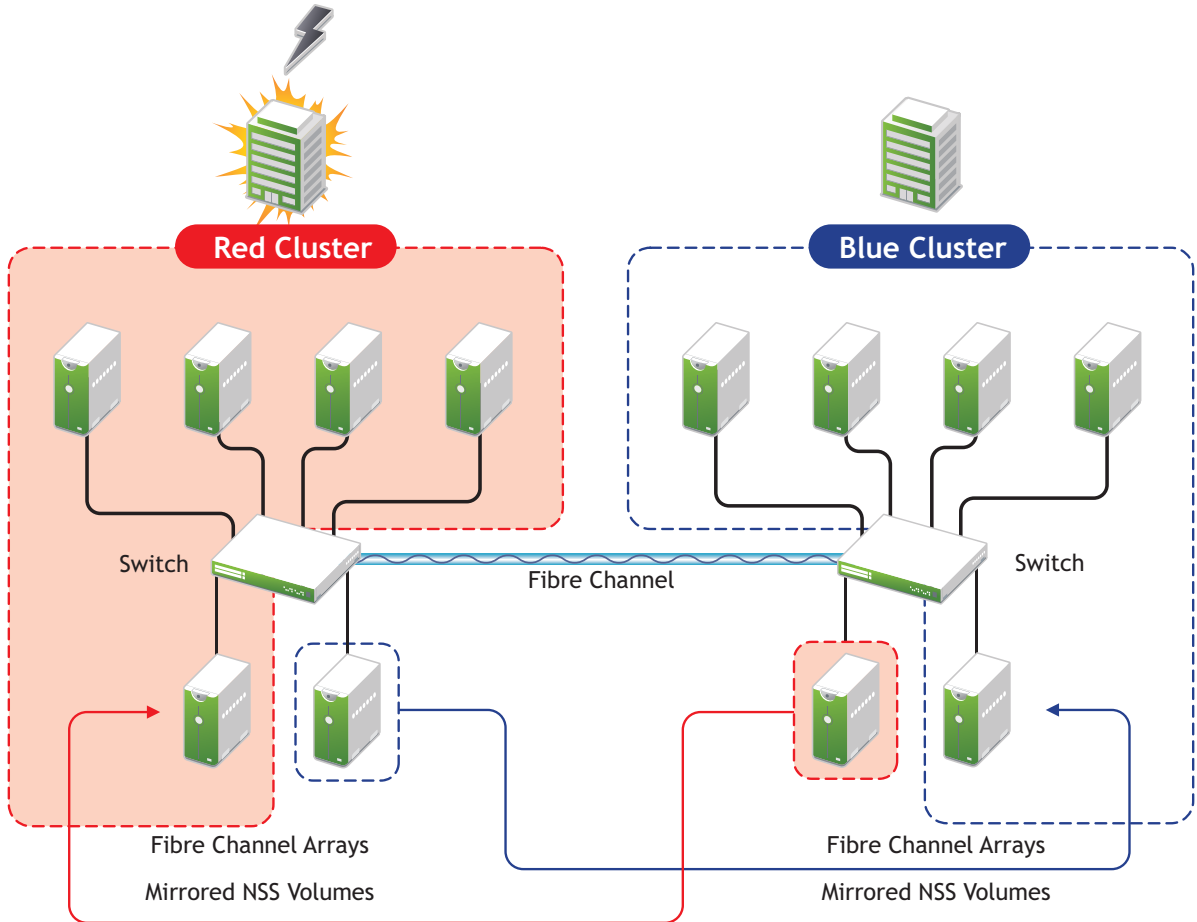
The maximum distance for a single stretch of fibre without optical repeaters to boost distance is about 10 kilometers. There are various devices available that convert Fibre Channel to IP in order to extend SANs to WAN scale.

NSS mirroring is synchronous, and so mirror I/Os must complete before the next I/O can be sent. Performance is a function of time to do an I/O over the longest link.

12.10.2   Requirements for NSS Mirroring

NSS partitions must be mirrored after they are created. If you have an existing partition that you want to mirror, you can either create another partition of equal size on another device to mirror the first partition to, or let the mirroring software automatically create another partition of equal size on another device.

Novell Cluster Services should be installed and running prior to creating and mirroring partitions on shared storage.

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

12.10.3 Creating and Mirroring NSS Pools on Shared Storage

Prior to creating and mirroring NSS partitions on shared storage, ensure that you have the following:

- OES 2 Linux is installed on all servers that will be part of the cluster
- All servers in the cluster are connected to a shared storage system
- One or more drive arrays are configured on the shared storage system
- At least 20 MB of free space on the shared storage system for a special cluster partition
- To ensure disaster recovery, the device you select to mirror should be in a different storage array.

Use one of the following methods to mirror clustered pool resources:

- "Creating a Pool on a Shared Mirrored Device" on page 197
- "Mirroring the Partition of an Existing NSS Pool" on page 198

Creating a Pool on a Shared Mirrored Device

To create an NSS pool resource on a shared mirrored device:

1 Start NSSMU by entering `nssmu` at the terminal console prompt of a cluster server.

2 Share the devices.

   2a Select *Devices* from the NSSMU main menu.

   2b Select the two devices that you want to use, then press F6 to mark the devices as *Sharable for Clustering*.

3 Create a mirrored RAID device using the shared devices.

   3a Select *RAID Devices* from the NSSMU main menu.

   3b Press Insert to create a software RAID, select RAID 1, then press Enter.

   3c Use the arrow keys to select the two shared devices that you want to contribute space to the RAID, specify the amount of space to use, then press Enter.

      The software RAID device automatically inherits the Sharable for Clustering setting.

4 Create a clustered NSS pool on the shared RAID device.

   4a Select *Pools* from the NSSMU main menu.

   4b Press Insert to create a new pool, then create the clustered pool as usual on the shared software RAID device as described in Section 12.3.3, "Creating a Cluster-Enabled Pool and Volume with NSSMU," on page 182.

Mirroring the Partition of an Existing NSS Pool

You can use NSSMU to mirror a clustered pool's partition. The second device must have as much space available as you know the pool's real allocated size to be, and be marked as Sharable for clustering. To mirror the partition for an existing clustered NSS pool resource:

1 Start NSSMU by entering `nssmu` at the terminal console prompt of a cluster server.

2 Share the second device.

    2a Select *Devices* from the NSSMU main menu.

    2b Select the device that you want to use, then press F6 to mark it as *Sharable for Clustering*.

    To ensure disaster recovery, the device you select to mirror should be in a different storage array.

3 Select *Partitions* from the NSSMU main menu.

4 Press F3, choose the partition that contains the pool (even though it might not look like it has sufficient free space) and select the second shareable partition, then select *Yes*.

This mirrors the existing pool and creates the RAID 1 device where each segment is the size of the pool.

5 Use one of the following methods to initiate mirroring for the newly created mirror:

    ◆ At the terminal console of a cluster node, enter the following to migrate the cluster resource to another node:

```
cluster migrate cluster_resource destination_node_name
```

    Migrating the pool causes load scripts to be executed and causes the mirroring to start on the new node.

    ◆ At the terminal console of the cluster node where the pool is currently active, enter

```
dmsetup message raid_device_name 0 remirror=on
```

    WARNING: Issue this command only on the node where the pool is currently active. Issuing the command on multiple nodes can corrupt the mirror.

6 Verify that the remirroring has begun by opening NSSMU on the node where the pool is currently active, open the RAID page, then select the RAID device.

The remirroring status shows a percentage that is greater than 0. It is fully synchronized at 100%.

12.10.4 Verifying the NSS Mirror Status in the Cluster

After you have configured NSS mirroring with Novell Cluster Services, you should check to ensure that it is working properly in a cluster environment.

1 Ensure that the volumes on the cluster-enabled pool are mounted on an assigned server by entering `volumes` at the terminal console.

2 Check the mirror status of the mirrored partition by entering `mirror status` at the terminal console of the server where the NSS pool on the mirrored partition is active.

After entering `mirror status`, you should see a message indicating that mirror status is 100 percent or a message indicating that the mirrored object is fully synchronized.

3 Migrate the pool to another server in the cluster and again check to ensure that the volumes on the pool are mounted by entering `volumes` at the terminal console.

4 Check the mirror status of the partition again by entering `mirror status` at the terminal console.

12.11 Renaming a Pool for a Pool Cluster Resource

When a cluster-enabled NSS pool is used as a Novell Cluster Services cluster resource, special steps must be followed if you decide to rename the pool. Renaming the Pool automatically changes the Pool Resource object name in eDirectory. It does not change the virtual server name (the NCS:NCP Server object).

Because renaming involves changing information for the Pool object in eDirectory, you must ensure that the shared pool is active on a cluster node that has its NCP server object in the same context as the Pool object of the pool you are going to rename.

1 If necessary, use the cluster migrate command to move the shared pool to a cluster node that has a Server object in the same context as the Pool object of the pool you are going to rename.

   The server must be in the resource's *Preferred Nodes* list.

   1a In iManager, select *Clusters > Cluster Manager*, then select the cluster that you want to manage.

   1b On the Cluster Manager page, select the check box next to the pool cluster resource, then click *Migrate*.

   1c Select the cluster node where you want to move the resource, then click *OK*.

      Wait for the resource to report a Running state on the target node before you continue.

2 Offline the pool cluster resource. On the Cluster Manager page, select the check box next to the pool cluster resource, then click *Offline*.

   Wait for the resource to report an Offline state before you continue.

3 Activate the pool locally on the cluster node:

   3a In iManager, select *Storage > Pools*.

   3b Select the server where you want to activate the pool.

   3c Select the pool, then click Activate.

4 In iManager on the Pools page, select the pool, then click Rename, specify the new name, then click OK.

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

5 Online the pool cluster resource:

   5a In iManager, select *Clusters > Cluster Manager*, then select the cluster that you want to manage.

   5b On the Cluster Manager page, select the check box next to the pool cluster resource, then click *Online*.

      The resource is brought online on a node in its Preferred Nodes list, according to availability. Wait for the resource to report a Running state before you continue.

      Novell Cluster Services automatically updates the pool resource load, unload, and monitor scripts to reflect the name change.

6 If you previously used the cluster rename command to use a customized name for the cluster resource, you must repeat that process to return to the customized name. For information, see Section 10.13, "Renaming a Cluster Resource," on page 166.

12.12 Adding a Volume to a Clustered Pool

When you use NSSMU or the Storage plug-in to iManager to add an NSS volume to a clustered pool, a `mount` command is automatically added to the load script for the pool cluster resource. Offline and online the primary pool cluster resource to apply the modified load script.

If you add a volume to the primary pool for a clustered DST shadow volume, the `mount` command is added twice in the primary pool's cluster load script, once after the primary pool's activation command and once after the secondary pool's activation command. You must manually delete the instance that occurs after the secondary pool's activation, then offline and online the primary pool cluster resource to apply the modified load script.

1 In iManager under *Roles and Tasks*, select *Storage > Volumes*.

2 Browse to select the Cluster object of the cluster where the pool is currently online.

  You can also select the server where the pool cluster resource is currently mounted.

3 Click *New* to open the *New Volume* wizard, then create the new volume on the shared pool.

  For information, see "Creating Unencrypted NSS Volumes" in the *OES 2 SP3: NSS File System Administration Guide for Linux*.

4 Under *Roles and Tasks*, select *Clusters > Cluster Manager*, then browse to select the Cluster object of the cluster.

5 If you added a volume to the primary pool for a DST shadow volume, modify the load script to remove the instance of the mount command that immediately follows the pool activation line for the secondary pool.

  5a Select *Clusters > Cluster Options*.

  5b Click the name link of the pool cluster resource to open its Properties page, then click the *Scripts* tab.

  5c Delete the `mount` command for the new volume that follows the pool activation command for the secondary pool.

  5d Click OK to save your changes.

  5e Click *Clusters > Cluster Manager* to return to the Cluster Manager page.

6 In the *Cluster Resource* list, select the check box next to the pool cluster resource, then click *Offline* to take the resource offline.

7 In the *Cluster Resource* list, select the check box next to the pool cluster resource, then click *Online* to bring the resource online.

  This applies the modified load script, which mounts the newly created volume.

12.13 Changing the Name Space for a Clustered NSS Volume

The load script for an NSS volume assumes that the default name space of LONG is used when mounting the clustered NSS volume. If you change the name space setting for shared volume by using NSSMU or the Storage plug-in for iManager, you must modify the load script for the pool cluster resource to add the name space to the `ncpcon mount` command for the volume. For information, see Section 12.5, "Configuring a Load Script for the Shared NSS Pool," on page 189.

1 Open Novell iManager in a Web browser, then log in as a cluster administrator user.

2 In *Roles and Tasks*, select *Clusters > Cluster Manager*, then select the cluster.

3 Modify the mount point path value in the load scripts for the pool cluster resource:

   3a On the Cluster Manager page, select the resource's name link to open its Cluster Properties page, then click the *Scripts* tab.

      The Scripts tab automatically displays the load script.

   3b Modify the load script:

      3b1 In the load script, add the `/opt=ns=<long|unix|dos|mac>` option for the `ncpcon mount` command and specify the custom name space to use for the volume. For example:

           `exit_on_error ncpcon mount /opt=ns=unix USERS=254`

      3b2 Click *Apply*, then click *OK* to acknowledge the confirmation message.

   3c At the bottom of the page, click *OK* to close the Properties page and save your changes.

      The changes do not take effect until the resource is taken offline and brought online.

4 Take the resource offline, and bring the resource online in order for the script changes to take effect.

   4a On the Cluster Manager page, select the check box next to the resource, then click *Offline*. Wait for the status to report that it is offline, then continue.

   4b Select the check box next to the resource, then click *Online*.

   4c Verify that the resource comes online and reports a *Running* state.

      If the resource goes into a Comatose state, it is probably because you made a mistake in the lines you added or modified in the script. Take the resource offline, then go back to correct the scripts, and try to bring it online again.

12.14 Changing an Assigned Volume ID

Novell Cluster Services supports NCP client access to cluster-enabled NSS volumes by using a unique volume ID to mount the volume in the cluster. The volume ID is used by an NCP client only for automatic reconnects to the volume after failover or migration of a cluster resource.

Valid volume ID values are 0 to 254 (up to 255 mounted volumes per server). When you create a new volume on a cluster-enabled pool, Cluster Services automatically assigns it a volume ID that is unique in the entire cluster and writes the value to the cluster resource load script for the pool. Values start at 254 for the first volume in the cluster and decrease for each new volume. You can view the volume IDs assigned on a node by using the `ncpcon volumes` command.

In older operating systems, there was a mounted volume limit of 64 volumes (values 0 to 63). Some older applications might have hard-coded the old maximum limit of 64 mounted volumes, and might not be able to handle volume IDs greater than 63. You can use the Clusters plug-in to iManager to modify the volume ID in the scripts for a given cluster resource in order to specify a value that works for the application.

Changing the volume ID does not affect the ability to log in to, back up, or access the data. However, there is a brief disruption of service as the cluster resource is offlined and onlined to apply the script changes. If you modify the volume ID for a volume in the cluster resource scripts, ensure that you do the following:

❒ Volume IDs that you manually assign must be unique across every volume on all servers in cluster.

❒ After the value is changed, you must offline and online the cluster resource for the volume in order to mount the volume with its new volume ID.

❐ After the volume is mounted with its new ID, the clients must log out and log in to the volume in order to reconnect to the volume with its new volume ID. Automatic reconnection after cluster resource failovers or migrations occurs properly after this one-time reset.

Some clients might cache the volume IDs. To reset the cached value, the client must be rebooted and reconnected to the volume.

❐ After the volume is mounted with its new ID, if the backup software is running on a client connection, you might need to restart the backup to reconnect to the volume with its new volume ID. Automatic reconnection after cluster resource failovers or migrations occurs properly after this one-time reset.

12.15 Expanding the Size of a Clustered Pool

NSS pools are made up of segments of space from one or more devices up to a maximum total size of 8 TB (terabytes). Each device can be up to 2 TB in size if a DOS partitioning scheme is used, or up to 8 TB in size if a GPT partitioning scheme is used. You can expand a pool by adding more space from the same device or from a different shared device. If you extend the size of a LUN device that is currently being used by the pool, you must make the extended size visible to all nodes in the cluster before you add the newly available space to the clustered pool. If you add a new shared LUN device, you must make the device visible to all nodes in the cluster before you add space from the device to the clustered pool.

12.15.1 Planning to Expand a Clustered Pool

Consider the guidelines in this section when you plan to expand the size of a clustered NSS pool.

Free Space on a Shared Device

To expand the size of a pool, you need free unpartitioned space on the same shared device or on other shared devices that can be failed over to the same node in a cluster. To extend the size of an existing LUN, the free unpartitioned space must immediately follow the LUN on the same shared device.

Extending a LUN Might Require a Service Outage

Novell Cluster Services and NSS allow you to perform an online extension of a LUN that is used by a clustered pool without taking the resource offline or stopping Cluster Services. However, not all SAN storage arrays and vendor-specific storage drivers fully support transparent online extension of

LUNs. Refer to the third-party vendor documentation for information about how to extend a LUN. We recommend that you test the specific hardware environment and hardware configuration to confirm that it behaves correctly before you perform an online LUN extension.

IMPORTANT: To prevent any chance of data loss or corruption, we recommend that you back up the volumes before you expand a LUN.

After you extend the LUN size on the SAN storage array, you must scan for devices on each node to update the node's storage map with its new size information before you expand the pool size on the active node, or allow the pool cluster resource to fail over to other nodes. If the device has multiple I/O paths, you must also update each node's multipath map. Depending on your SAN storage array, device driver, and server hardware, a server restart might be required to force the node to recognize the extended LUN size.

Adding a LUN Might Require a Service Outage

Novell Cluster Services and NSS allow you to add a LUN to a clustered pool without taking the resource offline or stopping Cluster Services. The LUN must be able to be failed over to the same node in a cluster as the pool cluster resource. Refer to the third-party vendor documentation for information about how to add a LUN.

After you create a new LUN on the SAN storage array and assign it to all nodes, you must scan for devices on each node to update the node's storage map with the new device before you expand the pool size on the active node, or allow the pool cluster resource to fail over to other nodes. If the device has multiple I/O paths, you must also update each node's multipath map. Depending on your SAN storage array, device driver, and server hardware, a server restart might be required to force the node to recognize the new LUN.

Multipath I/O Devices

Consider the following requirements and guidelines as you work with devices that have multiple I/O paths:

- The procedures in this section assume that you use the Linux Device Mapper - Multipath I/O (DM-MPIO) software to manage the multiple I/O paths for devices. You must modify the instructions accordingly if you are using a third-party multipath solution.

- Before you expand the pool size on the active node or allow the pool cluster resource to fail over to other nodes, you must update the multipath map on each node in the cluster so that DM-MPIO recognizes the LUN changes:

    - **Extend the LUN size:** On each node, you can restart the `multipathd` daemon, or you can use the `multipathd` command to resize the multipath map for the LUN.

    - **New LUN:** On each node, use the `multipath -v2` command to rebuild the node's multipath map.

- For a device with multiple I/O paths, use the device's multipath device name (such as `mpathb`) or device ID.

- Steps in this section that pertain only to multipath I/O devices are preceded by the keyword "MPIO".

12.15.2 Increasing the Size of a Clustered Pool by Extending the Size of Its LUN

One way to increase the size of a clustered pool is to add space to its LUN device on the SAN storage array, and then increase the size of the pool to use the newly available space. Before you extend the size of a LUN, ensure that you understand the cautions in "Extending a LUN Might Require a Service Outage" on page 202.

- "Extending the Size of a LUN" on page 204
- "Increasing the Pool Size by Using Free Space on the Same LUN" on page 206

Extending the Size of a LUN

You can extend the size of the LUN if there is free unpartitioned space on the device that immediately follows the LUN. If the LUN has multiple I/O paths, special handling is required to update the device size reported in the multipath map on each node. These steps are marked with the keyword "MPIO".

1 Identify the device that is used by the pool cluster resource:

  1a In iManager, go to *Storage > Pools*, then select the Cluster object of the cluster.

  1b In the Pools list, select the pool, then view the device name of the LUN under the pool *Details* area.

    For example, a device with a single path has a device node name such as `sdb`. A device with multiple I/O paths has a multipath device name such as `mpathb`.

2 (Offline LUN extension only) If the SAN storage array or vendor-specific storage driver does not support online LUN extension, take the pool cluster resource offline. Open a terminal as the `root` user, then at the command prompt, enter:

```
cluster offline <resource_name>
```

3 On the SAN storage array, use the third-party vendor tools to extend the size of the existing LUN.

Devices can be up to 2 TB in size with the DOS partitioning scheme. Devices can be up to 8 TB in size with the GPT partitioning scheme.

For example, increase the size of the existing LUN from 100 GB to 150 GB.

4 Beginning with the active node, perform the following steps on each node in turn in order to update the node's storage map with the new LUN size:

  4a Log in to the node as the `root` user, then open a terminal console.

  4b Scan for devices to recognize the new size for the LUN. At the command prompt, enter:

```
/bin/rescan-scsi-bus.sh -forcerescan [--luns=XX]
```

Use the `--luns=XX` option if you want to scan only for the extended LUN. You can use the `lsscsi(8)` command to locate the LUN number for the device. Information for each device is output on a single line. The first entry is four numbers separated by colons that represent the `Host:Bus:Target:LUN` of the SCSI device.

For information about other `rescan-scsi-bus.sh` options, see "Scanning for New Devices without Rebooting" (http://www.suse.com/documentation/sles11/stor_admin/data/scandev.html) in the *SUSE Linux Enterprise Server 11 Storage Administration Guide* (http://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html).

4c (MPIO) If the device has multiple I/O paths, do one of the following methods to update the multipath map with the new size information for the extended LUN:

- ◆ **Flush and rebuild the multipath maps:** At the command prompt, enter the following commands:

  ```
  multipath -F
  ```

  ```
  multipath -v2
  ```

- ◆ **Restart the multipathd daemon:** At the command prompt, enter one of the following commands:

  ```
  rcmultipathd restart
  ```

  ```
  /etc/multipathd restart
  ```

4d (MPIO) Verify that the new size of the LUN is correctly reported to DM-MPIO. At the command prompt, enter:

```
multipath -ll
```

The `-ll` option shows the current multipath topology from all available information (sysfs, the device mapper, path checkers, and so on).

For example, the `mpathb` device reports a new size of 150 GB.

```
mpathb (36001438005de9b8d0000800007520000) dm-1 HP,HSV450
[size=150G][features=1 queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=100][active]
 \_ 4:0:1:7 sdae       65:224 [active][ready]
 \_ 2:0:1:7 sdu        65:64  [active][ready]
\_ round-robin 0 [prio=20][enabled]
 \_ 4:0:0:7 sdv        65:80  [active][ready]
 \_ 2:0:0:7 sdg        8:96   [active][ready]
```

4e Scan for EVMS storage object changes. At the command prompt, enter:

```
evms_activate
```

4f Repeat Step 4a through Step 4e on each node in the cluster.

Depending on your SAN storage array, device driver, and server hardware, a server restart might be required to force a node to recognize the extended LUN size.

5 In iManager, verify that the new device size is reported correctly to NSS:

5a In iManager, go to *Storage > Devices*, then select the Cluster object of the cluster.

5b Select the device (such as mpathb), then view the device capacity under the device's *Details* area.

For example, the mpathb device's *Capacity* field reports a size of 150 GB.
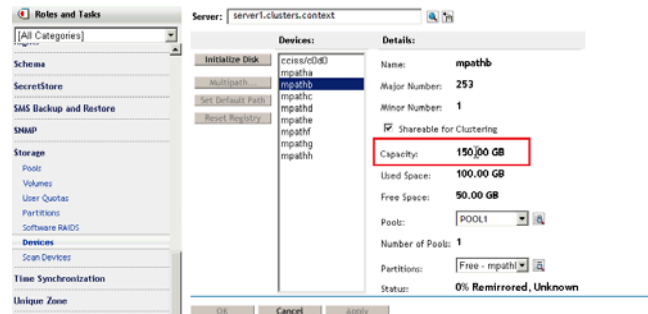
5c If the new LUN size is reported correctly, continue to Step 6. If the size is not reported correctly, repeat Step 4, then check again.

6 (Offline LUN extension only) If you took the pool cluster resource offline in Step 2, you can now bring pool cluster resource online. Open a terminal as the root user, then at the command prompt, enter:

```
cluster online <resource_name>
```

7 After the new LUN size has been recognized by all nodes, continue with "Increasing the Pool Size by Using Free Space on the Same LUN" on page 206.

Increasing the Pool Size by Using Free Space on the Same LUN

The pool size is not automatically expanded to use the new free space on the LUN. After you have expanded the size of the existing LUN and the new size is recognized by all nodes in the cluster, you can increase the size of the pool by allocating space from the newly available free space.

1 In iManager, go to *Storage > Pools*, then select the Cluster object of the cluster.

2 On the Pool page, select the pool, then view the pool's current size in the *Pool Details* area.

The *Total Space* for the pool reports only the amount of space that has been allocated to the pool. The pool size is not automatically expanded to use the new free space on the LUN.
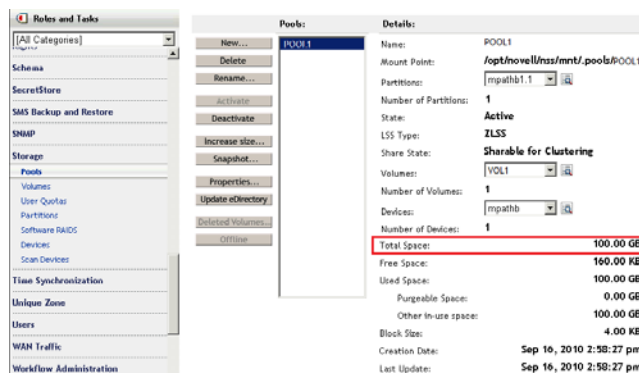
3 Increase the size of the pool by adding space from the same LUN:

    3a On the Pool page, select the pool, then click *Increase Size*.

        The Expand a Pool wizard opens and presents all devices with available free space.

    3b In the list of devices, select the check box for the newly expanded LUN, such as `mpathb`, then specify the amount of space to add in MB.

        Typically, this is the reported amount of free space available. For example:

	Used Size (MB)	Device Name	Free Size (MB)
☐	0	mpatha	1015
☑	51201	mpathb	51201
☐	0	mpathc	102399
☐	0	mpathd	20479
☐	0	mpathe	20479
☐	0	mpathf	20479
☐	0	mpathg	20479
☐	0	mpathh	20479

Pool Size: **153599**

    3c Click *Finish* to save and apply the change.

4 On the Pool page, select the pool, then verify that the *Total Space* field reports the new size.

For example, the new size of `POOL1` is 150 GB.

5 Go to *Storage > Partitions*, then verify that a new partition exists for the LUN and is assigned to the pool.

For example, the partitions `mpathb1.1` and `mpathb1.2` are allocated to `POOL1`.

Server: server1.clusters.context

Partitions
New | Edit | Delete | Details

	Name ▾	Type ▾	Status ▾	Device Name ▾	RAID Name ▾	Pool Name ▾	Size
☐	cciss/c0d0_freespace1	Free Space	Available	cciss/c0d0			27.07 MB
☐	cciss/c0d0p1	Linux	In Use	cciss/c0d0			509.84 MB
☐	cciss/c0d0p2	Linux Swap	In Use	cciss/c0d0			15.99 GB
☐	cciss/c0d0p3	Linux	In Use	cciss/c0d0			256.88 GB
☐	cluster1.sbd	Clustering	In Use	mpatha			7.98 MB
☐	mpatha1.nwfreespace1	Free Space	Available	mpatha			7.66 MB
☐	mpatha_freespace1	Free Space	Available	mpatha			1008.31 MB
☐	mpathb1.1	NSS	In Use	mpathb		POOL1	100.00 GB
☐	mpathb1.2	NSS	In Use	mpathb		POOL1	50.00 GB
☐	mpathb_freespace1	Free Space	Available	mpathb			2.34 MB

6 For each volume on the pool, if the volume quota is set to *No Quota*, the volume can grow to the size of the pool. Go to *Storage > Volumes*, select the volume, then verify that its *Available Space* field has increased by the amount of space that was added to the pool.

For example, by increasing the pool size by 50 GB, the available space for VOL1 increased from about 82 GB to about 132 GB.
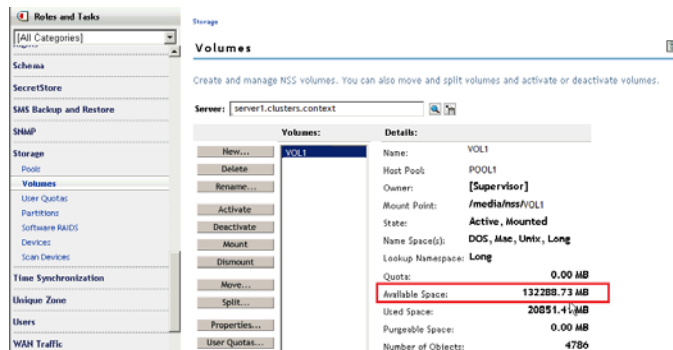
7 Continue with .

12.15.3 Increasing the Size of a Clustered Pool by Adding Space from a Different Shared LUN

You can increase the size of a clustered pool by adding space to the pool from a different shared LUN device that can be failed over with the resource. Before you add a LUN, ensure that you understand the cautions in "Adding a LUN Might Require a Service Outage" on page 203.

- ◆ "Creating and Sharing a New LUN Device" on page 208
- ◆ "Increasing the Pool Size by Adding Free Space from a Different LUN" on page 210

Creating and Sharing a New LUN Device

If you do not have a shared device that can be failed over with the resource, you can create a new LUN and share it with all nodes in the cluster. If the LUN has multiple I/O paths, special handling is required to add the device to the multipath map on each node. These steps are marked with the keyword "MPIO".

1 Use third-party tools for the SAN storage array to create a new shared LUN device, and assign it to each cluster node.

2 Beginning with the active node, perform the following steps on each node in turn in order to update the node's storage map and multipath map with the new device:

2a Log in to the node as the root user, then open a terminal console.

2b Scan for devices to help the OS recognize the new LUN. At the command prompt, enter:

```
/bin/rescan-scsi-bus.sh -forcerescan
```

For information about other rescan-scsi-bus.sh options, see "Scanning for New Devices without Rebooting" (http://www.suse.com/documentation/sles11/stor_admin/data/scandev.html) in the *SUSE Linux Enterprise Server 11 Storage Administration Guide* (http://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html).

2c Use the `lsscsi` command to verify that the LUN is seen by the OS. At the command prompt, enter:

```
lsscsi
```

Information for each device is output on a single line. The first entry is four numbers separated by colons that represent the `Host:Bus:Target:LUN` of the SCSI device. In the sample output below, the `/dev/sdc` and `/dev/sde` devices are two individual paths for LUN 2.

```
# lsscsi
[4:0:0:0]    disk    DGC        RAID 5            0216  /dev/sdb
[4:0:0:1]    disk    DGC        RAID 5            0216  /dev/sdf
[4:0:0:2]    disk    DGC        RAID 5            0216  /dev/sdc
[4:0:1:0]    disk    DGC        RAID 5            0216  /dev/sdd
[4:0:1:1]    disk    DGC        RAID 5            0216  /dev/sdg
[4:0:1:2]    disk    DGC        RAID 5            0216  /dev/sde
```

If the LUN is not seen, repeat this step using different `lsscsi` command line parameters to scan for all devices (such as `-w -c -l`). For information about command options, see the `lsscsi(8)` man page.

If the LUN is still not seen by the OS, a server restart might be required.

2d Scan for EVMS storage object changes. At the command prompt, enter:

```
evms_activate
```

2e (MPIO) If the device has multiple I/O paths, rebuild the multipath map. At the command prompt, enter:

```
multipath -v2
```

2f (MPIO) Verify that the new LUN is reported to DM-MPIO. At the command prompt, enter:

```
multipath -ll
```

The `-ll` option shows the current multipath topology from all available information (`sysfs`, the device mapper, path checkers, and so on).

2g Verify that the LUN is now visible to NSS on the node. In NSSMU, go to the Devices page, then verify that the new LUN is listed.

For example, the new device is `mpathj` on the active node. Remember that a multipath device might have a different multipath device name on each node.

```
                              NSS Management Utility
 NSS Version 4.12a   Build 3018
      Devices                              Device Information
    cciss/c0d0
    mpatha             Name: mpathj
    mpathb             ObjectID:  mpathj
    mpathc             Removable: No
    mpathd
    mpathe             Capacity:           25.00 GB
    mpathf             Partitioned Space:   0.00 KB
    mpathg             Unpartitioned Space: 0.00 KB
    mpathh
  * mpathj             Sharable for Clustering: No
```
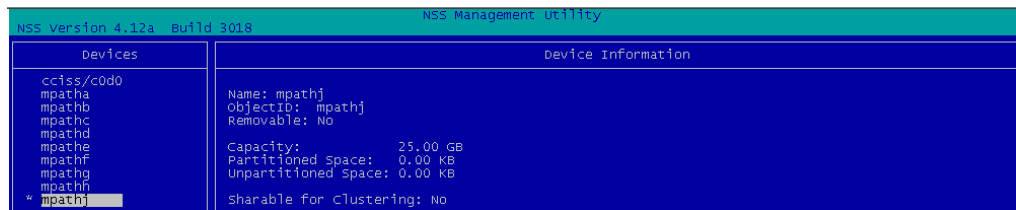
2h Repeat Step 2a through Step 2g on each node in the cluster.

Depending on your SAN storage array, device driver, and server hardware, a server restart might be required to force a node to recognize the new LUN.

3 In iManager, initialize the new LUN, then mark it as shareable for clustering:

3a In iManager, go to *Storage > Devices.*

3b On the Devices page, select the Cluster object of the cluster.

Selecting the Cluster object connects you to the current master node of the cluster.

Configuring Cluster Resources for Shared NSS Pools and Volumes **209**

3c In the Devices list, select the newly added LUN device.

> **WARNING:** Ensure that you select the correct device. Initializing a device destroys all data on it.

3d Click *Initialize*, then click *OK* to confirm that you want to initialize the device.

If you are prompted for the partitioning scheme, specify your preferred partitioning scheme as DOS (up to 2 TB in size) or GPT (up to 8 TB in size).

3e In the Devices list, select the new LUN device.

3f In the *Details* area, select the *Shareable for Clustering* check box, then click *Apply*.
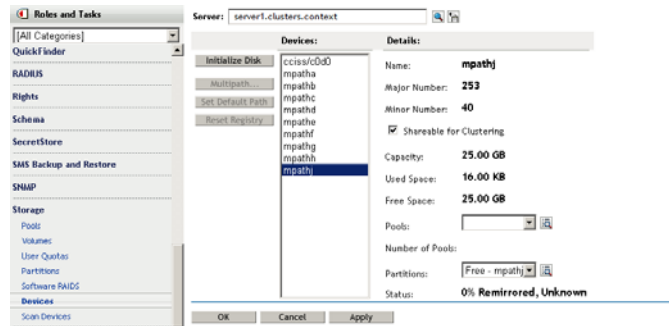
4 After the new LUN is recognized by all nodes, and has been initialized and shared, continue with "Increasing the Pool Size by Adding Free Space from a Different LUN" on page 210.

Increasing the Pool Size by Adding Free Space from a Different LUN

After you have set up a different shared device that can be failed over with the pool cluster resource, you can increase the size of the pool by allocating space from it.

1 In iManager, go to *Storage > Pools*, then select the Cluster object of the cluster.

Selecting the Cluster object automatically connects you to the current master node of a cluster.

2 Select the pool, then view the pool's current size.

The *Total Space* for the pool reports only the amount of space that has been allocated to the pool. For example, Pool7 is 20 GB in size.
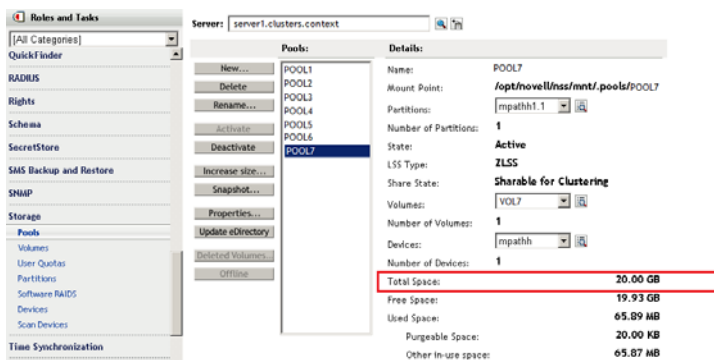
3 Increase the size of the pool by adding space from a different LUN:

3a On the Pool page, select the pool, then click *Increase Size*.

The Expand a Pool wizard opens and presents all devices with available free space.

210 OES 2 SP3: Novell Cluster Services 1.8.8 Administration Guide for Linux

3b In the list of devices, select the check box for the new LUN, such as `mpathj`, then specify the amount of space to add in MB.

Typically, this is the reported amount of free space available. For example, `mpathj` has 25599 MB available:

	Used Size (MB)	Device Name	Free Size (MB)
☐	0	mpatha	1015
☐	0	mpathb	2
☐	0	mpathc	1
☑	25599	mpathj	25599

Pool Size: **46077**

3c Click *Finish* to save and apply the change.

4 On the Pool page, select the pool, then verify that the *Total Space* field reports the new size.

For example, `POOL7` has increased in size from 20 GB to about 45 GB.

5 For each volume on the pool, if the volume quota is set to *No Quota*, the volume can grow to the size of the pool. Go to *Storage > Volumes*, select the volume, then verify that a volume's *Available Space* field has increased by the amount of space that was added to the pool.

6 After the pool size is increased, continue with Section 12.15.4, "Verifying the Expanded Pool Cluster Resource on All Nodes," on page 211.

12.15.4 Verifying the Expanded Pool Cluster Resource on All Nodes

Verify that the expanded pool size is reported correctly on each node in the cluster.

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

2 Cluster migrate the pool cluster resource to each node in turn, and use NSSMU to verify that the devices, partitions, pool, and volumes show the correct information on that server:

2a Cluster migrate the pool cluster resource.

```
cluster migrate resource_name node_name
```

2b Launch NSSMU by entering

```
nssmu
```

2c On the Devices page, select the device and verify that the new size is reported.

2d On the Partitions page, select each partition for the device, then verify that the partitions are shown for the expanded pool.

2e On the Pools page, select the pool, then verify the total space available.

2f On the Volumes page, select each volume that has no quota set, then verify that the available space has increased by the amount of space added to its pool.

2g If the information is incorrect, wait a few minutes to allow time for the system to recognize the new partitions, then try again.

3 Repeat the previous validation for each node in the cluster.

12.16 Deleting NSS Pool Cluster Resources

Ensure that you offline the cluster resource before attempting to delete either the cluster resource or the clustered pool. For example, if you want to unshare a pool, offline the cluster resource for the pool before you mark the pool or the device as Not Shareable for Clustering, then you can delete the eDirectory object for the cluster resource.

WARNING: If you attempt to delete a cluster resource without first offlining it, deletion errors occur, and the data associated with the clustered pool is not recoverable.

All resource configuration must happen from the master node. On the Cluster Options page for Novell iManager, connect to the Cluster object, not to Cluster Node objects. On the *Storage > Pools* page for iManager, connect to the master node. Run NSSMU only on the master node.

For information, see Section 10.14, "Deleting Cluster Resources," on page 167.

13 Configuring and Managing Cluster Resources for Shared Linux POSIX Volumes

After you have installed and configured Novell Cluster Services on a server, you can configure storage resources. This section describes how to create and cluster-enable shared Linux POSIX volumes as cluster resources with Novell Cluster Services.

IMPORTANT: For information about configuring NCP volumes on Linux POSIX cluster resources, see "Configuring NCP Volumes with Novell Cluster Services" in the *OES 2 SP3: NCP Server for Linux Administration Guide*.

- Section 13.1, "Requirements for Creating Linux POSIX Volume Cluster Resources," on page 213
- Section 13.2, "Gathering Information for Clustering a Linux POSIX Volume," on page 215
- Section 13.3, "Creating Linux POSIX Volumes on Shared Devices," on page 216
- Section 13.4, "Cluster-Enabling a Linux POSIX Volume on a Shared Device," on page 221
- Section 13.5, "Verifying the Load and Unload Scripts," on page 228
- Section 13.6, "Creating a Virtual Server Object for the Linux POSIX Volume Cluster Resource," on page 228
- Section 13.7, "Sample Scripts for a Linux POSIX Volume Cluster Resource," on page 231
- Section 13.8, "Expanding EVMS Volumes on Shared Disks," on page 234
- Section 13.9, "Deleting Shared Storage," on page 235
- Section 13.10, "Known Issues for Working with Cluster Resources for Linux POSIX Volumes," on page 235

13.1 Requirements for Creating Linux POSIX Volume Cluster Resources

Your system must meet the requirements in this section in addition to the cluster requirements described in Chapter 4, "Planning for Novell Cluster Services," on page 55.

- Section 13.1.1, "Novell Cluster Services," on page 214
- Section 13.1.2, "EVMS," on page 214
- Section 13.1.3, "Cluster Segment Manager," on page 214
- Section 13.1.4, "Shareable Device," on page 214
- Section 13.1.5, "Journaled Linux POSIX File Systems," on page 215
- Section 13.1.6, "Resource IP Address," on page 215

13.1.1 Novell Cluster Services

Novell Cluster Services must be installed and running when you create and manage the shared Linux POSIX volume. The cluster must be active.

13.1.2 EVMS

The EVMS (Enterprise Volume Management System) is installed automatically when you install Novell Cluster Services. When you configure Novell Cluster Services on a server, the Cluster Segment Manager (CSM) plug-in for EVMS is installed in EVMS. This setup is required to be completed before you can configure Linux POSIX volumes for use with Novell Cluster Services.

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

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

For information about using EVMS, see "Using EVMS to Manage Devices" (http://www.suse.com/documentation/sles10/stor_admin/data/evmsconfig.html) in the *SLES 10 SP3/SP4 Storage Administration Guide* (http://www.suse.com/documentation/sles10/stor_admin/data/bookinfo.html), or see the EVMS man page by entering `man evms` at a terminal console.

13.1.3 Cluster Segment Manager

Novell Cluster Services uses the Cluster Segment Manager (CSM) plug-in to EVMS in order to manage shared devices. The CSM allows only one node at a time to access a shared cluster resource. This helps prevent the potential data corruption that can occur if multiple nodes concurrently access the same data.

The CSM plug-in to EVMS is required to create a container on the shared device where you will create the Linux POSIX volume.

The Cluster Segment Manager requires use of the entire disk (or LUN).

CSM containers require Novell Cluster Services to be running on all nodes that access the cluster resources on OES 2 Linux.

IMPORTANT: Do not create or make to modifications to any shared EVMS objects unless Novell Cluster Services is running.

13.1.4 Shareable Device

You should carefully plan how you want to configure your shared storage prior to configuring the Linux POSIX cluster resource. For planning information, see the following:

For the Linux POSIX volume, you need an unpartitioned disk or LUN that is connected via Fibre Channel or iSCSI protocols to the OES 2 Linux server. The disk or LUN must be able to be managed by EVMS.

You must install and configure Novell Cluster Services on a server before you can configure the storage device as shareable for clustering.

13.1.5 Journaled Linux POSIX File Systems

Novell Cluster Services supports cluster-enabling EVMS partitions and volumes for any journaled Linux POSIX file systems (such as Ext3, ReiserFS, and XFS).

13.1.6 Resource IP Address

Each Linux POSIX cluster resource needs a unique static IP address that is in the same subnet as the IP addresses that are used for the cluster and cluster nodes. The IP address is used to provide access and failover capability for the cluster-enabled volume.

13.1.7 Mixed-Node Clusters

A mixed-node cluster is supported as a temporary configuration scenario for converting a cluster from NetWare 6.5 SP8 to OES 2 Linux or for upgrading a cluster from OES 1 Linux to OES 2 Linux. You can create new Linux POSIX cluster resources on the OES 2 Linux nodes in order to support the migration of services and data from the old nodes to the new OES 2 Linux (or later) nodes. However, the new cluster resources should be cluster migrated or failed over only between the OES 2 Linux nodes.

13.1.8 Novell Samba

Novell Samba is supported for Linux POSIX volume cluster resources. It can be set up after you cluster-enable the volume. For information about setting up Novell Samba for storage resources, see "Configuring Samba for Novell Cluster Services" in the *OES2 SP3: Samba Administration Guide*.

13.2 Gathering Information for Clustering a Linux POSIX Volume

Gather the information that you will use as you follow the steps to cluster a Linux POSIX volume.

IMPORTANT: On Linux, all names are case sensitive in the tools and cluster scripts.

Cluster Information	Example	Caveats
cluster name	cluster1	Name of the cluster.
cluster context	ou=clusters,ou=city,o=mycompany	Context where you created the cluster.
RESOURCE_IP	10.10.10.44	IP address of the Linux POSIX cluster resource.
container_name	csm44	Name of the EVMS Cluster Segment Manager (CSM) container.

Cluster Information	Example	Caveats
evms_volumename	`lxvol44`	Name of the Linux POSIX volume that you create on the CSM.

MOUNT_DEV	`/dev/evms/`
	`container_name/`
	`evms_volumename`
		The Linux path for the EVMS volume. For example:
		
		`/dev/evms/csm44/lxvol44`

MOUNT_FS	ext3 resiserfs xfs	The file system type you specify when you mount the volume.

MOUNT_POINT	`/path_to_mount_point`	The mount location for the EVMS volume. You can mount the EVMS volume anywhere. It can have the same or different name than the underlying EVMS volume. For example:
		
		`/mnt/lxvol44`
		`/mnt/users`
		`/home`

13.3 Creating Linux POSIX Volumes on Shared Devices

Using EVMS, you can add the Cluster Segment Manager to a shared device and create a Linux volume and file system on it for use with Novell Cluster Services.

Complete the tasks in this section to create a Linux POSIX volume and file system on a shared device:

13.3.1 Initializing and Sharing a Device

After Novell Cluster Services is configured on the server, you can mark the device as shareable for clustering. If the device is newly added to the server, it might also be necessary to initialize the device in order to get the system to recognize its free space.

You can perform the device management tasks by using the Novell Storage Services (NSS) Management Utility or the Storage plug-in to iManager.

Initializing and Sharing a Disk with NSSMU

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

2 At the console prompt, enter

```
nssmu
```

3 In the NSSMU, select *Devices* and press Enter.

4 In the *Devices* list, select the device where you want to create the Linux POSIX cluster resource.

5 If the device has not been initialized, press F3 to initialize the disk.

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

This initializes the disk and lays down a DOS Segment Manager on it. You will remove this segment manager later.

6 Press F6 to mark the device as shareable for clustering.

The *Shareable for Clustering* value changes from *No* to *Yes*.

7 Exit NSSMU.

8 Continue with Section 13.3.2, "Removing Formatting and Segment Managers," on page 217.

Initializing and Sharing a Disk with iManager

1 In iManager, select *Storage > Devices*, then select a server.

2 From the *Devices* list, select the device where you want to create the Linux POSIX cluster resource.

3 If the device has not been initialized, click *Initialize Disk.*

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

This initializes the disk and lays down a DOS Segment Manager on it. You will remove this segment manager later.

4 In the *Details* area, select the *Shareable for Clustering* check box, then click *OK* to apply the change.

5 Exit iManager.

6 Continue with Section 13.3.2, "Removing Formatting and Segment Managers," on page 217.

13.3.2 Removing Formatting and Segment Managers

The EVMS Cluster Segment Manager must be the first segment manager laid down on the device that you want to use for the shared Linux volume. Before you can assign the CSM to a device, you must Ensure that the space is free of other formatting and segment managers.

WARNING: Clearing the existing formatting for a device destroys all data on it.

EVMS by default shows unformatted free space on disks as compatibility volumes. CMS requires the entire device. You must delete any existing compatibility volumes in order to free the space.

If the device contains existing segment managers, you must also delete them so that the CSM container is the first one on the device.

1 At a Linux terminal console, log in as the root user, then enter

```
evmsgui
```

2 Locate and delete any compatibility volumes on the device that you want to use for the shared volume:

2a Click the *Volumes* tab, then locate any compatibility volumes that are on the space you want to use.

2b Right-click the volume of interest, select *Display details*, click the *Page 2* tab, verify from the *Status* field that the volume is a compatibility volume, then click *OK* to dismiss the dialog box.

2c If the volume is a compatibility volume, continue with Step 2d. Otherwise, go to Step 3.

2d Click the *Volumes* tab, right-click the volume, then select *Delete*.

> **WARNING:** All data on the selected volume is destroyed.

2e Select the volume, then click *Recursive Delete*.

2f If a *Response Required* pop-up appears, click the *Write zeros* button.

2g If another pop-up appears, click *Continue* to write 1024 bytes to the end of the volume.

2h Click *Save*, then click *Save* again to save your changes.

3 Locate and delete any existing segment managers on the device that you want to use for the shared volume, such as NetWare Segment Manager or DOS Segment Manager:

3a Click the *Disks* tab, then locate and select the device that you want to use for the shared volume, such as device `sdb`.

3b Right-click, then select *Remove segment manager from Object*.

This option appears only if there is an existing segment manager for the selected disk.

3c Select the listed non-CSM segment manager, click *Remove*, then click *OK*.

> **WARNING:** All data on the selected disk space is destroyed.

3d Click *Save*, then click *Save* again to save your changes.

3e Repeat Step 3a to Step 3d until the *Remove segment manager from Object* option is no longer available when you right-click the device.

4 Continue with Section 13.3.3, "Creating a Cluster Segment Manager Container," on page 218.

13.3.3 Creating a Cluster Segment Manager Container

A Cluster Segment Manager container provides exclusive access to shared storage for a node. You must assign an EVMS CSM to the shared device where you want to create the shared volume. The EVMS Cluster Segment Manager must be the first segment manager laid down on the space you want to use for the shared Linux volume. You use the Cluster Segment Manager plug-in for EVMS to create a CSM container.

> **IMPORTANT:** CSM containers require Novell Cluster Services to be running on all nodes that access the CSM container. Ensure that Novell Cluster Services is running before you continue.

1 In evmsgui, click *Actions > Create > Container*.

2 On the Create Storage Container page, select *Cluster Segment Manager*, then click *Next*.

3 On the Select Plugin Acceptable Objects page, select the disks (storage objects) you want to place in the container (such as `sdb`), then click *Next*.

4 On the Configuration Options page, select the cluster server node where you are creating the container, specify *Private* as the type, then specify a name for the container.

Specify a name that is easy to associate with the cluster resource you plan to create for the container. For example, if the IP address you plan to assign to the resource is 10.10.10.44, you might name the container `csm44`.

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

Container
Disk
EVMS
Plugin
Region
Segment
Volume

5 Click *Create*, then click *OK*.

6 Click *Save*, then click *Save* again to save your changes.

7 Continue with Section 13.3.4, "Adding a Segment Manager to the CSM Container," on page 219.

13.3.4 Adding a Segment Manager to the CSM Container

As a best practice, a cluster resource uses only one volume per shared device. However, if you need multiple volumes on the same device, you should add a segment manager to the CSM container.

A CSM container uses the entire EVMS disk. Without segment managers on top of the CSM container, creating additional volumes or expanding or shrinking volumes is not possible. Also, because only one EVMS volume is possible in the container, only one file system type is used.

After creating a CSM container, you can optionally add a non-CSM segment manager container on top of it. This allows you to create one volume or multiple smaller volumes on your shared device. You can expand or shrink existing EVMS volumes to utilize the space on your shared device. In addition, this means that you can also have different file system types on your shared device.

IMPORTANT: If you create multiple volumes, they must be managed in the same cluster resource because the device moves from node to node as a single entity.

Use the following procedure to add a non-CSM segment manager on the CSM container, or skip ahead to Section 13.3.5, "Creating an EVMS Volume," on page 220 if you do not want to add another segment manager.

1 In evmsgui, click *Actions > Add > Segment Manager to Storage Object*.

2 On the Add Segment Manager to Storage Object page, choose the desired segment manager (such as *DOS Segment Manager*), then click *Next*.

Most of the segment managers will work. The DOS segment manager is added by default for some EVMS operations.

3 On the Select Plugin Acceptable Objects page, choose the CSM container storage object where you want to add the segment manager, then click *Next*.

4 On the Configurable Options page, select the disk type (Linux is the default), click *Add*, then click *OK*.

5 Click *Save*, then click *Save* again to save your changes.

6 If you added a DOS segment manager, create a segment for it:

> **IMPORTANT:** Some segment managers such as the DOS segment manager require you to create a segment before creating an EVMS volume. Without a segment, the additional segment manager does not appear when you attempt to create an EVMS volume.

    6a In evmsgui, click *Actions > Create > Segment*.

    6b On the Create Disk Segment page, select *DOS Segment Manager*, then click *Next*.

    6c On the Select Plugin Acceptable Objects page, choose the CSM container storage object (such as `csm44/sdb_freespace1`) where you want to add the segment, then click *Next*.

    6d Specify the size of the segment, the partition type (such as Linux LVM), click *Create*, then click *OK*.

    6e Click *Save*, then click *Save* again to save your changes.

7 Continue with

13.3.5 Creating an EVMS Volume

1 In evmsgui, click *Actions > Create > EVMS Volume*.

2 On the Create EVMS Volume page, select the container you just created (either the CSM container or the additional segment manager container), then specify a name for the volume (such as `lxvol44`).

3 Click *Create*, then click *OK*.

4 Click *Save*, then click *Save* again to save your changes.

5 Click the *Volumes* tab to verify that the EVMS volume was created.

For example, a volume named `lxvol44` would be listed as `/dev/evms/csm44/lxvol44`.

6 Continue with

13.3.6 Making a File System on the EVMS Volume

1 In evmsgui, click the *Disks* tab, then activate the CSM container:

    1a On the Disks page, right-click the CSM container, then select *Activate*.

    1b On the Activate page, select the CSM container, click *Activate*, then click *OK*.

    1c Click *Save*, then click *Save* again to save your changes.

2 Click the *Volumes* tab, then activate the EVMS volume:

    2a On the Volumes page, right-click the EVMS volume, then select *Activate*.

    2b On the Activate page, select the volume, click *Activate*, then click *OK*.

    2c Click *Save*, then click *Save* again to save your changes.

3 Make the file system on the EVMS volume:

    3a On the Volumes page, right-click the volume, then select *Make File System*.

    3b On the Make File System page, choose a Linux POSIX file system interface module from the list (such as *Ext2/3 File System Interface Module*), then click *Next*.

    3c Specify a volume label, click *Make*, then click *OK*.

    3d Click *Save*, then click *Save* again to save your changes.

The file system type is now listed under the *Plugin* column.

4 Mount the volume:

   4a On the Volumes page, right-click the volume, then select *Mount*.

   4b On the Mount File System page, select the volume, then specify the Linux path to use for the mount point, such as `/mnt/mount_point`.

   4c (Optional) Click *Options*, specify any desired mount options, then click *OK*.

   4d Click *Mount*, then click *OK*.

   The mount point now appears in the *Mount Point* column.

5 Continue with Section 13.4, "Cluster-Enabling a Linux POSIX Volume on a Shared Device," on page 221.

13.4 Cluster-Enabling a Linux POSIX Volume on a Shared Device

After you have created the shared Linux POSIX volume, you cluster-enable it by creating a cluster resource for it. The cluster resource allows the shared disk and its contents to be moved or mounted on different servers in the cluster. The cluster resource provides a way for clients to automatically reconnect to the volume regardless of which server is hosting it.

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

Perform the tasks in this section to create and configure a new cluster resource. You can modify the cluster resource settings afterwards by using the *Clusters > Cluster Options > Properties* option in iManager.

- Section 13.4.1, "Logging in to iManager," on page 221
- Section 13.4.2, "Creating a Cluster Resource for a Linux POSIX Volume," on page 222
- Section 13.4.3, "Configuring a Load Script for the Linux POSIX Volume Cluster Resource," on page 223
- Section 13.4.4, "Configuring an Unload Script for a Linux POSIX Volume Cluster Resource," on page 224
- Section 13.4.5, "Enabling Monitoring and Configuring a Monitor Script for a Linux POSIX Volume Cluster Resource," on page 225
- Section 13.4.6, "Configuring Policies for a Linux POSIX Cluster Resource," on page 226

13.4.1 Logging in to iManager

When you create a new cluster resource, a Cluster Resource object is added in eDirectory. Log in to iManager as an administrator user with the credentials necessary to extend the schema for this object.

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

`http://`*server_ip_address*`/nps/iManager.html`
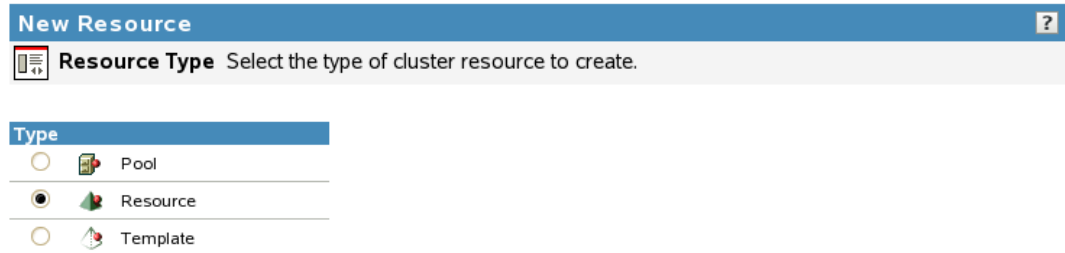
Replace *server_ip_address* with the IP address or DNS name of an OES 2 Linux server in the cluster that has iManager installed or with the IP address for Apache-based services.

2 Specify your user name and password.

3 Specify the tree where the cluster is located, the click *Login*.

4 Continue with "Creating a Cluster Resource for a Linux POSIX Volume" on page 222.

13.4.2 Creating a Cluster Resource for a Linux POSIX Volume

1 In iManager, click *Clusters*, then click *Cluster Options*.

2 Browse and select the Cluster object of the cluster where you want to add a cluster resource.

   When the page refreshes, a list of existing cluster resources is displayed.

3 Click *New*.

4 Specify Resource as the resource type you want to create by clicking the *Resource* radio button, then click *Next*.
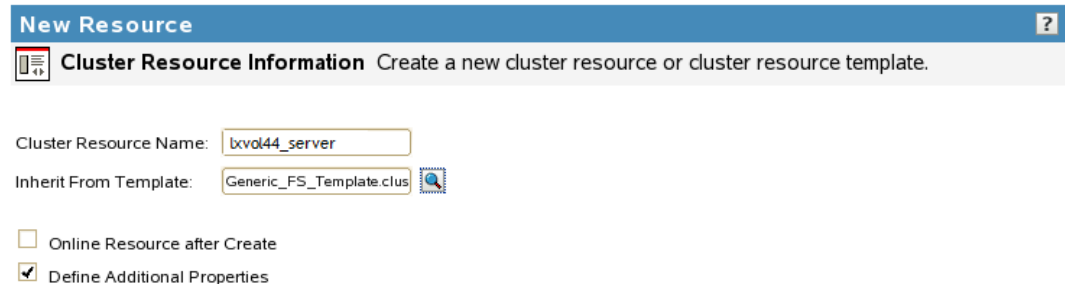
5 Specify the name of the resource you want to create.

   This is the name of the resource for the cluster-enabled Linux POSIX volume. For example, `lxvol44_server`.

   Do not use periods in cluster resource names. Novell clients interpret periods as delimiters. If you use a space in a cluster resource name, that space is converted to an underscore.

6 In the *Inherit From Template* field, browse to the Cluster container object, then locate and select the *Generic_FS_Template* in that container.

7 Deselect *Online on creation*.

   This option is deselected by default.

   Whenever you modify cluster resource settings, you must offline, then online the resource in order for the changes to take effect.

8 Select the *Define Additional Properties* check box, then click *Next*.

   This takes you through several pages where you can define the load, unload, and monitoring scripts, then configure the cluster resource properties. Each of these are described in later sections as you continue with the cluster resource configuration.

9 Continue with "Configuring a Load Script for the Linux POSIX Volume Cluster Resource" on page 223.

13.4.3 Configuring a Load Script for the Linux POSIX Volume Cluster Resource

The resource load script specifies the commands to start the resource (including mounting the file system) on a server in the cluster.

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

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

If you later need to modify the IP address, administrator credentials, or other attributes of an existing resource, follow the procedure in Section 8.9, "Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster," on page 116.

The generic file system template (selected in Step 6 of Section 13.4.2, "Creating a Cluster Resource for a Linux POSIX Volume," on page 222) contains a load script, and unload script, and a monitoring script that you must customize for your specific configuration.

Gather the following values for each of the scripts. Replace the sample values with actual values. Parameters and values in the script are case sensitive.

Resource Definition Information	Example Value
A unique static IP address that you want to assign to the cluster resource	10.10.10.44
The name of the Cluster Segment Manager container that you created in Section 13.3.3, "Creating a Cluster Segment Manager Container," on page 218	`csm44`
The Linux POSIX volume name that you created in Section 13.3.5, "Creating an EVMS Volume," on page 220	`lxvol44`
The file system type you made on the EVMS volume in Section 13.3.6, "Making a File System on the EVMS Volume," on page 220	`ext3`
The Linux path of the mount point where you want to mount the EVMS volume	`/mnt/lxvol44`

For an example of the customized load script, see Section 13.7.1, "Sample Load Script for the Linux POSIX Volume Cluster Resource," on page 232.

To configure the load script for new cluster resource:

1 On the Load Script page, view the definition section that was pulled from the Generic_FS_Template:

```
# 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
```

2 Modify the variable values in the definition lines with the values for your Linux POSIX volume.

For example:

```
# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
#define the container name
container_name=csm44
# define the device
MOUNT_DEV=/dev/evms/$container_name/lxvol44
# define the mount point
MOUNT_POINT=/mnt/lxvol44
```

3 (Optional) Modify the mount command to add mount options that are available for the file system.

Check the function `mount_fs` in the `/opt/novell/ncs/lib/ncsfuncs` file for additional mount options that are available for the file system you are using.

For example, the Ext3 file system supports mount options for access control lists (ACL) and extended attributes (USER_XATTR). You could add the following to the load script for a shared Ext3 volume to take advantage of those options:

```
#define mount options
MOUNT_OPTIONS=acl,user_xattr

# mount the file system
exit_on_error mount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS $MOUNT_OPTIONS
```

4 If the path to the mount point does not exist on other nodes, you can add the following lines in the script before the line to mount the file system:

```
# create the mount point path when loading on a new node
ignore_error mkdir -p $MOUNT_POINT
```

5 Below the script editing area, specify the *Load Script Timeout* value, then click *Next*.

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

6 Continue with Section 13.4.4, "Configuring an Unload Script for a Linux POSIX Volume Cluster Resource," on page 224.

13.4.4 Configuring an Unload Script for a Linux POSIX Volume Cluster Resource

The cluster resource unload script specifies the commands to stop the resource (including unmounting the file system) on a server in the cluster.

The generic file system template (selected in Step 6 of Section 13.4.2, "Creating a Cluster Resource for a Linux POSIX Volume," on page 222) contains a load script, and unload script, and a monitoring script that you must customize for your specific configuration. Use the same values for the defined variables that you used in the load script in Section 13.4.3, "Configuring a Load Script for the Linux POSIX Volume Cluster Resource," on page 223.

To configure the unload script for the new cluster resource:

1 On the Unload Script page, view the definition section that was pulled from the Generic_FS_Template:

```
# 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
```

2 Modify the variable values in the definition lines with the values for your Linux POSIX volume.

For example:

```
# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
#define the container name
container_name=csm44
# define the device
MOUNT_DEV=/dev/evms/$container_name/lxvol44
# define the mount point
MOUNT_POINT=/mnt/lxvol44
```

3 Below the script editing area, specify the *Unload Script Timeout* value, then click *Next*.

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

4 Continue with Section 13.4.5, "Enabling Monitoring and Configuring a Monitor Script for a Linux POSIX Volume Cluster Resource," on page 225.

13.4.5 Enabling Monitoring and Configuring a Monitor Script for a Linux POSIX Volume Cluster Resource

The cluster resource monitor script specifies the commands to monitor the cluster resource.

The generic file system template (selected in Step 6 of Section 13.4.2, "Creating a Cluster Resource for a Linux POSIX Volume," on page 222) contains a load script, and unload script, and a monitoring script that you must customize for your specific configuration. Use the same values for the defined variables that you used in the load script in Section 13.4.3, "Configuring a Load Script for the Linux POSIX Volume Cluster Resource," on page 223.

To configure the monitoring script for the new cluster resource:

1 On the Monitor Scripts page, select the *Enable Resource Monitoring* check box to enable resource monitoring for the selected resource.

Resource monitoring is disabled by default.

2 On the Monitoring Script page, view the definition section that was pulled from the Generic_FS_Template:

```
# 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
```

3 Modify the variable values in the definition lines with the values for your Linux POSIX volume.

For example:

```
# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
#define the container name
container_name=csm44
# define the device
MOUNT_DEV=/dev/evms/$container_name/lxvol44
# define the mount point
MOUNT_POINT=/mnt/lxvol44
```

4 Below the script editing area, specify the *Monitor Script Timeout* value, then click *Next*.

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

5 Continue with to configure the cluster resource behavior for the Linux POSIX volume cluster resource.

13.4.6 Configuring Policies for a Linux POSIX Cluster Resource

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

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

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

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

The preferred node is the first server in the *Assigned Nodes* list for the resource.

IMPORTANT: Resources fail back only to the first node in their *Assigned Nodes* list. So if a resource has failed over to three servers since it originally ran on its preferred node, and the second server the resource was running on comes back up, the resource will not failback to that second server.

Resources do not automatically move from node to node just because a node higher in the *Assigned Nodes* list rejoins the cluster, unless the *Failback* mode is set to AUTO and the first node in the *Assigned Nodes* list rejoins the cluster.

To configure cluster policies for the new cluster resource:

1 (Optional) Select the *Resource Follows Master* check box if you want to ensure that the resource runs only on the master node in the cluster.

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

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

The quorum default values were set when you installed Novell Cluster Services. You can change the quorum default values by accessing the properties page for the Cluster object.

Selecting this box ensures that the resource is launched immediately on any server in the Assigned Nodes list as soon as any server in the list is brought online.

3 Specify the *Start*, *Failover*, and *Failback* modes for this resource.

The default for both *Start* and *Failover* modes is AUTO, and the default for *Failback* mode is DISABLE.

4 Click *Next*.

5 Assign nodes to the resource as described in Section 10.8, "Configuring Preferred Nodes for a Resource," on page 160, then click *Finished*.

6 When you are done, verify that the resource was created by viewing its entry in the *Cluster Objects* list on the Cluster Options page.

Clusters

Cluster Options

Cluster: cluster1.clusters.city.mycompany

Properties...   Repair...

View cluster resource configuration information and administer cluster resources for the selected cluster.

Cluster Objects

New | Delete | Details   15 Item(s

	Type	Name	IP Address	Distinguished Name	Pool Name
☐		Master_IP_Address_Resource	137.65.67.44	cn=cluster1,ou=clusters,ou=city,o=mycompany	
☐		avalon	137.65.67.37	cn=avalon,cn=cluster1,ou=clusters,ou=city, o=mycompany	
☐		POOLC1_SERVER	137.65.67.41	cn=CLUSTER1_POOLC1_SERVER,ou=clusters, ou=city,o=mycompany	POOLC1
☐		lxvol44_server			

7 Configure resource priorities for load order as described in Section 10.9, "Configuring Resource Priorities for Load Order," on page 161.

8 (Optional) Assign the resource to a Resource Mutual Exclusion Group as described in Section 10.10, "Configuring Resource Mutual Exclusion Groups," on page 162.

9 Continue with Section 13.5, "Verifying the Load and Unload Scripts," on page 228.

13.5 Verifying the Load and Unload Scripts

After you have created the resource, ensure that the load and unload scripts are working.

1 In iManager, select *Clusters > Cluster Manager*, then select the cluster.

2 Select the check box next to the Linux POSIX volume cluster resource, then click *Online*.

Ensure that the resource successfully enters the *Running* state. If the resource goes comatose instead, it is most likely because you made an error when typing values in the script definitions. Take the resource offline, go to the resource's *Properties > Scripts* page to review and modify its scripts as needed to correct errors, then try again to online the resource.

3 (Optional) Continue with Section 13.6, "Creating a Virtual Server Object for the Linux POSIX Volume Cluster Resource," on page 228.

13.6 Creating a Virtual Server Object for the Linux POSIX Volume Cluster Resource

When you cluster-enable the Linux POSIX volume, a virtual server object is not created automatically as it is when you cluster-enable an NSS pool. You can assign a virtual server name to the cluster resource by using the `/opt/novell/ncs/bin/ncs_ncpserv.py` script to create an `NCS:NCP Server` object for it. Having a virtual server object allows the Linux POSIX volume cluster resource to be viewed in the *Browse* panel in iManager. You can also bind the virtual server name to the IP address to allow users to access the resource via an assigned name in addition to its IP address

The `NCS:NCP Server` object does not give users NCP access to the data on a Linux POSIX volume. An NCP volume is required to do that.

IMPORTANT: If you are creating NCP volumes on the shared Linux POSIX volume, do not follow the procedure in this section. Go to "Creating a Shared NCP Volume on the Linux POSIX Cluster Resource" in the *OES 2 SP3: NCP Server for Linux Administration Guide*, then follow the instructions there.

The virtual server name is stored in Novell eDirectory in an `NCS:NCP Server` object under the Cluster object where you created the resource. You must add a line to the load and unload scripts that identifies the name of this virtual server and a line that binds or unbinds the name to the IP address of the Linux POSIX cluster resource.

13.6.1 Creating the Virtual Server Object

You use the `/opt/novell/ncs/bin/ncs_ncpserv.py` script to create a virtual server object (`NCS:NCP Server`) in eDirectory for the Linux POSIX volume cluster resource. If the resource does not have NCP volumes on it, you do not use the -v option. For information about the `ncs_ncpser.py` script, see Section A.5, "ncs_ncpserv.py Script," on page 270.

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

2 In the console, use the `cd` command to go to the `/opt/novell/ncs/bin` directory.

3 At the terminal console prompt, enter

```
./ncs_ncpserv.py -c lx_volumename -i resource_ip_address
```

Replace the *lx_volumename* and *resource_ip_address* with the information for your particular solution.

Do not use periods in cluster resource names. Novell clients interpret periods as delimiters. If you use a space in a cluster resource name, that space is converted to an underscore.

For example, to create the NCS:NCP Server object for the `lxvol44` cluster resource where the IP address is 10.10.10.44 and the cluster context is `ou=clusters,ou=city,o=mycompany`, enter

```
./ncs_ncpserv.py -c lxvol44 -i 10.10.10.44
```

The confirmation message is displayed:

```
NCP Server 'cn=cluster1_lxvol44_server,ou=clusters,ou=city,o=mycompany'
created.
```

4 Continue with Section 13.6.2, "Modifying the Load Script," on page 229

13.6.2 Modifying the Load Script

After you have created an NCS:NCP Server object, you must modify the load script so that it binds the NCS:NCP Server object to the resource.

1 In iManager, select *Clusters > Cluster Options*, then select the cluster.

2 Click the name link of the Linux POSIX cluster resource to open its Cluster Resource Properties page.

3 Click the *Scripts* tab to open the *Load* script.

4 In the definition area, add the following lines to define the virtual NCP server name:

```
# define NCP server name
NCP_SERVER=cluster1_lxvol44_server
```

Replace the NCP server name with the name for your virtual NCP server.

5 Under the mount line, add a line to bind the NCP server name to the resource IP address:

```
# bind the NCP volume
exit_on_error ncpcon bind --ncpservername=$NCP_SERVER --ipaddress=$RESOURCE_IP
```

6 Click *Apply* to save your changes.

The script changes are not active until the next time the cluster resource is taken offline, and then brought online. Do not active the script changes at this time.

7 Continue with

13.6.3 Modifying the Unload Script

After you have created an NCS:NCP Server object, you must modify the unload script so that it unbinds the NCS:NCP Server object from the resource.

1 In iManager, select *Clusters > Cluster Options*, then select the cluster.

2 Click the name link of the Linux POSIX cluster resource to open its Cluster Resource Properties page.

3 Click the *Scripts* tab, then click *Unload* to open the *Unload* script.

4 In the definition area, add the following lines to define the virtual NCP server name:

```
# define NCP server name
NCP_SERVER=cluster1_lxvol44_server
```

Replace the NCP server name with the name for your virtual NCP server. Use the same value for variable that you did in the load script.

5 Under the definition, add a line to unbind the NCP server name from the resource IP address:

```
# unbind the NCP server name
ignore_error ncpcon unbind --ncpservername=$NCP_SERVER --
ipaddress=$RESOURCE_IP
```

6 Click *Apply* to save your changes.

The script changes are not active until the next time the cluster resource is taken offline, and then brought online.

7 Continue with

13.6.4 Activating the Script Changes

The script changes are not active until the next time the cluster resource is taken offline, and then brought online.

1 In iManager, select *Clusters > Cluster Manager*, then select the cluster.

2 Select the check box next to the Linux POSIX cluster resource, then click *Offline*.

Wait until the resource is reports an Offline status before continuing.

3 Select the check box next to the Linux POSIX cluster resource, then click *Online*.

Wait until the resource is reports an Online status before continuing.

4 Continue with .

13.6.5 Verifying the Virtual Server

Verify that an NCS:NCP Server object appears in the *Browse* panel in iManager.

1 In the iManager toolbar, click the *View Objects* icon.

2 In the left panel, click *Browse*.

3 Browse to the Cluster container to see the virtual server object for the cluster resource, such as `cluster1_lxvol44_server`.
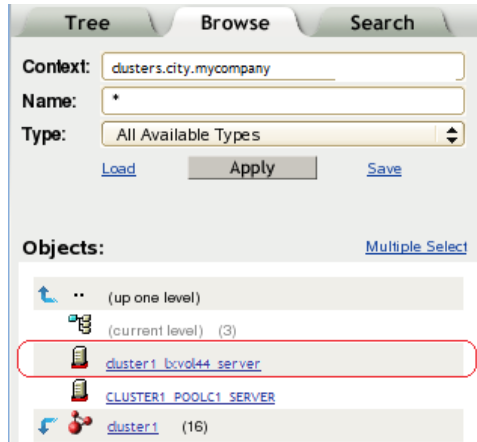
13.7 Sample Scripts for a Linux POSIX Volume Cluster Resource

The scripts in this section are based on the sample values in the following table. Ensure that you replace the sample values with the ones you used in your solution.

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

Variable	Template Value	Sample Value	Description
RESOURCE_IP	a.b.c.d	10.10.10.44	IP address of the virtual cluster server for this cluster resource
Container_name	name	`csm44`	The name you gave to the cluster segment manager.
MOUNT_DEV	`/dev/evms/ $container_name/ volume_name`	`/dev/evms/ $container_name/ lxvol44`	The Linux path for the EVMS volume you created, such as `lxvol44`.
MOUNT_FS	reiserfs	ext3	The file system type you made on the EVMS volume.

Variable	Template Value	Sample Value	Description
MOUNT_POINT	/mnt/mount_point	/mnt/lxvol44	The mount location for the EVMS volume you created. This example shows a mount location with a directory named the same as the EVMS volume name (lxvol44). You can mount the EVMS volume anywhere.

After you create a resource, you can modify its IP address, administrator credentials, the order or IP addresses of the LDAP servers, or other attributes by using the procedure in Section 8.9, "Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster," on page 116.

13.7.1 Sample Load Script for the Linux POSIX Volume Cluster Resource

The following is an example load script for the cluster resource for a Linux POSIX volume:

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
#define the container name
container_name=csm44
# define the device
MOUNT_DEV=/dev/evms/$container_name/lxvol44
# define the mount point
MOUNT_POINT=/mnt/lxvol44

#activate the container
exit_on_error activate_evms_container $container_name $MOUNT_DEV $NCS_TIMEOUT

# mount the file system
ignore_error mkdir -p $MOUNT_POINT
exit_on_error mount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

exit 0
```

13.7.2 Sample Unload Script for the Linux POSIX Volume Cluster Resource

The following is an example unload script for the cluster resource for a Linux POSIX volume:

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
#define the container name
container_name=csm44
# define the device
MOUNT_DEV=/dev/evms/$container_name/lxvol44
# define the mount point
MOUNT_POINT=/mnt/lxvol44

# unmount the volume
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error umount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# deactivate the container
ignore_error deactivate_evms_container $container_name $NCS_TIMEOUT

# return status
exit 0
```

13.7.3 Sample Monitor Script for a Linux POSIX Volume Cluster Resource

The following is an example monitor script for the cluster resource for a Linux POSIX volume:

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
#define the container name
container_name=csm44
# define the device
MOUNT_DEV=/dev/evms/$container_name/lxvol44
# define the mount point
MOUNT_POINT=/mnt/lxvol44

# test the file system
exit_on_error status_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# status the IP address
exit_on_error status_secondary_ipaddress $RESOURCE_IP

# (optional) status of the eDirectory service
#exit_on_error rcndsd status

# (optional) status of the Linux User Management service
# exit_on_error namcd status

exit 0
```

13.8 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.

> **IMPORTANT:** If you did not add a segment manager on top of the CSM container when you set up the resource, creating additional volumes or expanding or shrinking volumes is not possible. See Section 13.3.4, "Adding a Segment Manager to the CSM Container," on page 219.

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

- Section 13.8.1, "Expanding a Volume to a Separate Disk," on page 234
- Section 13.8.2, "Moving a Volume to a Larger Disk," on page 235

13.8.1 Expanding a Volume to a Separate Disk

1 Enter `evmsgui` at a terminal console prompt as the `root` user.

2 In evmsgui, unmount the file system for the shared Linux POSIX volume you want to expand.

3 In evmsgui, click the *Volumes* tab, right-click the volume you want to expand then select *Add Feature*.

4 Select *Drive Linking Feature*, then click *Next*.

5 Provide a name for the drive link, click *Add*, then save your changes.

6 Click *Actions*, select *Create*, and then select *Container*.

7 Ensure that the target disk contains only unpartitioned free space.

   The entire target disk is used for the expansion, so you must make a disk available does not have other volumes or segments on it.

8 Select the *Cluster Segment Manager*, click *Next*, then select the disk you want to expand the volume to.

   If there are no disks without a Segment Manager, then Cluster Segment Manager does not appear as an option.

   The entire disk is used for the expansion, so you must select a disk that does not have other volumes or segments on it.

9 Provide the same settings information (same name, type (Private), owning node, and so on) as the existing container for the volume, the save your changes.

10 Re-mount the volume to be expanded.

11 Click the *Volumes* tab, right-click the volume, then click *Expand*.

12 Select the volume that you are expanding, then click *Next*.

13 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.

14 Select the storage device the volume is being expanded to, select *Expand*, and save your changes.

15 Click *Save* twice, then exit evmsgui.

13.8.2 Moving a Volume to a Larger Disk

1 Enter `evmsgui` at a terminal console prompt as the `root` user.

2 In evmsgui, unmount the file system for the volume you want to move by right-clicking the volume and selecting *Unmount*.

3 Add a larger disk to the CSM container:

   3a In evmsgui, click *Actions*, select *Create*, then click *Container*.

   3b Select the *Cluster Segment Manager*, then click *Next*.

   3c Select the larger disk you want to move the volume to.

   The entire disk is used for the expansion, so you must select a disk that does not have other volumes on it.

   3d Provide the same settings information (name, type (Private), owning node, and so on) as the existing container for the volume, then save your changes.

   3e Click *Save* and exit evmsgui.

4 Restart evmsgui, click the *Containers* tab, then expand the container so that the objects under the container appear.

   The new disk should appear as part of the container.

5 Right-click the object for the disk where the volume resides and select *Replace*.

6 Select the object for the disk where the volume will be moved, then click *Next*.

7 Save your changes.

   Saving your changes could take a while, depending on volume size and other factors.

8 Click *Save*, exit evmsgui, then restart evmsgui.

9 Click the *Volumes* tab, right-click the volume, then select *Check/Repair filesystem*.

   This runs the repair process and ensures that no problems exist on the moved volume.

10 Click the *Disks* tab, right-click the disk the volume was moved from, then select *Remove from container*.

11 Click *Save* twice, then exit evmsgui.

13.9 Deleting Shared Storage

Ensure that you offline the cluster resource before attempting to delete either the cluster resource or the shared storage resource.

WARNING: If you attempt to delete a cluster resource or the shared storage without first offlining it, deletion errors occur, and the data associated with the shared storage is not recoverable.

13.10 Known Issues for Working with Cluster Resources for Linux POSIX Volumes

This section describes known issues for working with EVMS.

13.10.1 Dismount Volumes before Onlining a Comatose Resource

If a Linux POSIX volume resource goes comatose, you must verify that the volume is not mounted on any node in the cluster before bringing the resource online again. Failure to do so might result in corruption.

WARNING: To avoid corruption, ensure that the Linux POSIX volume in a comatose resource is dismounted from any node before attempting to online it.

13.10.2 Cluster Services Must Be Running When Using EVMS

Novell Cluster Services must be running on all nodes in the cluster whenever you make any modifications to the cluster or cluster resources by using the EVMS management tools or the Clusters plug-in in iManager.

13.10.3 Close EVMS Utilities When They Are Not In Use

The EVMS management utilities (evms, evmsn, evmsgui) lock the EVMS engine while running, potentially blocking other EVMS-related actions from taking place. This affects NSS pool and volume actions and Linux POSIX volume actions.

Do not run the EVMS management utilities when they are not being actively used. Do not concurrently open multiple instances of the utilities. Do not open the utilities when using NSSMU or the Clusters or Storage plug-ins for iManager.

13.10.4 Do Not Migrate Resources When EVMS Tools Are Running

Cluster resources for NSS pools and Linux POSIX volumes should not be migrated while the EVMS management utilities (evms, evmsn, and evmsgui) are running.

14 Configuring Novell Cluster Services in a Xen Virtualization Environment

Novell Cluster Services is installed and configured in a virtualization environment by using the same methods and processes as those used on a physical Novell Open Enterprise Server (OES) 2 SP3 server. It can be used in the host and guest virtual environments.

You can install Novell Cluster Services on an OES 2 SP3 Xen virtualization host server (Dom0). You can create cluster resources for the virtual machines by using the Xen and XenLive resource templates. You can fail over or cluster migrate these virtual machine cluster resources to different physical nodes in your cluster. Only the Xen and XenLive templates can be used in the OES 2 SP3 Xen host server environment.

For a virtual machine, you can use any virtualization hypervisor (such as Xen, KVM, VMware, and Hyper-V) that supports the SUSE Linux Enterprise Server 10 SP3 or later operating system, which is the platform used for OES 2 SP3 services. All templates except Xen and XenLive are valid on guest servers (DomU). When Novell Cluster Services is installed on an OES 2 SP3 guest server environment, no additional changes or special configuration are required. You can create clusters consisting of all virtual nodes, or use a combination of virtual and physical nodes.

Although many different cluster virtualization scenarios are possible, only those outlined in the sections below have been tested:

14.1 Prerequisites for Xen Host Server Environments

When using Novell Cluster Services to create cluster resources in a Xen host server environment, one network adapter on each cluster node must be available without Xen bridging to allow Novell Cluster Services to monitor for network state changes in the cluster. You need at least one other network adapter on each cluster node to use as a bridge for the virtual servers that will be hosted on the Xen host.

For each cluster node, determine which network adapter you want to use for Novell Cluster Services communications and network state monitoring, then do the following:

- Disable bridging on the network adapter.
- Assign the IP address of the host server to the network adapter.

For information about why this configuration is necessary, see "TID 7004595: Novell Cluster Services is unable to detect a network link down when installed on a Xen domain 0 (dom0)" (http://www.novell.com/support/viewContent.do?externalId=7004595&sliceId=1) on Novell Support (http://www.novell.com/support).

14.2 Virtual Machines as Cluster Resources

In this scenario, you have an OES 2 Linux cluster configured on physical machines. OES 2 Linux and Xen are installed and configured on each node along with Novell Cluster Services. This part of the Novell Cluster Services configuration does not differ from that of an OES 2 Linux cluster without virtualization.

You then create either Linux or NetWare virtual machines on each cluster node and configure those virtual machines to be cluster resources. You can then fail over or migrate virtual machine cluster resources (entire virtual machines) to different physical nodes in your cluster.

Figure 14-1 depicts how this setup might look. Novell Cluster Services (NCS) is installed and running on the virtual machine (VM) host server.

Figure 14-1 *Virtual Machines as Cluster Resources*
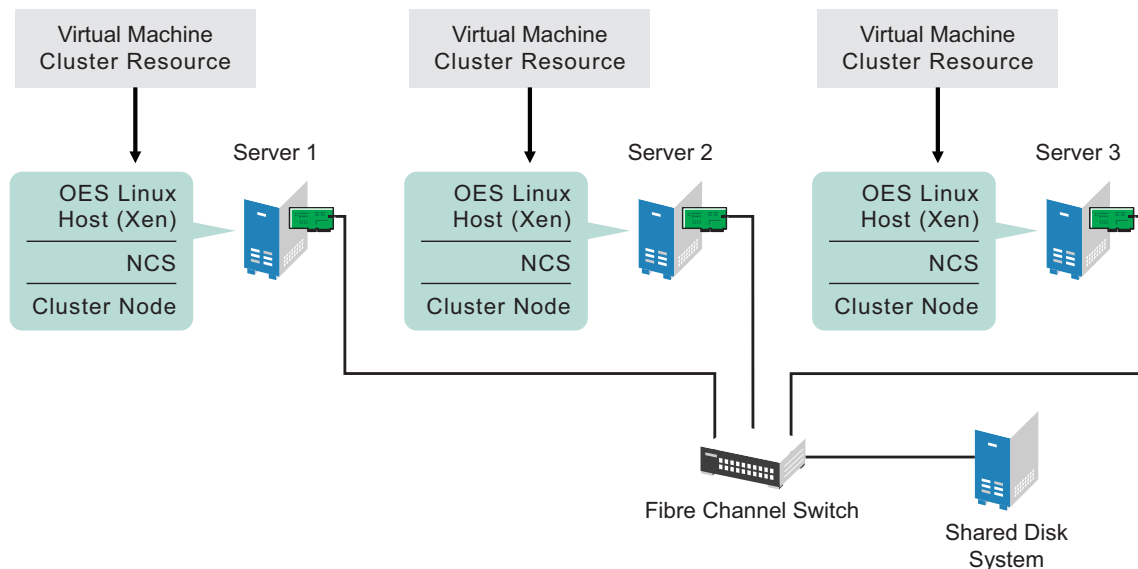
The following sections describe how to create a cluster resource and its cluster scripts for each virtual machine:

14.2.1 Creating a Xen Virtual Machine Cluster Resource

Novell Cluster Services includes two Xen (virtual machine) resource templates, which greatly simplify the process for creating a virtual machine cluster resource. Much of the virtual machine cluster resource configuration is performed automatically by the Xen resource templates. The two templates are named Xen_Template and XenLive_Template. Both templates perform similar functions to automatically configure the cluster resource for the virtual machine.

The XenLive template provides an additional function to allow a manual virtual machine resource migration without the need to boot or bring up the virtual machine on the cluster node where the virtual machine has been migrated. This lets clients continue to access a virtual machine that has been migrated without reconnecting or waiting for the virtual machine to boot or load on the target node.

IMPORTANT: The live migrate function is only useful for a manual virtual machine resource migration, and does not work for a virtual machine resource failover or failback.

Ensure that your Xen setup is working properly before you attempt to set up the Novell Cluster Services clustering for your virtual machines in the Xen host environment. Refer to the *Virtualization with Xen* (http://www.suse.com/documentation/sles10/book_virtualization_xen/data/book_virtualization_xen.html) to find out how to set up XEN and XEN virtual machines.

To configure a virtual machine as a cluster resource:

1 Open your Internet browser and enter the URL for iManager.

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

2 Specify your user name and password, specify the tree where you are installing the cluster, then click *Login*.

3 In iManager, select *Clusters*, then click *Cluster Options*.

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

4 Browse and select the cluster name.

5 On the Cluster Options page, click *New*.

6 Click the *Resource* radio button to specify *Resource* as the resource type you want to create, then click *Next*.

7 Specify a name for the virtual machine resource.

8 In the *Inherit From Template* field, browse to the Cluster object container, then select the desired Xen template name from the list of templates in the container.

The Xen templates are named *Xen_Template* and *XenLive_Template*.

9 Select the *Define Additional Properties* check box, click *Next*, then continue with "Configuring Virtual Machine Load, Unload, and Monitor Scripts" on page 240.

14.2.2 Configuring Virtual Machine Load, Unload, and Monitor Scripts

The Xen resource templates configure the virtual machine resource by automatically creating load, unload, and monitor scripts, setting failover and failback modes, and assigning the virtual machine as a resource to all nodes in the cluster.

The load, unload, and monitor scripts for virtual machine cluster resources do not need to be modified if all the following are true:

- The resource name is the same as the virtual machine name.
- The configuration file name is the same as the virtual machine name.
- The mount point directory name is the same as the virtual machine name.
- You are using the Reiser file system.

If you are not modifying the scripts, continue the setup by configuring the resource policies and the resource server assignments. For information, see Section 10.7, "Configuring the Start, Failover, and Failback Modes for Cluster Resources," on page 159 and Section 10.8, "Configuring Preferred Nodes for a Resource," on page 160.

If you are modifying the scripts, continue with the following sections:

- "Configuring the Load Script" on page 240
- "Configuring the Unload Script" on page 241
- "Configuring the Monitor Script" on page 243

Configuring the Load Script

The virtual machine resource load script page should already be displayed. The load script contains commands to start the virtual machine. You can customize some commands for your specific configuration.

1 View and, if necessary, edit the lines in the script for your specific directory structure, mount point, configuration file, and file system type (in the Xen_Template).

See the following examples of the default Xen_Template and XenLive_Template load scripts:

- "Sample Xen_Template Load Script" on page 241
- "Sample XenLive_Template Load Script" on page 241

2 Click *Next* and continue with "Configuring the Unload Script" on page 241.

Sample Xen_Template Load Script

The Xen_Template load script appears similar to the following example:

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

# filesystem settings
export OCF_RESKEY_device=/dev/evms/$OCF_RESOURCE_INSTANCE
export OCF_RESKEY_directory=/mnt/$OCF_RESOURCE_INSTANCE
export OCF_RESKEY_fstype=reiserfs

#export OCF_RESKEY_options=

# service settings
export OCF_RESKEY_xmfile=$OCF_RESKEY_directory/$OCF_RESOURCE_INSTANCE

# mount the file system
exit_on_error ocf_start Filesystem

# start the service
exit_on_error ocf_start Xen

# return status
exit 0
```

Sample XenLive_Template Load Script

The XenLive_Template load script appears similar to the following example:

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

# filesystem settings
export OCF_RESKEY_directory=/mnt/$OCF_RESOURCE_INSTANCE

# service settings
export OCF_RESKEY_xmfile=$OCF_RESKEY_directory/$OCF_RESOURCE_INSTANCE

# start the service
if [ -n "$NCS_TOFROM" ]
then
  exit_on_error ocf_migrate_from Xen
else
  exit_on_error ocf_start Xen
fi

# return status
exit 0
```

Configuring the Unload Script

The virtual machine resource unload script page should now be displayed. The unload script contains commands to stop the virtual machine. You can customize some commands for your specific configuration.

1 View and, if necessary, edit the lines in the script for your specific directory structure, mount point, configuration files, and file system type (in the Xen_Template).

   Use the same values that you specified in the load script.

   See the following examples of the default Xen_Template and XenLive_Template unload scripts:

2 Click *Next*, then continue the setup by configuring the resource policies and the resource server assignments.

For information, see Section 10.7, "Configuring the Start, Failover, and Failback Modes for Cluster Resources," on page 159 and Section 10.8, "Configuring Preferred Nodes for a Resource," on page 160.

3 If you want to enable monitoring for the resource, continue with "Configuring the Monitor Script" on page 243.

Sample Xen_Template Unload Script

The Xen_Template unload script appears similar to the following example:

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

# filesystem settings
export OCF_RESKEY_device=/dev/evms/$OCF_RESOURCE_INSTANCE
export OCF_RESKEY_directory=/mnt/$OCF_RESOURCE_INSTANCE
export OCF_RESKEY_fstype=reiserfs

#export OCF_RESKEY_options=

# service settings
export OCF_RESKEY_xmfile=$OCF_RESKEY_directory/$OCF_RESOURCE_INSTANCE

# stop the service
ignore_error ocf_stop Xen

# umount the file system
ignore_error ocf_stop Filesystem

# return status
exit 0
```

Sample XenLive_Template Unload Script

The XenLive_Template unload script appears similar to the following example:

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

# filesystem settings
export OCF_RESKEY_directory=/mnt/$OCF_RESOURCE_INSTANCE

# service settings
export OCF_RESKEY_xmfile=$OCF_RESKEY_directory/$OCF_RESOURCE_INSTANCE
export OCF_RESKEY_CRM_meta_migrate_target=$NCS_TOFROM

RC=0

# stop the service
if [ -n $NCS_TOFROM ]
then
  RC = `ocf_migrate_to Xen`
  if [ $RC -ne 0 ]
  then
    ignore_error ocf_stop Xen
  fi
else
  ignore_error ocf_stop Xen
fi

# return status
exit $RC
```

Configuring the Monitor Script

The Xen_Template and XenLive Template each include a resource monitoring script that you can customize. You use the script to monitor the health of a virtual machine cluster resource.

Resource monitoring is disabled by default. If you want to enable resource monitoring for a virtual machine cluster resource, you must enable it prior to customizing the resource monitoring script.

Enabling Resource Monitoring

To enable resource monitoring for a virtual machine cluster resource:

1 In iManager, click *Clusters*, then click *Cluster Options*.

2 Browse and select the Cluster object.

3 Select the check box next to the virtual machine resource, then click the *Details* link.

4 Click the *Monitoring* tab, then select the *Enable Resource Monitoring* check box to enable resource monitoring for the resource.

Resource monitoring is disabled by default.

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

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

6 Specify the number of failures (*Maximum Local Failures*) for the specified amount of time (*Time Interval*).

If the resource monitor detects that the resource fails the number of times specified in the amount of time specified, a failure action initiates.

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

If the failure action initiates and you chose the option to migrate the resource to another server, the resource migrates to the next server in its *Assigned Nodes* list. The resource remains on the server it has migrated to unless you migrate it to another server or the failure action initiates again, in which case it again migrates to the next server in its *Assigned Nodes* list.

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

With resource monitoring, the *Failover*, *Failback*, and *Start* modes have no effect on where the resource migrates. This means that a resource that has been migrated by the resource monitoring failure action does not migrate back to the node it migrated from unless you manually migrate it back.

Viewing or Modifying the Monitor Script

To view or customize the monitor script for the virtual machine's cluster resource:

1 In iManager, click *Clusters*, then click *Cluster Options*.

2 Browse and select the Cluster object.

3 Select the check box next to the virtual machine resource that you created, then click the *Details* link.

4 Click the *Scripts* tab, then click the *Monitor Script* link.

5 View or edit the commands in the script that monitor the resource on the server.

You can use the same commands that would be used at the Linux terminal console.

See the following examples of the default Xen_Template and XenLive_Template monitor scripts:

- ◆ "Sample Xen_Template Monitor Script" on page 244
- ◆ "Sample XenLive_Template Monitor Script" on page 245

6 Specify the *Monitor Script Timeout* value, then click *Apply* to save the script.

The timeout value determines how much time the script is given to complete. If the script does not complete within the specified time, the failure action initiates based on your settings in Step 7 of "Enabling Resource Monitoring" on page 243. Cluster Services marks the monitor process as failed right after the defined timeout expires, but it must wait for the process to conclude before it can start other resource operations.

Sample Xen_Template Monitor Script

The Xen_Template monitor script appears similar to the following example:

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

# filesystem settings
export OCF_RESKEY_device=/dev/evms/$OCF_RESOURCE_INSTANCE
export OCF_RESKEY_directory=/mnt/$OCF_RESOURCE_INSTANCE
export OCF_RESKEY_fstype=reiserfs

#export OCF_RESKEY_options=

# service settings
export OCF_RESKEY_xmfile=$OCF_RESKEY_directory/$OCF_RESOURCE_INSTANCE

# status of the file system
exit_on_error ocf_status Filesystem

# status of the service
exit_on_error ocf_status Xen

# return status
exit 0
```

Sample XenLive_Template Monitor Script

The XenLive_Template monitor script appears similar to the following example:

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

# service settings
export OCF_RESKEY_xmfile=$OCF_RESKEY_directory/$OCF_RESOURCE_INSTANCE

# status of the service
exit_on_error ocf_status Xen

# return status
exit 0
```

14.2.3 Setting Up Live Migration

Live migrations use the XenLive template. You can manually copy the virtual machine configuration files to the same path on each node of the cluster, or you can set up an OCFS2 file system for the configuration files. Do one of the following:

- Manually copy the configuration file for the virtual machine to the same directory (the path must be the same) on each cluster node where the virtual machine will run.
- Configure the OCFS2 file system on a shared disk system and copy the virtual machine configuration file to a directory on the file system. You also must ensure that all cluster nodes where the virtual machine will run have access to the OCFS2 file system on the shared disk system.

    Ensure that your OCFS2 file system is working properly before you attempt to use it with Novell Cluster Services.

    An overview of OCFS2 is available in "Oracle Cluster File System 2" (http://www.suse.com/documentation/sles10/book_sle_reference/data/cha_ocfs2.html) in the *SUSE Linux Enterprise Server 10 SP4 Administration Guide* (http://www.suse.com/documentation/sles10/book_sle_reference/data/book_sle_reference.html). For detailed information about using OCFS2, see the *OCFS2 Project* (http://oss.oracle.com/projects/ocfs2/) on the Oracle Web site.

For information about setting up live migration, see *Configuring a Xen VM for Live Migration with a Cluster* (http://www.novell.com/coolsolutions/feature/19676.html) in Novell Cool Solutions (http://www.novell.com/communities/coolsolutions).

14.3 Virtual Machines as Cluster Nodes

In this scenario, you have OES 2 Linux (Xen) installed and configured on each node (physical machine). You then create either a NetWare or a Linux virtual machine on each physical machine and install and configure Novell Cluster Services on each virtual machine. The combined virtual machines (cluster nodes) comprise one cluster.

IMPORTANT: All virtual cluster nodes in the same cluster should be either Linux or NetWare. Do not mix Linux and NetWare cluster nodes in the same cluster.

You can then create and configure cluster resources on each virtual cluster node. The process for creating and configuring cluster resources on a virtual cluster node is the same as on a physical cluster node.

Cluster resources can be failed over or migrated between virtual cluster nodes that are on the same physical node or on separate physical nodes.

Figure 14-2 depicts using virtual machines as cluster nodes.

Figure 14-2   *Virtual Machines as Cluster Nodes*
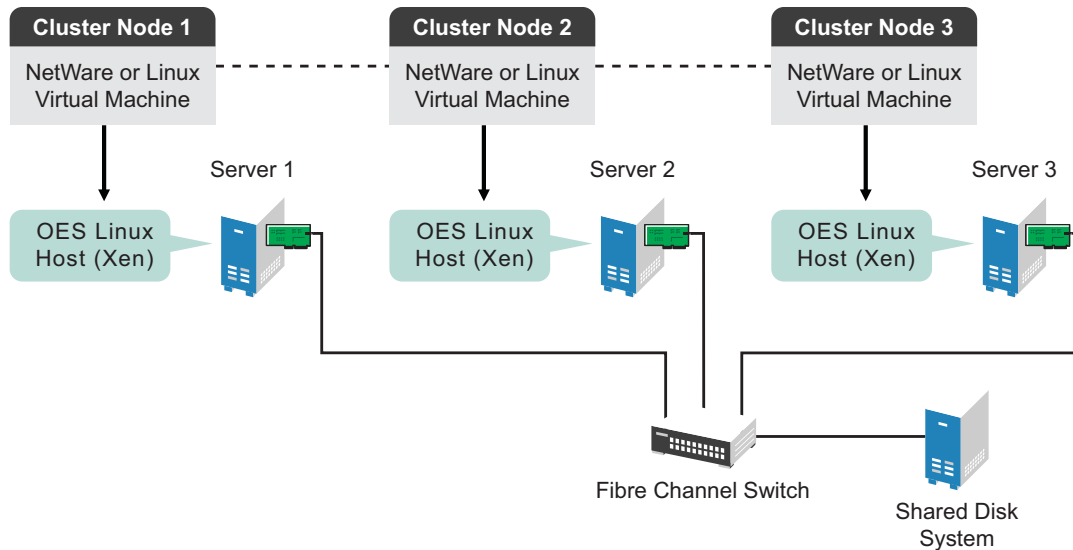
14.4 Virtual Cluster Nodes in Separate Clusters

In this scenario, you have OES 2 Linux (Xen) installed and configured on each node (physical machine). You then create multiple NetWare or Linux virtual machines on each physical machine and install and configure Novell Cluster Services on each virtual machine. During the Novell Cluster Services installation, you create separate clusters of virtual cluster nodes, with each virtual cluster node residing on a separate physical machine. This way you have multiple clusters of virtual cluster nodes on fewer physical machines.

IMPORTANT: All virtual cluster nodes in the same cluster should be either Linux or NetWare. Do not mix Linux and NetWare cluster nodes in the same cluster.

You can then create and configure cluster resources on each virtual cluster node and cluster. The process for creating and configuring cluster resources on a virtual cluster node is the same as on a physical cluster node.

Figure 14-3 depicts using virtual cluster nodes in separate clusters.

Figure 14-3   *Virtual Cluster Nodes in Separate Clusters*
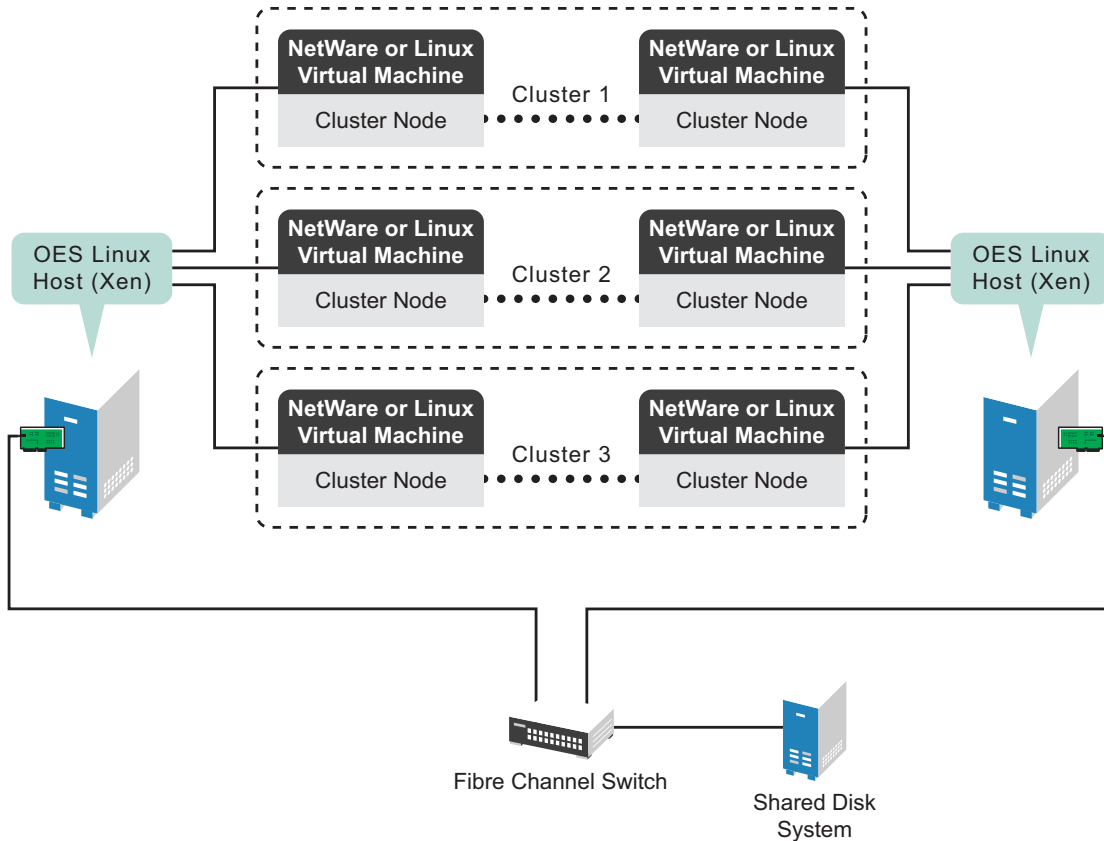
14.5   Mixed Physical and Virtual Node Clusters

This is a temporary scenario that is used for upgrading cluster nodes or converting clusters from physical to virtual cluster nodes.

This can be done through several different methods. One method is to add a virtual NetWare cluster node to an existing physical NetWare cluster. To do this, you install an OES 2 Linux (Xen) server (physical machine). You then create a NetWare virtual machine on the physical machine and install and configure Novell Cluster Services on the NetWare virtual machine. During the Novell Cluster Services installation, you add the NetWare virtual cluster node to your existing NetWare cluster (NetWare physical nodes).

You can then migrate the desired resources from NetWare physical cluster nodes to the NetWare virtual cluster node. This lets you offload resources from physical nodes so you can upgrade hardware and software and then replace the physical NetWare cluster nodes with virtual NetWare cluster nodes.

Figure 14-4 depicts how this setup might look.

Figure 14-4   *Mixed Physical and Virtual Node Cluster*
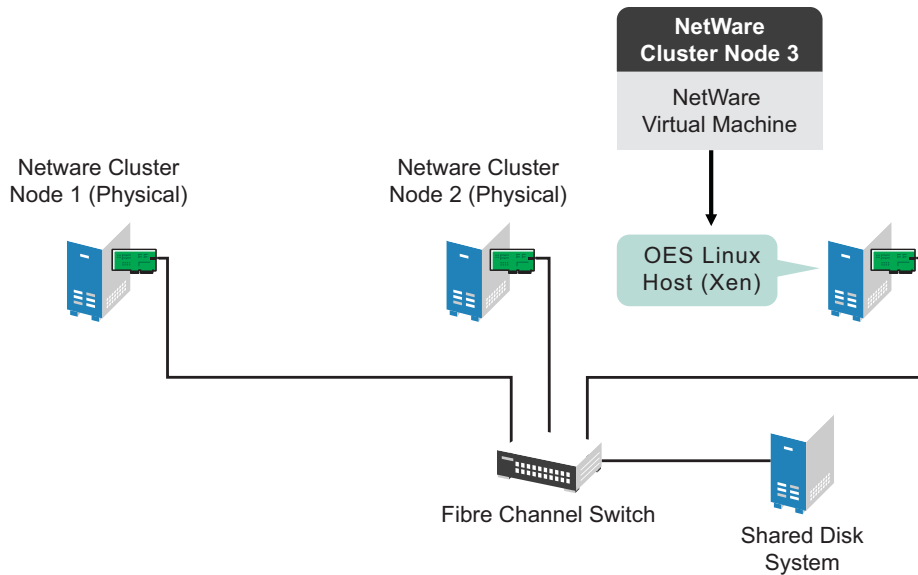
Another method is to install Novell Cluster Services on physical NetWare nodes and create a separate cluster for each node. You then install an OES 2 Linux (Xen) server (physical machine) and create NetWare virtual machines and install Novell Cluster Services on each virtual machine. You can then add one virtual NetWare cluster node to each cluster to create multiple two-node clusters, each containing one physical and one virtual cluster node.

This allows you to migrate the desired resources from each physical NetWare cluster node to the NetWare virtual cluster node in the same cluster. Using this setup, you offload resources from physical nodes so you can upgrade hardware and software and then replace the physical NetWare cluster nodes in each cluster with virtual NetWare cluster nodes.

Figure 14-5 depicts how this setup might look.

Figure 14-5 *Separate Mixed Physical and Virtual Node Clusters*
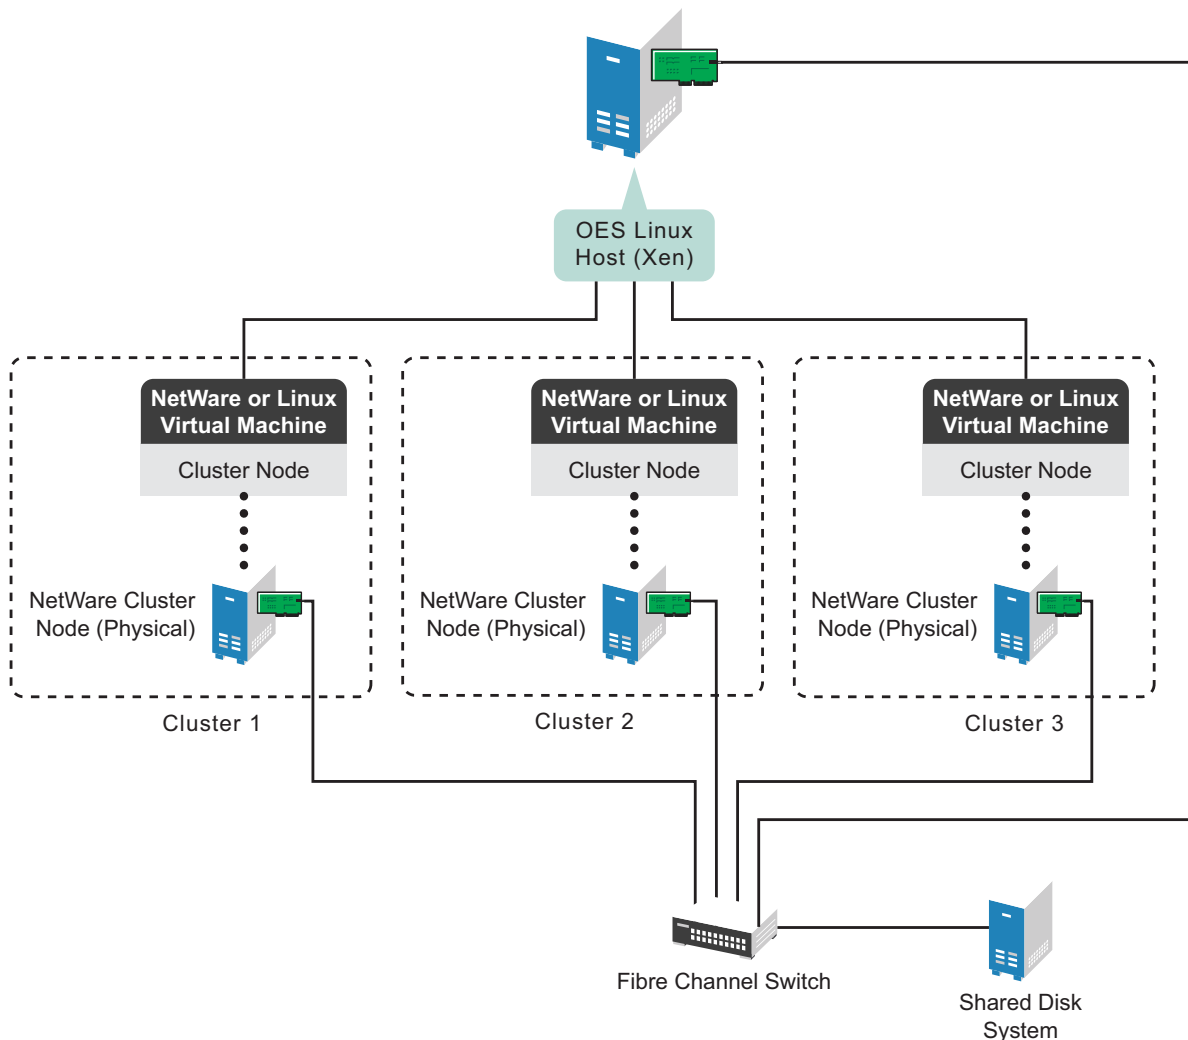
14.6 Additional Information

To get started with virtualization, see "Introduction to Xen Virtualization" (http://www.novell.com/documentation/sles10/book_virtualization_xen/data/sec_xen_basics.html) in the *Virtualization with Xen* (http://www.suse.com/documentation/sles10/book_virtualization_xen/data/book_virtualization_xen.html) guide.

For information on setting up virtualized NetWare, see "Installing, Upgrading, or Updating OES on a Xen-based VM" in the *OES 2 SP3: Installation Guide* guide.

For information on setting up virtualized OES 2 Linux, see "Installing, Upgrading, or Updating OES on a Xen-based VM" in the *OES 2 SP3: Installation Guide* guide.

15 Troubleshooting Novell Cluster Services

This section describes known issues and frequently asked questions for managing Novell Cluster Services.

15.1 File Location Information

Information about Novell Cluster Services scripts and configuration can be found in the following locations:

Information	File Location
Load scripts Unload scripts Cluster Services configuration files	`/var/opt/novell/ncs`
Load and unload script output files (.out) Most resource load and unload issues can be solved with the information in the .out files.	`/var/opt/novell/log/ncs`
Cluster Services scripts, such as ldncs and unldncs	`/opt/novell/ncs/bin`
Resource templates	`/opt/novell/ncs/templates`

15.2 Version Issues

Knowing the location and purpose of the files that make up Novell Cluster Services can be useful in helping you troubleshoot problems and resolve version issues. For information, see Appendix B, "Files for Novell Cluster Services," on page 273.

15.3 Diagnosing a Poison Pill Condition

A poison pill is given to a node if it stops sending out the Novell Cluster Services heartbeat packages, including the case when the node hangs or reboots.

To evaluate the poison pill condition on the node, look at the `/var/log/messages` log from the node that rebooted (was given the poison pill). Check the messages right before the restart. Normally, you can spot the process that caused the server to hang or reboot.

You can run the `cluster stats display` command on the surviving master node to show when the Novell Cluster Services heartbeat packages stopped coming from the rebooted node. To be sure that is the case, you can also run a LAN trace to record the packages among the nodes.

Other events that might cause a pause in sending out heartbeat packages include the following:

- Antivirus software scanning the `/admin` file system
- Network traffic control (including the firewall)
- `check_oesnss_vol.pl` running under Nagios
- Packages that have been recently updated through the YaST Online Update or patch channels

15.4 Diagnosing Cluster Problems

To help Novell Technical Support diagnose your cluster problems, some or all of the following information might be requested:

- Cluster statistics. A report created by the cluster statistics gathering tool.

- Cluster configuration. An HTML report created by the *Clusters > Cluster Manager > Run Report* option in the Clusters plug-in to iManager.
- SAN configuration:
  - Host bus adapter, hub or switch type
  - Device driver name and revision
  - Storage controller type
  - Storage controller firmware revision
- SBD configuration - Single or mirrored partition?
- LAN configuration:
  - Network interface card, hub, or switch type
  - Device driver name and revision
  - Dedicated heartbeat or shared public LAN?
- Server configuration:
  - Type, memory, and number of CPUs
  - Software revisions and patches
  - List of RPMs
  - `/opt/novell/ncs/bin/ldncs` file
- LAN packet trace
- Console log files
- Abend log files
- Server coredump files
- Observable Coincidences? For example:
  - Are poison pills experienced at the same time of day as scheduled maintenance activities?
  - Are server inventory or health checks, virus scans, backups, or periodic eDirectory activity?

15.5   Why did a resource not migrate to the node I specified?

In a failover situation, the cluster looks for the best solution for the node. It goes to the most preferred node if it can. If that node is not available, it attempts to go to the next node in the resource's preferred nodes list. It considers if the node is available and if it has an Resource Mutual Exclusion rules that must be obeyed for resources that are already mounted on the target node. If all of the nodes in the resource's preferred nodes list are not available or have RME conflicts, the resource goes comatose instead of failing over to one of the nodes in its preferred nodes list.

In a cluster migration situation, the resource is currently online and you specify a target node. The resource cannot go to the specified node if any of the following conditions exist, and it stays online on the original node. You can view the target node's log to understand why the resource was not available.

- The specified node is not in the resource's preferred nodes list.
- The specified node is in the resource's preferred nodes list, but the node is currently running another resource that is a Resource Mutual Exclusion conflict for the resource you are attempting to migrate.
- The specified node is in the resource's preferred nodes list, but the node is not available.

15.6 Error Bad XML: Cluster Search Times Out

If you are using in Novell eDirectory 8.7.3*x*, time-outs are possible when you search from iManager for eDirectory objects (such as NCP Server objects, Volume objects, and Cluster objects) because the Object Class attribute is not indexed by default. The LDAP sub-tree search can take over 30 seconds, which causes the query to time out. For example, a Cluster object search from the Cluster Options page returns the error:

```
Bad XML found during parsing when accessing cluster options
```

We recommend that you create a value index on the objects' Object Class attribute. (Object Class is considered an attribute for indexing purposes.) This helps to reduce the time needed for the sub-tree search from over 30 seconds to 10 to 50 milliseconds. For instructions, see "Creating an Index" in the *Novell eDirectory 8.8 SP7 Administration Guide*.

Building indexes speeds up the sub-tree search, even if some partitions being searched do not contain these types of objects. For example, searching for a Cluster object in a context that contains only users is not expected to return results; however, the Object Class search is still performed, and benefits from having an index present.

The sub-tree search performance issue is resolved in the eDirectory 8.8.*x* release with the addition of the AncestorID feature.

15.7 Error NFS Stale Handle Reported on Host when Cluster Failover Occurs

The host I/O reports an `NFS Stale Handle` error immediately after failover is performed on the cluster.

This problem can occur if a user attempts to access the IP address of the resource before the volume is mounted or after the volume is dismounted. You can modify the order of commands in a load script to: NSS, NFS, IP, NCP, [AFP], and [CIFS | Samba]. Reverse the order in an unload script. This allows storage to be mounted before the IP address is bound and to be dismounted after the IP address is unbound. Adjust your scripts to this order, then take the resource offline and bring it online to apply the changes. For sample scripts, see Section 12.9, "Adding NFS Export for a Clustered Pool Resource," on page 193.

15.8 A Device Name Is Required to Create a Cluster Partition

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

15.9 Error 20897 - This node is not a cluster member

If Novell Cluster Services is installed on a node, but an SBD does not exist, NSS tools (including NSSMU, NSS console commands, and the Storage plug-in to Novell iManager) return the following error:

```
Error 20897 - This node is not a cluster member.
```

In a Novell Cluster Services cluster, NSS tools use the cluster's SBD to detect if a node is a cluster member and to lock against concurrent changes to physically shared storage. Without an SBD, NSS tools cannot detect whether a node is a member of the cluster and cannot acquire the locks it needs to execute tasks. In this state, to minimize the risk of corruption, you must ensure that nobody else is changing any storage on any nodes at the same time.

For information about creating an SBD partition, see Section 9.14, "Creating or Deleting Cluster SBD Partitions," on page 136.

15.10 Problems with Clusters Plug-In After Update for Novell iManager 2.7.5

The Clusters plug-in for Novell iManager 2.7.5 in OES 11 SP1 was reorganized to display two tasks in the left panel: *My Clusters* and *My Resources*. If you have problems displaying or using the plug-in after updating to this version, ensure that you performed the following tasks after the update:

- Step 5 in Section 5.7.3, "Uninstalling and Reinstalling the Clusters Plug-In," on page 93
- Section 5.7.4, "Updating Role-Based Services for the Clusters Plug-In for OES 11 SP1," on page 94

15.11 Cluster Resource Goes Comatose Immediately After Migration or Failover

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

For information, see "SLP" on page 61.

15.12 Error 603 - smdr.novell is not registered with SLP for a new cluster resource

You might get an error after you create a cluster resource, indicating that smdr.novell is not registered with SLP for cluster resources, but the smdr.novell service for the node is registered.

```
Error: "cluster--<212>: Read ResVol error -603"
```

The first time a cluster resource is created, smdrd cannot figure it out. The smdrd service is designed to work with resources that contain a data volume.

Ensure that the resource contains a volume and is configured properly, then restart smdrd. Thereafter, smdrd is aware of the cluster resource, and advertises it correctly.

1 Log in to the server as the root user, open a terminal console, then enter

```
rcnovell-smdrd restart
```

15.13 Migration Authentication Error Occurs when Connecting to a New Cluster Resource on the Target Server

You might get a Migration Authentication error when you attempt to connect to a newly created cluster resource on the target server.

```
2013-05-03 12:34:20,095 INFO  - Migration Framework:Authentication:nbackup:
Unable to retrieve the Target Service Name list from <resource_IP_address>, trying
with DNS name <resource_DNS_name>
2013-05-03 12:34:20,096 INFO  - Migration Framework:Authentication:
2013-05-03 12:34:20,096 INFO  - Migration Framework:Authentication:nbackup:
Unable to retrieve the Target Service Name list from <resource_DNS_name>
```

There are two possible causes:

◆ The Storage Management Service (SMS) `smdr.novell` service running on the target server is not aware of the newly created resource.

The Migration tool uses SMS to transfer data between nodes. When a cluster resource is first created, the `smdrd` service is not automatically aware of the new resource. For information, see Section 15.12, "Error 603 - smdr.novell is not registered with SLP for a new cluster resource," on page 255.

Restart `smdrd`. Thereafter, `smdrd` is aware of the cluster resource, and advertises it correctly.

1. Log in to the target server as the `root` user, open a terminal console, then enter

   `rcnovell-smdrd restart`

◆ The cluster resource does not contain a volume.

A properly created pool cluster resource or LVM cluster resource contains at least one volume. Even if you have restarted the `smdrd` service, the service by design does not recognize a resource unless it contains a data volume.

Ensure that the resource contains a volume and is configured properly, then restart `smdrd`. Thereafter, `smdrd` is aware of the cluster resource, and advertises it correctly.

1. Create a volume on the pool or LVM volume group, update the resource scripts as needed, then take the resource offline and bring it online to apply the changes.

2. Log in to the target server as the `root` user, open a terminal console, then enter

   `rcnovell-smdrd restart`

15.14 Cluster View Displays the Wrong Cluster Node Name

In OES 2 SP1, a behavior change was made to address a deadlock defect. After adding a new node to the cluster, you must run the `/opt/novell/ncs/bin/ncs-configd.py -init` script or the `rcnovell-ncs restart` command in order to make `cluster view` and the Clusters plug-in for iManager display the new node's name correctly.

1 Open a terminal console on the master node, then log in as the root user.

2 At the console prompt, enter

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

15.15 NSS Takes Up to 10 Minutes to Load When the Server Is Rebooted (Linux)

In some environments, a timing problem prevents the NDP user space application (`ndpapp`) from loading, and in turn, NSS cannot be loaded until the problem resolves itself. You can increase the UDEV event handling limits to 1024 to circumvent this problem. For instructions, see "NSS Takes Up to 10 Minutes to Load When the Server Is Rebooted" in the *OES 2 SP3: NSS File System Administration Guide for Linux*.

15.16 Could Not Delete This Resource Data_Server (Error 499)

When you try to delete a cluster resource, you get a message "Could Not Delete This Resource Data_Server (Error 499). The error means that the object does not exist in eDirectory. This error can occur if you try to delete an orphaned cluster resource. An orphaned cluster resource does not have a pool or volume associated with it, does not have an IP address, and the scripts are empty.

To delete an orphaned cluster resource:

1 On the master node of the cluster, open a terminal console prompt as the `root` user, then enter the following commands:

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

echo -n "exec rm -f /var/opt/novell/ncs/data_server.*" > /proc/ncs/cluster
```

15.17 Problem Authenticating to Remote Servers during Cluster Configuration

During the Novell Open Enterprise Server (OES) 2 Linux cluster installation and configuration, if you choose *Remote System* on the Novell Cluster Services LDAP Configuration page and you have LDAP configured to point to a NetWare 6.0 or earlier NetWare server, the cluster configuration fails.

To work around this problem, you must edit the `/etc/openldap/ldap.conf` file. Either disable certificates (`TLS_REQCERT <level>` line) or change the file that contains the certificates (`TLS_CACERT <filename>` line). See the `ldap.conf` man page for more information.

15.18 Problem Connecting to an iSCSI Target

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

For instructions on configuring an OES 2 Linux server as an iSCSI initiator and connecting to an iSCSI target, go to "Mass Storage over IP Networks--iSCSI" (http://www.suse.com/documentation/sles10/book_sle_reference/data/cha_inst_system_iscsi.html) in the *SUSE Linux Enterprise Server 10 Administration Guide (http://www.suse.com/documentation/sles10/book_sle_reference/data/book_sle_reference.html).*

15.19   Problem Deleting a Cluster Resource or Clustered Pool

If you attempt to delete a cluster resource or clustered pool without first offlining the cluster resource, deletion errors occur, and the data associated with the clustered pool is not recoverable.

To avoid this problem, offline the cluster resource before attempting to delete either the cluster resource or the clustered pool.

15.20   Where Is the 'Prevent Cascading Failovers' Option?

Novell Cluster Services for NetWare provides a Preventing Cascading Failovers option. Its behavior was controlled using the `clstrlib /hmo=off` option in `ldncs.ncf`. Is there an equivalent option for Linux?

Cascading failover prevention was not implemented for Linux because user space applications (cluster resources) typically cannot cause a kernel failure and cascade failure of other kernels due to failing over user space code, as compared to ring0 NLMs on NetWare.

15.21   Is there a way to uninstall Novell Cluster Services from a server?

There is no utility to remove Novell Cluster Services and its related eDirectory objects. For information, see Section 9.13, "Deleting a Cluster Node from a Cluster," on page 135.

15.22   Cluster Resource Is Stuck in "NDS Sync" or "eDirectory Sync" State

In OES 2, a cluster resource can become stuck in an `NDS Sync` or `eDirectory Sync` state if the character case for cluster information stored in the `/etc/opt/novell/ncs/clstrlib.conf` file does not match the case used for information stored in eDirectory. The `clstrlib.conf` file is automatically populated with information that you enter as you mange a cluster or cluster resource with tools such as iManager, YaST, and commands. If you use the wrong case when typing cluster information, that information is written to the file.

For example, you might enter the wrong case when performing the following cluster management tasks:

- In the YaST Software Installer, you type the cluster name and context when adding a node to an existing cluster.

  IMPORTANT: In OES 2 SP1 and later releases, the installer has been modified to ensure that the information that is added to the `/etc/opt/novell/ncs/clstrlib.conf` file matches the information in eDirectory.

- You modify the cluster configuration for an existing cluster as described in Section 8.9, "Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster," on page 116.

- You modify the case for a container name or object name for an eDirectory object.

For information about how to make corrections in the `/etc/opt/novell/ncs/clstrlib.conf` file to resolve the `NDS Sync` or `eDirectory Sync` state, see *Cluster Resource Stuck in "NDS SYNC" or "eDirectory Sync" State* (TID 7001394) (http://www.novell.com/support/viewContent.do?externalId=7001394&sliceId=2) on the Novell Support (http://www.novell.com/support) Web site.

A patch that helps prevent the occurrence of mismatched cluster information is available in the OES 2 SP1 patch channel and OES 2 SP2 patch channel in their respective March 2010 patch releases.

IMPORTANT: After the patch is applied, the `/etc/opt/novell/ncs/clstrlib.conf` file is case-insensitive.

15.23 Running supportconfig Concurrently on Multiple Nodes Causes Cluster Problems

Running the `supportconfig` command concurrently on multiple nodes can cause the multiple nodes to lose communications with the others and go down.

Novell Support provides the `supportconfig` package to customers in order to gather configuration, log, and error information from both the installation of the OS and from varying applications. The `supportconfig` command was designed without clustering considerations. The command is normally run after a server experiences problems. It is not a good idea to run the command concurrently on multiple nodes in a healthy cluster. If you need to run the `supportconfig` command on multiple nodes in a cluster, run it on one node at a time.

15.24 Nodes Reboot Randomly after January or March 2010 Updates

After applying the January or March 2010 maintenance patches for OES 2 SP1 or OES 2 SP2, some Novell Cluster Services nodes reboot randomly.

This problem was caused by a code error in the NCS timer interruption handler. A fix is available. For information, see *Novell Cluster Services Nodes Rebooting Randomly after March 2010 Updates (TID 7005916)* (http://www.novell.com/support/search.do?usemicrosite=true&searchString=7005916) on the Novell Support (http://www.novell.com/support) Web site.

This issue has been resolved in OES 2 SP3.

16 Security Considerations

This section describes security issues and recommendations for Novell Cluster Services for Novell Open Enterprise Server 2 Linux. It is intended for security administrators or anyone who is responsible for the security of the system. It requires a basic understanding of Novell Cluster Services. It also requires the organizational authorization and the administrative rights to effect the configuration recommendations.

16.1 Cluster Administration Rights

For information about the rights needed to install and manage Novell Cluster Services, see the following sections:

16.2 Ports

For each cluster, you can specify the port used for cluster communication. The default cluster port number is 7023, and is automatically assigned when the cluster is created. You might need to modify this value if there is a port conflict. For information, see Section 8.8, "Modifying the Cluster IP Address and Port Properties," on page 115.

You must specify the eDirectory port when you install clusters. Port 636 is the default port number for eDirectory communications in the tree.

If you are using a firewall, the port must be opened for CIMOM communications between the cluster and iManager. Port 5989 is the default setting for secure HTTP (HTTPS) communications with OpenWBEM.

16.3   Email Alerts

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

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

For information, see Section 8.4, "Configuring Cluster Event Email Notification," on page 109.

16.4   Log Files

NCS normally writes log messages to `/var/log/messages`. You can view events in the *NCS Event Log* in the Clusters plug-in to Novell iManager.

NCS installation writes log messages to `/var/opt/novell/install/ncslog`.

NCS writes resource runtime messages to files under `/var/opt/novell/log/ncs/`. Output log files are named with the resource name, load or unload to indicate which script, and the `.out` extension. For example: `/var/opt/novell/log/ncs/myresource.unload.out` reports the log for the `myresource` unload script messages.

A   Console Commands for Novell Cluster Services

Novell Cluster Services provides several console commands to help you perform certain cluster-related tasks.

- Section A.1, "Cluster Management Commands," on page 263
- Section A.2, "Business Continuity Clustering Commands," on page 268
- Section A.3, "extend_schema Command," on page 268
- Section A.4, "SBD Utility," on page 269
- Section A.5, "ncs_ncpserv.py Script," on page 270

A.1   Cluster Management Commands

To execute a cluster console command, enter `cluster` followed by the command. For example, if you want to display cluster statistics, enter `cluster stats display` at the terminal console. You can also enter `cluster help` at the console prompt to get information on the commands and their functions.

The functions of many of the commands can also be performed using iManager. See the other sections of this document for additional information.

Table A-1 lists the cluster-related terminal console commands and gives a brief description of each command.

Table A-1   *Cluster Console Commands*

Cluster Console Command	Description
`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 to fail over or fail back to another node.
	Specify the resource name in the command and use the YES or NO switch to accept or reject an alert on a resource to control fail over, fail back, or manual start.
	Example
	`cluster alert res1 yes`

Cluster Console Command	Description

CONVERT preview [resource]
 	
CONVERT commit
 	
Beginning in OES 2 SP3, you can use cluster preview without specifying a resource to preview all resources.
 	
CONVERT preview
	Finalizes the cluster conversion from NetWare to Linux after all nodes in a mixed cluster have been converted to Linux. The CLUSTER CONVERT command can be executed only on Linux cluster nodes.
	
	Specify a resource name with the Preview option to view the resource load and unload script changes prior to finalizing the conversion. If you do not provide a resource name, it displays the information for all resources.
	
	Use the Commit switch without specifying a resource to finalize the conversion for all cluster resources.
	
	Examples
	
	`cluster convert preview res1`
	`cluster convert commit`

DOWN	Removes all cluster nodes from the cluster. This command has the same effect as executing the CLUSTER LEAVE command on every server in the cluster. Beginning in OES 2 SP3, you are prompted to confirm the cluster down.
EXEC "path_to_script"	Executes the specified script on all nodes in the cluster. **Example** `cluster exec "command"`
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 email notification settings. Priority displays the resource priority list. Protocol displays the cluster protocol settings. Summary displays the cluster protocol summary. **Example** `cluster info protocol`
JOIN	Adds the node where the command is executed to the cluster and makes the node visible to other servers in the cluster. Novell Cluster Services software must already be installed and running on a node for it to join the cluster.
LEAVE	Removes the node where the command is executed from the cluster. The node will not be visible to other servers in the cluster.

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. Running the command without the {ON|OFF} parameter reports the maintenance status of the cluster (that is, whether maintenance mode is on or off). **Example** `cluster maintenance on`
`MIGRATE {resource} {node}`	Migrates the specified resource from the node where it is currently running to the node that you specify in the command. The node that you specify must be running in the cluster and also be in the resource's *Assigned Nodes* list. **Example** `cluster migrate res1 node1`
`MONITOR {resource} <start | stop | status>`	Manually starts, stops, or checks the status of resource monitoring. The resource must be running on the node where you issue the command. Resource monitoring must be enabled for the resource in order to use this command. **Example** `cluster monitor sh_pool_res_01 status`
`OFFLINE {resource}`	Unloads the specified resource from the node where it is currently running. **Example** `cluster offline res1`

`ONLINE {resource}{node name}`	Starts the specified resource on the most preferred node that is currently active. You can start the resource on a different node by specifying that node in the command. The node that you specify must be running in the cluster and also be in the resource's *Assigned Nodes* list.
	
	Example
	
	`cluster online res1`
	`cluster online res1 node1`

`POOLS`	Lists the NSS pools on the shared disk system that are accessible by Novell Cluster Services.
`RENAME <old_resource_name> <new_resource_name>`	Renames a pool cluster resource. This command must be issued from the master node. The resource must be in offline state to be renamed. The new name must not exist prior to the renaming. Renaming the resource does not modify the resource's virtual server name. **Example** `cluster rename POOL1_SERVER customized_name22`
`RESOURCES`	Lists all resources that currently exist in the cluster. The resources do not need to be online or running.

Cluster Console Command	Description
RESTART [seconds]	Restarts Novell Cluster Services software on all servers in the cluster. The cluster leave process begins immediately. Specify a value of 60 or more seconds as the time to wait before the cluster join begins. The default setting is 60 seconds. The default setting is used if a value is not provided, if the value provided is smaller than 60 seconds, or if the input is invalid. Beginning in OES 2 SP3, you are prompted to confirm the cluster restart. **Example** `cluster restart 60`
SCAN FOR NEW DEVICES	**IMPORTANT:** This command is obsolete on Linux. You can use the `rescan-scsi-bus.sh` script to scan for the new devices on Linux without rebooting. For information, see "Scanning for New Devices without Rebooting" (http://www.suse.com/documentation/sles10/stor_admin/data/scandev.html) in the *SLES 10 SP4: Storage Administration Guide* (http://www.suse.com/documentation/sles10/stor_admin/data/bookinfo.html).

Cluster Console Command	Description
SET {Parameter} {Value}	Sets cluster parameters individually for the cluster. See Chapter 8, "Configuring Cluster Policies and Priorities," on page 105 for more information on cluster parameters.
	Specify one of the following parameters and a value for that parameter:
	IPADDRESS sets the cluster master IP address to the specified value. If you change the cluster IP address, you must restart cluster software on all cluster nodes.
	PORT sets, or lets you change, the IP port number of the cluster master.
	QUORUMWAIT sets the amount of time in seconds that the cluster waits before resources start to load.
	QUORUM sets the number of nodes that must be running in the cluster before resources will start to load.
	HEARTBEAT sets the amount of time in seconds between transmits for all nodes in the cluster except the master.
	TOLERANCE sets the amount of time in seconds that the master node gives all other nodes in the cluster to signal that they are alive.
	MASTERWATCHDOG sets the amount of time in seconds between transmits for the master node in the cluster.
	SLAVEWATCHDOG sets the amount of time in seconds that the slave nodes give the master node in the cluster to signal that it is alive.
	MAXRETRANSMITS sets the maximum number of times transmits will be attempted between the master node and slave nodes.
	ENABLEEMAIL enables and disables email notification. You can set the value to OFF to disable email notification, or either CRITICAL or VERBOSE to enable email notification.
	EMAILADDRESSES lets you specify the email addresses used for email notification. The addresses should be separated by spaces. Using this parameter without specifying any addresses clears existing addresses that have been set previously.
	EMAILOPTIONS sets the email notification options. Specify XML as the value to receive email notification in XML format. Not specifying any value with this parameter turns notification in XML format off.
	RESOURCEPRIORITY sets the resource priority for the cluster.
	Example
	`cluster set ipaddress 10.1.1.1`
STATS {Display, Clear}	The Display parameter reports the node number, node name, and heartbeat information to the console screen.
	The Clear parameter resets the various stats counters. This command is useful for troubleshooting network problems related to the cluster heartbeat protocol.
	Example
	`cluster stats display`
STATUS {resource}	Reports the status of the specified resource. This includes the number of times the resource has been migrated or failed over to another server, the resource state, and the node where the resource is currently running.
VIEW	Displays the node name, cluster epoch number, master node name, and a list of nodes that are currently members of the cluster.

A.2 Business Continuity Clustering Commands

If Novell Business Continuity Clustering (BCC) 1.2 for OES 2 SP1 Linux is installed and running on the cluster, you can use the following additional `cluster` commands to manage the BCC cluster environment.

```
connections
credentials
disable
enable
nsmi
resetresources
```

For information about using these commands, see "Console Commands for BCC" (http://www.novell.com/documentation/bcc/bcc12_admin_lx/data/bcccommands.html) in the *BCC 1.2: Administration Guide for Linux* (http://www.novell.com/documentation/bcc/bcc12_admin_lx/data/bookinfo.html).

A.3 extend_schema Command

/opt/novell/oes-install/util/extend_schema --port port_num admin_username admin_password server_ip_address schema_file

A tree administrator user with credentials to do so must use the `extend_schema` command to extend the eDirectory schema before a cluster is installed anywhere in a tree. This allows container administrators (or non-administrator users) to install a cluster in a container in that same tree without needing full administrator rights for the tree. You need to extend the schema only one time in the tree where you will be installing clusters.

To extend the schema, the tree administrator user modifies the following schema files in the given order:

```
/opt/novell/ncs/schema/ncs.ldif
```

```
/opt/novell/ncs/schema/ncpserver.preldif
```

Replace the parameters with the credentials to access and location of the eDirectory schema files.

Parameter	Description	Example
port_num	The port number you assigned for eDirectory communications in the tree where you plan to install clusters. The default port is 636.	`636`
admin_username	The typeful fully distinguished user name of the administrator who has the eDirectory rights needed to extend the schema.	`cn=admin,o=example`
admin_password	The password of the administrator user.	`pas5W0rd`
server_ip_address	The IP address of the eDirectory server that contains the schema files.	`10.10.10.1`

For example, enter the following commands in the order shown, using the values for your particular solution:

```
/opt/novell/oes-install/util/extend_schema --port 636 cn=admin,o=example
pas5W0rd 10.1.1.1 /opt/novell/ncs/schema/ncs.ldif

/opt/novell/oes-install/util/extend_schema --port 636 cn=admin,o=example
pas5W0rd 10.1.1.1 /opt/novell/ncs/schema/ncpserver.preldif
```

A.4 SBD Utility

The SBD utility (sbdutil) allows you to create, find, or view a Novell Cluster Services SBD partition.

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

We recommend that you carve out a LUN/disk of 20 MB in size to use for the SBD. If you mirror the SBD, you need to carve out a separate LUN/disk of equal size to use. Before you begin, each of these small disks must be initialized and marked as shareable for clustering. You can initialize the device by using NSSMU or the Storage plug-in to iManager. The NSS utility called ncsinit is available for initializing a device and setting it to a shared state.

Syntax

```
sbdutil [-c|-f|-i|-v] [-s] [-d device] [-d device] [-p partition] [-n
cluster_name]

sbdutil -c -d device [-s [size]] [-n cluster_name]

sbdutil -f [-s] [-n cluster_name]

sbdutil -i -p partition [-s] [-n cluster_name]

sbdutil -v [-p partition] [-s] [-n cluster_name]
```

Enter the command at a terminal console prompt as the root user or any other user in admin or ncsgroup.

Options

-c

Create an SBD partition. This option requires at least one device to be specified. You can create a mirrored SBD by supplying more than one device with multiple instances of the -d option.

IMPORTANT: Do not create an SBD partition for a cluster that already has an SBD partition. If you need to re-create the SBD for a cluster, delete its existing SBD first.

To delete an SBD:

1. Enter cluster down at the server console of one cluster server.

   This causes all cluster servers to leave the cluster.

2. Delete the SBD partition.

   You can use nssmu, evmsgui, or other utilities to delete the SBD partition. If the partition is mirrored, delete the RAID device.

-f

Find the SBD partition.

-i

Initialize the SBD partition. If there is information in the `sbdutil -v` view from old nodes that have been removed from the cluster, you can run `sbdutil -i` to initialize the view. This can be done without any issues while the cluster is running.

-v

View the SBD partition.

-d *device*

The device where you want to create an SBD partition. You can create a mirrored SBD by supplying more than one device with multiple instances of the -d option. Use the EVMSGUI (or EVMSN or EVMS) tool to find the names of the devices you want to use, and only use the base (leaf) names (such as sdb or mpathd) with the -d option.

-p *partition*

Use this partition instead of searching for one.

-n *cluster_name*

Use this cluster name instead of getting it from `cluster.xml`.

-s

Assume the device is a shared disk system instead of checking `cluster.xml`. An optional partition size (in MB) can also be specified when creating a partition (-c). The default size is 8 MB. Some of the allocated space is used for storing metadata such as the partition table and MBR (master boot record) for the partition.

Notes

For an iSCSI SAN, a very small iSCSI device might rarely have uncommon CHS (cylinder-head-sector) geometry values. NSS prefers at least 32 sectors per track. If there are fewer than 32 sectors per track, NSS tools using EVMS can fail to create the partition or to mark the device as shareable. You can use fdisk to check for valid CHS values before you initialize and share the device. If necessary, you can use fdisk to set the sectors for the device (such as `/dev/sdx`) to 32:

```
fdisk -H 64 -S 32 /dev/sdx
```

A.5   ncs_ncpserv.py Script

The `/opt/novell/ncs/bin/ncs_ncpserv.py` script creates a virtual NCP Server object (`NCS:NCP Server`) in eDirectory, and associates it with none, one, or multiple NCP volumes that you specify. It automatically renames the NCP Volume objects to use the cluster name instead of the server name where the NCP volume was created. NCP clients access files on the Linux POSIX volume via the virtual server name.

A.5.1   Syntax

At the terminal console prompt on the master node of the cluster, enter the following command as the `root` user:

```
./ncs_ncpserv.py -c lx_volumename -i resource_ip_address [-v <ncp_volumename |
"ncp_volumename1:ncp_volumename2:...">]
```

Replace the *lx_volumename, resource_ip_address,* and *ncp_volumename* with the information for your particular solution.

If the -v option is not specified, all of the NCP volumes that currently exist on the Linux POSIX cluster resource are bound to the IP address.

If you enter multiple volume names, use colons to delimit the names and put quotation marks around the list of names. Volume names can be listed by the volume name (`MY_NNCP_VOL06`) or by the volume distinguished name (`cn=CLUS_02_MY_NNCP_VOL06,o=novell`), or any combination of the two methods.

A.5.2 Examples

In the following examples, the resource IP address is 10.10.10.44, the cluster name is `cluster1` and the cluster context is `ou=clusters,o=mycompany`.

Example 1

To specify a single NCP volume named USERS on the `lxvol44` cluster resource, enter

```
./ncs_ncpserv.py -c lxvol44 -i 10.10.10.44 -v USERS
```

The following confirmation message is displayed:

```
NCP Server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' created.

Object 'cn=servername_USERS,ou=clusters,o=mycompany' renamed to
'cn=cluster1_USERS,ou=clusters,o=mycompany'.
The volume name you need to use in the scripts is: USERS
NCP server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' and volume
'cn=cluster1_USERS,ou=clusters,o=mycompany' are linked with each other.
```

Example 2

To specify multiple NCP volumes on the `lxvol44` cluster resource, enter

```
./ncs_ncpserv.py -c lxvol44 -i 10.10.10.44 -v
"USERS:MY_NCP_VOL06:cn=servername_MY_NCP_VOL07,ou=clusters,o=novell"
```

The following confirmation message is displayed:

```
NCP Server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' created.

Object 'cn=servername_USERS,ou=clusters,o=mycompany' renamed to
'cn=cluster1_USERS,ou=clusters,o=mycompany'.
The volume name you need to use in the scripts is: USERS
NCP server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' and volume
'cn=cluster1_USERS,ou=clusters,o=mycompany' are linked with each other.

Object 'cn=servername_MY_NCP_VOL06,ou=clusters,o=mycompany' renamed to
'cn=cluster1_MY_NCP_VOL06,ou=clusters,o=mycompany'.
The volume name you need to use in the scripts is: MY_NCP_VOL06
NCP server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' and volume
'cn=cluster_MY_NCP_VOL06,ou=clusters,o=mycompany' are linked with each other.

Object 'cn=servername_MY_NCP_VOL07,ou=clusters,o=mycompany' renamed to
'cn=cluster1_MY_NCP_VOL07,ou=clusters,o=mycompany'.
The volume name you need to use in the scripts is: MY_NCP_VOL07
NCP server 'cn=cluster1_lxvol44_server,ou=clusters,o=mycompany' and volume
'cn=cluster_MY_NCP_VOL07,ou=clusters,o=mycompany' are linked with each other.
```

B Files for Novell Cluster Services

Knowing the location and purpose of the files that make up Novell Cluster Services can be useful in helping you troubleshoot problems and resolve version issues. Table B-1 lists the path and purpose for some of the files that are part of Novell Cluster Services.

Table B-1   *Location and Purpose of Novell Cluster Services Files*

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

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

NCS File Name and Path	Purpose
/lib/modules/*kernel_dir*/ncs/clstrlib.ko	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
/lib/modules/*kernel_dir*/ncs/cma.ko	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
/lib/modules/*kernel_dir*/ncs/cmsg.ko	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
/lib/modules/*kernel_dir*/ncs/crm.ko	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
/lib/modules/*kernel_dir*/ncs/css.ko	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
/lib/modules/*kernel_dir*/ncs/cvb.ko	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
/lib/modules/*kernel_dir*/ncs/gipc.ko	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
/lib/modules/*kernel_dir*/ncs/sbd.ko	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
/lib/modules/*kernel_dir*/ncs/sbdlib.ko	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
/lib/modules/*kernel_dir*/ncs/vipx.ko	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.
/lib/modules/*kernel_dir*/ncs/vll.ko	Kernel module. Replace *kernel_dir* with the current kernel directory. Use `uname -r` to see the current kernel directory.

C Comparing Novell Cluster Services for Linux and NetWare

Table C-1 compares the features and capabilities of Novell Cluster Services 1.8.5 for Linux on Novell Open Enterprise Server (OES) 2 Linux to Novell Cluster Services 1.8.5 for NetWare on NetWare 6.5 Support Pack 7 or later.

Table C-1 *Comparison of Novell Cluster Services for OES 2 Linux and NetWare 6.5 SP7 or Later*

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
Operating system	OES 2 Linux (OES 2 built on SUSE Linux Enterprise Server 10 SP1) or later	NetWare 6.5 SP7 or later
Two-node cluster with OES 2 license	Yes	Yes
Up to 32 nodes in a single cluster	Yes	Yes
Guest servers on Xen VMs as cluster nodes	Yes	Yes
Business Continuity Clustering support	BCC 1.2 for OES 2 SP1 Linux and later	BCC 1.1 SP2 for NetWare 6.5 SP8 and later
Administrator users	The administrator user whose credentials you provide during the install is the cluster administrator. The tree admin is not automatically given rights. Rights must be granted manually. For information, see Section 5.6, "Configuring Additional Administrators," on page 92.	The administrator user whose credentials you provide during the install is the cluster administrator. For NetWare, rights are automatically extended to the tree administrator user.
Directory-based cluster configuration	Common schema for NetWare and Linux	Common schema for NetWare and Linux

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
Schema extension during the Novell Cluster Services install	In OES 2 Linux, the user who installs Novell Cluster Services must have schema extension rights.	Yes; the user who installs Novell Cluster Services must have schema extension rights.
	In OES 2 SP1 Linux and later, the schema extension is performed separately from the Novell Cluster Services install. A user with schema extension rights extends the schema. For information, see Section 5.2, "Extending the eDirectory Schema to Add Cluster Objects," on page 78.	
	Afterwards, any administrator with sufficient rights can install Novell Cluster Services. For information, see Section 5.3, "Assigning Install Rights for Container Administrators (or Non-Administrator Users)," on page 80.	
Forward migration for Novell Cluster Services	OES 1 SP2 Linux to OES 2 Linux or later. Down cluster and rolling cluster upgrade are supported.	NetWare 6.5 SP6 to NetWare 6.5 SP7 or later. Down cluster and rolling cluster upgrade are supported.
		NetWare 6.0 to NetWare 6.5 SP7 or later. Down cluster and rolling cluster upgrade are supported.
		NetWare 5.1 to NetWare 6.5 SP7 or later. Only the down cluster upgrade is supported.
Cluster conversion from NetWare to Linux	NetWare 6.5 SP7 or later to OES 2 Linux or later	Not applicable
	NetWare 6.5 SP6 to OES 2 Linux or later	
	NetWare 6.0 with latest service packs and patches to OES 2 Linux. Uses the same process as NetWare 6.5 to OES 2 Linux.	

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
Mixed-node Linux and NetWare clusters	Supported only for rolling cluster conversions from NetWare to Linux.	Supported only for rolling cluster conversions from NetWare to Linux.
	NSS pools created on Linux cannot fail over to NetWare nodes.	NSS pools created on NetWare can fail over to Linux nodes.
	In a mixed-cluster environment, you cannot add storage to the cluster or modify the existing storage pools.	
	After adding a Linux node to a NetWare cluster, it is no longer possible to add new NetWare nodes. Only Linux nodes can be added.	
SBD (split-brain detector)	Yes; during the install on the first server in the cluster, or by using the sbdutil after the install and before adding a second node to the cluster.	Yes; during the install on the first server in the cluster.
Mirrored SBD	Yes; by using evmsgui after the install and before adding a second node to the cluster.	Yes; during the install on the first server in the cluster.
Shared disks	Fibre Channel SAN LUNs	Fibre Channel SAN LUNs
	iSCSI SAN LUNs	iSCSI SAN LUNs
	SCSI disks (shared external drive arrays)	SCSI disks (shared external drive arrays)
Cluster-aware shared devices	Yes; requires using the EVMS Cluster Segment Manager to provide cluster-awareness similar to NetWare Media Manager.	Yes; provided by NetWare Media Manager.
Requires Novell Storage Services (NSS)	Required only if you create NSS pools as cluster resources.	Yes; NSS is the default file system on NetWare.
		Novell Cluster Services is not supported on NetWare traditional volumes.
Disk format when using shared NSS pools	NetWare Segment Manager	NetWare Segment Manager
Requires NCP (NetWare Core Protocol)	Requires NCP Server to be installed if you create cluster resources that use NCP, such as NSS pools, NCP volumes, and Dynamic Storage Technology shadow volume pairs.	Yes; NCP is the default file access protocol on NetWare.

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
NSS pools as cluster resources	Yes For information, see Chapter 12, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 173. Shareable for Clustering Multiple-Server Activation Prevention (MSAP) Cluster volume broker; Linux kernel module handles NSS pool events.	Yes For information, see "Setting Up Cluster Resources for Novell Cluster Services" in the *NW6.5 SP8: Novell Cluster Services 1.8.5 Administration Guide*. Shareable for Clustering Multiple-Server Activation Prevention (MSAP) Cluster volume broker
Linux POSIX file systems as cluster resources	Yes For information, see Chapter 13, "Configuring and Managing Cluster Resources for Shared Linux POSIX Volumes," on page 213.	Not applicable
NCP volumes on Linux POSIX file systems as cluster resources	Yes For information, see "Configuring NCP Volumes with Novell Cluster Services" in the *OES 2 SP3: NCP Server for Linux Administration Guide*.	Not applicable
Dynamic Storage Technology shadow volume pairs as cluster resources	Yes; by combining the load and unload scripts for shared NSS pools and managing the pair as a single cluster resource. For information, see "Configuring DST Shadow Volume Pairs with Novell Cluster Services" in the *OES 2 SP3: Dynamic Storage Technology Administration Guide*.	Not supported by Dynamic Storage Technology.
Xen virtual machines as cluster resources	Yes For information, see Section 14.2, "Virtual Machines as Cluster Resources," on page 238.	Not applicable
iManager 2.7	Yes	Yes
Clusters plug-in and Storage Management plug-in for iManager	Yes	Yes
Cluster-enabling NSS pools by using the NSS plug-in for iManager	Yes	Yes
Cluster-enabling NSS pools by using the NSS Management Utility (NSSMU)	Yes	Yes
Command line interface	Yes; using the terminal console	Yes; using the terminal console

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
XML-based API	Yes; same as NetWare except that it uses the /_adminfs path on Linux.	Yes; same as for Linux except that it uses the _admin: volume on NetWare.
Load, unload, and monitor scripts	Yes Script commands differ. Scripts are automatically translated from NetWare commands to Linux commands during the cluster conversion from NetWare to Linux. For a comparison of script commands, see "Planning the Conversion of Load and Unload Scripts" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*.	Yes
NCP support for accessing files on shared NSS pools	Yes	Yes
NCP support for accessing files on shared NCP volumes on Linux POSIX file systems	Yes	Not applicable
Novell AFP support for accessing files on shared NSS pools	Yes; in OES 2 SP1 Linux and later	Yes
Novell CIFS support for accessing files on shared NSS pools	Yes; in OES 2 SP1 Linux and later. Cross-protocol locking is not supported in OES 2 SP1.	Yes; with cross-protocol locking
Linux Samba/CIFS support for accessing files on shared NSS pools on Linux	Yes; with support for cross-protocol locking that allows both NCP users and Linux Samba/CIFS users to concurrently access data. Requires users to be Linux-enabled with Linux User Management. Requires Universal Password.	Not applicable
Linux Samba/CIFS support for accessing files on shared NCP volumes on Linux POSIX file systems	Yes; with support for cross-protocol locking that allows both NCP users and Linux Samba/CIFS users to concurrently access data. Requires users to be Linux-enabled with Linux User Management. Requires Universal Password.	Not applicable
Linux Samba/CIFS support for accessing files on shared Linux POSIX file systems	Yes Requires users to be Linux-enabled with Linux User Management. Requires Universal Password.	Not applicable
Leverage Heartbeat 2 resource agents	Yes	Not applicable

Feature or Capability	Cluster Services for Linux	Cluster Services for NetWare
LAN fault tolerance	Channel bonding	NIC teaming
	For information, see `/usr/src/ linux/Documentation/ bonding.txt`	For information, see "NIC Teaming" in the *NW 6.5 SP8: TCP/ IP Administration Guide*
Multipath I/O	Device Mapper - Multipath I/O, or third-party MPIO solutions	Media Manager Multipath I/O, or third-party MPIO solutions
	For information, see "Managing Multipath I/O for Devices" (http:// www.suse.com/documentation/ sles10/stor_admin/data/ multipathing.html) in the *SLES 10 SP3/SP4 Storage Administration Guide* (http://www.suse.com/ documentation/sles10/stor_admin/ data/bookinfo.html).	For information, see "Managing Multipath I/O to Devices (NetWare)" in the *NW 6.5 SP8: NSS File System Administration Guide*.

D Comparing Clustering Support for OES 2 Services on Linux and NetWare

Table D-1 compares clustering support for Novell Open Enterprise Server (OES) 2 services when using Novell Cluster Services 1.8.5 on OES 2 Linux and NetWare 6.5 SP7 or later.

NSS pool cluster resources can be cluster migrated from NetWare to Linux as part of a cluster conversion. If the resource contains data only, no additional steps are required. However, clustered services can require special handling. For information, see the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*.

Table D-1 *Comparison of Clustering Support for OES 2 Services on Linux and NetWare*

Service	OES 2 NetWare	OES 2 Linux	Comments
AFP (Apple Filing Protocol)	Yes See "Setting Up for Macintosh" in the *NW 6.5 SP8: AFP, CIFS, and NFS (NFAP) Administration Guide*.	Yes; for OES 2 SP1 Linux and later See "Configuring AFP with Novell Cluster Services for an NSS File System" in the *OES 2 SP3: Novell AFP For Linux Administration Guide*.	For information about cluster migrating an AFP service cluster resource from NetWare to Linux, see "Novell AFP" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*.
Apache Web Server	Yes See "Apache with Novell Cluster Services" in the *NW6.5 SP8: Novell Cluster Services 1.8.5 Resource Configuration Guide*.	Yes; use the standard Apache Web Server for Linux.	For information about cluster migrating an Apache service cluster resource from NetWare to Linux, see "Apache HTTP Server" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*.

Service	OES 2 NetWare	OES 2 Linux	Comments
Archive and Version Services (Novell)	Yes See "Installing and Configuring New Archive Servers" in the *NW 6.5 SP8: Novell Archive and Version Services 2.1 Administration Guide*.	Yes See "Configuring Archive and Version Service for Novell Cluster Services" in the *OES 2 SP3: Novell Archive and Version Services 2.1 Administration Guide for Linux*.	On NetWare, it uses the MySQL database. On Linux, it uses PostgreSQL instead of MySQL. The archive database must reside on a Linux POSIX file system. For information about converting an Apache service cluster resource from NetWare to Linux, see "Novell Archive and Version Services" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*.
CIFS (Windows File Services)	Yes; Novell CIFS See "Setting Up for Windows" in the *NW 6.5 SP8: AFP, CIFS, and NFS (NFAP) Administration Guide*.	Yes; Novell CIFS See "Configuring CIFS with Novell Cluster Services for an NSS File System" in the *OES 2 SP3: Novell CIFS for Linux Administration Guide*.	For information about converting a cluster resource from NetWare to Linux, see "Novell CIFS" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*.
DFS (Novell Distributed File Services) Volume location database (VLDB)	Yes See "Clustering Novell Distributed File Services" in the *NW 6.5 SP8: Novell Distributed File Services Administration Guide*.	Yes See "Clustering Novell Distributed File Services" in the *OES 2 SP3: Novell Distributed File Services Administration Guide for Linux*.	For information about converting a cluster resource from NetWare to Linux, see "Novell Distributed File Services VLDB" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*.
DHCP	Yes See "Clustering in NetWare 6.5" in the *NW 6.5 SP8: Novell DNS/ DHCP Services Administration Guide*.	Yes See "Configuring DHCP with Novell Cluster Services for the Linux File System" in the *OES 2 SP3: Novell DNS/DHCP Administration Guide*.	DHCP for Linux supports using a shared Linux POSIX file system or a shared NSS (Novell Storage Services) pool (supported as of OES 2 SP1) for the cluster resource. For information about converting a cluster resource from NetWare to Linux, see "DHCP Server" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*.

Service	OES 2 NetWare	OES 2 Linux	Comments
DNS	Yes See "Creating a Cluster-Enabled DNS Server" in the *NW 6.5 SP8: Novell DNS/DHCP Services Administration Guide*.	Yes See "Configuring DNS with Novell Cluster Services" in the *OES 2 SP3: Novell DNS/DHCP Administration Guide*.	The association of the DNS Server object with the NCP Server must be changed manually after the cluster conversion. It is not done as part of the simple cluster migration. For information about converting a cluster resource from NetWare to Linux, see "DNS Server" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*.
Dynamic Storage Technology service	Not available	Can be used in a cluster, but the service is not clustered. See also Storage, DST shadow volume pairs.	DST runs on each OES 2 Linux node and you set the global server-level parameters to be the same on each one.
eDirectory 8.8	No	No	eDirectory has its own redundancy built in (multiple replicas) and would not benefit from being clustered.
eDirectory Certificate Server	Yes See "Server Certificate Objects and Clustering" in the *Novell Certificate Server 3.3.1 Administration Guide* (http://www.novell.com/documentation/crt33/).	Yes See "Server Certificate Objects and Clustering" in the *Novell Certificate Server 3.3.8 Administration Guide*.	For information about converting a cluster resource from NetWare to Linux, see "eDirectory Server Certificates" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*.
exteNd Application Server and MySQL	Yes; NetWare 6.5 SP2 or earlier. See "Configuring Novell exteNd Application Server and MySQL with Novell Cluster Services" in the *NW6.5 SP8: Novell Cluster Services 1.8.5 Resource Configuration Guide*.	Not available on Linux.	This install option was discontinued beginning in NetWare 6.5 SP3. See also MySQL.
FTP Server	Yes See "Cluster-Enabling NetWare FTP Server" in the *NW 6.5 SP8: Novell FTP Administration Guide*.	No	Use the standard Linux FTP solution.

Service	OES 2 NetWare	OES 2 Linux	Comments
Identity Manager 3.6 Bundle Edition	Can be used in a cluster, but is not clustered. Requires Identity Manager 3.5.	Can be used in a cluster, but is not clustered.	In OES 2 SP1 Linux, it is 32-bit only and requires a 32-bit operating system.
iPrint	Yes See the *NW 6.5 SP8: iPrint Administration Guide*.	Yes See the *OES 2 SP3: iPrint for Linux Administration Guide*.	For information about converting a cluster resource from NetWare to Linux, see "Novell iPrint" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*.
MySQL	Yes See "Configuring MySQL on Novell Clustering Services" in the *NW 6.5 SP8: Novell MySQL Administration Guide*.	Yes; use the standard MySQL service for Linux. MySQL 5.0.*x* on OES2 Linux is offered under the GPL.	For Linux, use a procedure similar to the one on NetWare. Use the Linux commands for MySQL in the load and unload scripts. Use a Linux path on a shared Linux POSIX file system for the MySQL database.
NCP Server	Can be used in a cluster, but is not clustered.	Can be used in a cluster, but is not clustered. See also Storage, NCP volumes on Linux POSIX file systems.	NCP Server runs on each server node in the cluster. It should have the same configuration on each node of the cluster.
NetStorage	Yes See "Configuring NetStorage with Novell Cluster Services" in the *NW 6.5 SP8: NetStorage Administration Guide*.	Yes See "Configuring NetStorage with Novell Cluster Services" in the *OES 2 SP3: NetStorage Administration Guide*.	No known issues.
NFS	Yes See "Cluster-Enabling Native File Access for UNIX" in the *NW 6.5 SP8: AFP, CIFS, and NFS (NFAP) Administration Guide*.	No; native to Linux	Use the standard Linux solution for NFS.
Novell iFolder 2.1.*x*	Yes See "Configuring iFolder on Novell Cluster Services" in the *OES 2 SP1: Novell iFolder 2.1 Installation and Administration Guide*.	Not applicable	iFolder 2.1x is replace by iFolder 3.6 and later in OES 2 Linux.

Service	OES 2 NetWare	OES 2 Linux	Comments
Novell iFolder 3.x	Not applicable	Yes	See "Clustering iFolder 3.8 Servers with Novell Cluster Services for Linux"in the *Novell iFolder 3.8 Administration Guide*.
Printing	Yes	Yes	See iPrint.
QuickFinder	Yes	Yes	See Search.
Search (QuickFinder)	Yes; requires QuickFinder 4.2.0. Install QuickFinder on each server in your cluster, and use virtual search servers. See "Configuring QuickFinder Server for Novell Cluster Services" in the *NW 6.5 SP8 Novell QuickFinder Server 5.0 Administration Guide*.	Yes; requires QuickFinder 5.0.1. Install QuickFinder on each server in your cluster, and use virtual search servers. See "Configuring QuickFinder Server for Novell Cluster Services" in the *OES 2 SP3: Novell QuickFinder Server 5.0 Administration Guide*.	For information about converting a cluster resource from NetWare to Linux, see "QuickFinder Server" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*.
Storage, DST shadow volume pairs	Not applicable	Yes See "Configuring DST Shadow Volume Pairs with Novell Cluster Services" in the *OES 2 SP3: Dynamic Storage Technology Administration Guide*.	DST shadow volumes are on shared NSS pools that are created separately, then managed in the same load/unload scripts.
Storage, NCP volumes on Linux POSIX file systems	Not applicable	Yes For information, see "Configuring NCP Volumes with Novell Cluster Services" in the *OES 2 SP3: NCP Server for Linux Administration Guide*.	The NCP Server service is not clustered; its volumes can be clustered.
Storage, NetWare Traditional volumes	No	Not applicable	

Service	OES 2 NetWare	OES 2 Linux	Comments
Storage, NSS pools and volumes	Yes See "Setting Up Cluster Resources for Novell Cluster Services" in the *NW6.5 SP8: Novell Cluster Services 1.8.5 Administration Guide*.	Yes See Chapter 12, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 173.	NSS software is not clustered. The NSS service must be installed and running on each server node where you want to load the shared pools. Server-level NSS parameter changes must be made manually on each node. Shared NSS pools and volumes that were created on NetWare can be migrated or failed over from a NetWare server to a Linux server during a cluster conversion. For information, see the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*. Most NSS features are available on both platforms. For information see "Cross-Platform Issues for NSS" in the *OES 2 SP3: NSS File System Administration Guide for Linux*.
Tomcat	Yes See "Configuring Tomcat and Novell Cluster Services" in the *NW6.5 SP8: Novell Cluster Services 1.8.5 Resource Configuration Guide*.	Yes; native to Linux Use a similar procedure to the one outlined for Tomcat on NetWare, but use the Linux locations and files.	You cannot convert the NetWare Tomcat configuration for a Linux server.
Xen guest servers as nodes in a cluster	Virtualized NetWare nodes can be used in NetWare clusters. Nodes can be any combination of virtual and physical servers.	Virtualized OES 2 Linux nodes can be used in OES 2 Linux clusters. Nodes can be any combination of virtual and physical servers.	See Chapter 14, "Configuring Novell Cluster Services in a Xen Virtualization Environment," on page 237.
Xen virtual machines on the host server	Not applicable	Yes; use the Xen and XenLive templates.	See Section 14.2, "Virtual Machines as Cluster Resources," on page 238.

E Documentation Updates

This section contains information about documentation content changes made to the *Novell Cluster Services for Linux Administration Guide* since the initial 1.8.5 release for Novell Open Enterprise Server 2.

This document was updated on the following dates:

E.1 December 4, 3013

The new site for iManager 2.7.6 and earlier is at https://www.netiq.com/documentation/imanager27.

Updates were made to the following section. The changes are explained below.

E.1.1 Managing Clusters

Location	Change
Section 8.8, "Modifying the Cluster IP Address and Port Properties," on page 115	A cluster restart is required before iManager can continue to manage the cluster with its new master IP address.
Section 8.9, "Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster," on page 116	This section was relocated here from the "Managing a Cluster" section.
"Changing the IP Addresses of Cluster Resources" on page 118	This section was modified to also address Linux volume cluster resources.
"Changing the IP Address of the Master IP Resource" on page 120	This section is new.

E.2 September 6, 2013

Updates were made to the following sections. The changes are explained below.

E.2.1 Configuring and Managing Cluster Resources

Location	Change
Section 10.1.2, "Using Parameter Values with Spaces in a Cluster Script," on page 148	In cluster scripts, if a parameter value contains spaces, the preferred solution is to enclose the value with both double-quotes (`"`) and a single-quotes (`'`), such as `"'quoted text'"`.
Section 10.1.3, "Using Double Quotation Marks in a Cluster Script," on page 148	This section is new.
Section 10.8, "Configuring Preferred Nodes for a Resource," on page 160	By default, a new resource is configured to run on all nodes, including future nodes. If its preferred node list is ever changed (directly or indirectly) by administrators, future nodes are not automatically added to the resource's preferred nodes list. **IMPORTANT:** Ensure that you prepare the node for the services in the resource before you migrate or fail over the resource to it.

E.2.2 Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 12.9, "Adding NFS Export for a Clustered Pool Resource," on page 193	Before you use the `exportfs(8)` command in resource scripts, ensure that you install the `nfs-kernel-server` package on all of the cluster nodes in the resource's preferred nodes list. This contains the support utilities for the kernel `nfsd` daemon. It is not installed by default. You can use the *YaST > Software > Software Management* tool to find and install the package.

Section 12.9.1, "Sample NFS Load Script," on page 194	Add the following command in the load script before the `exportfs` command:
	
	`# Start the NFS service`
	`exit_on_error rcnfsserver start`

E.3 August 6, 2013

Updates were made to the following section. The changes are explained below.

- Section E.3.1, "Configuring Cluster Policies and Priorities," on page 292

E.3.1 Configuring Cluster Policies and Priorities

Location	Change
Section 8.5, "Configuring Cascade Failover Prevention," on page 110	You can create a configuration file in the `/etc/modprobe.d` file that contains the flag setting to enable and disable the Cascade Failover Prevention function for the server.

E.4 July 25, 2013

Updates were made to the following section. The changes are explained below.

- Section E.4.1, "Configuring Cluster Policies and Priorities," on page 292

E.4.1 Configuring Cluster Policies and Priorities

Location	Change
Section 8.5, "Configuring Cascade Failover Prevention," on page 110	A resource might be quarantined if it is likely responsible for several consecutive server failures, unrelated to interference from failures of other resources.

E.5 May 2013 Scheduled Maintenance

Updates were made to the following sections. The changes are explained below.

- Section E.5.1, "Planning for Novell Cluster Services," on page 292
- Section E.5.2, "What's New or Changed in Novell Cluster Services," on page 293

E.5.1 Planning for Novell Cluster Services

Location	Change
Section 4.5.3, "SLP," on page 61	The SLP refresh interval is now based on the eDirectory advertise-life-time setting.

E.5.2 What's New or Changed in Novell Cluster Services

Location	Change
Section 2.1, "What's New (May 2013 Patches)," on page 29	This section is new.

E.6 May 22, 2013

Updates were made to the following sections. The changes are explained below.

- Section E.6.1, "Managing Clusters," on page 293
- Section E.6.2, "What's New or Changed in Novell Cluster Services," on page 293

E.6.1 Managing Clusters

Location	Change
Section 8.5, "Configuring Cascade Failover Prevention," on page 110	This section is new. Cascading failover prevention is supported in OES 11 and later.

E.6.2 What's New or Changed in Novell Cluster Services

Location	Change
Section 2.10.16, "Cascading Failover Prevention," on page 41	This section is new.

E.7 May 7, 2013

Updates were made to the following section. The changes are explained below.

E.7.1 Troubleshooting Novell Cluster Services

Location	Change
"Migration Authentication Error Occurs when Connecting to a New Cluster Resource on the Target Server" on page 256	This section is new.

E.8   April 2013 Scheduled Maintenance

Updates were made to the following sections. The changes are explained below.

- Section E.8.1, "Troubleshooting Novell Cluster Services," on page 294
- Section E.8.2, "What's New or Changed for Novell Cluster Services," on page 294

E.8.1   Troubleshooting Novell Cluster Services

Location	Change
"Error NFS Stale Handle Reported on Host when Cluster Failover Occurs" on page 254	This section is new.

E.8.2   What's New or Changed for Novell Cluster Services

Location	Change
Section 2.2, "What's New (April 2013 Patches)," on page 30	This section is new.

E.9   April 12, 2013

Updates were made to the following section. The changes are explained below.

- Section E.9.1, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 294

E.9.1   Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
"Creating a Cluster-Enabled Pool and Volume with iManager" on page 180 "Creating a Cluster-Enabled Pool and Volume with NSSMU" on page 182	Ensure that the SAN device is attached to all of the nodes in the cluster. Ensure that you log in to the master node to create clustered pools and volumes. Typically, the pool creation takes less than a minute. The volume creation takes less than 10 seconds. However, if you have a large tree or the server does not hold an eDirectory replica, the create time can take up to 3 minutes.
Chapter 12, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 173	In the `ncpcon mount` command, place the `/opt` switch before the volume information. `ncpcon mount /opt=ns=long <VOLUMENAME>=<VOLUMEID>`

E.10   February 15, 2013

Updates were made to the following section. The changes are explained below.

- Section E.10.1, "What's New or Changed for Novell Cluster Services," on page 295

E.10.1   What's New or Changed for Novell Cluster Services

Location	Change
Section 2.3, "What's New (January 2013 Patches)," on page 30	Added information about Novell iManager 2.7.6 and Novell Client 2 SP3 for Windows.

E.11   January 2013 Scheduled Maintenance

Updates were made to the following section. The changes are explained below.

- Section E.11.1, "What's New or Changed for Novell Cluster Services," on page 295

E.11.1   What's New or Changed for Novell Cluster Services

Location	Change
Section 2.3, "What's New (January 2013 Patches)," on page 30	This section is new. It describes the changes to the Clusters plug-in for Novell iManager 2.7.5, and compares how to perform common tasks in the old and new versions of the plug-in. The comparison of old and new versions of the Clusters plug-in can also be found in "Clusters Plug-In Changes for Novell iManager 2.7.5" in the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide*.

E.12   November 30, 2012

Updates were made to the following section. The changes are explained below.

- Section E.12.1, "Console Commands for Novell Cluster Services," on page 295

E.12.1   Console Commands for Novell Cluster Services

Location	Change
Section A.4, "SBD Utility," on page 269	For an iSCSI SAN, a very small iSCSI device might rarely have uncommon CHS (cylinder-head-sector) geometry values. NSS prefers at least 32 sectors per track.

E.13   October 23, 2012

Updates were made to the following sections. The changes are explained below.

- Section E.13.1, "Managing Clusters," on page 296
- Section E.13.2, "Planning for Novell Cluster Services," on page 296

E.13.1   Managing Clusters

Location	Change
Section 9.14.1, "Prerequisites for Creating an SBD Partition," on page 137	For an iSCSI SAN, a very small iSCSI device might rarely have uncommon CHS (cylinder-head-sector) geometry values. NSS prefers at least 32 sectors per track.

E.13.2   Planning for Novell Cluster Services

Location	Change
Section 4.7.2, "SBD Partitions," on page 70	For an iSCSI SAN, a very small iSCSI device might rarely have uncommon CHS (cylinder-head-sector) geometry values. NSS prefers at least 32 sectors per track.

E.14   September 2012 Scheduled Maintenance

Updates were made to the following sections. The changes are explained below.

- Section E.14.1, "Installing and Configuring Novell Cluster Services," on page 296
- Section E.14.2, "Troubleshooting Novell Cluster Services," on page 297
- Section E.14.3, "What's New or Changed for Novell Cluster Services," on page 297

E.14.1   Installing and Configuring Novell Cluster Services

Location	Change
Section 5.7.2, "Installing the Clusters Plug-In," on page 93 Step 5 in Section 5.7.3, "Uninstalling and Reinstalling the Clusters Plug-In," on page 93	Verify that the `/var/opt/novell/iManager/nps/WEB-INF/lib` directory contains the class loader file `commons-lang-2.6.jar` (or later version), then remove all earlier versions of the file. You must remove the earlier versions (such as `commons-lang.jar` or `commons-lang-2.4.jar`) to ensure that iManager loads the 2.6 or later version of the class loader file.
Section 5.7.4, "Updating Role-Based Services for the Clusters Plug-In for OES 11 SP1," on page 94	This section is new.

E.14.2   Troubleshooting Novell Cluster Services

Location	Change
Section 15.9, "Error 20897 - This node is not a cluster member," on page 254	This section is new.
Section 15.10, "Problems with Clusters Plug-In After Update for Novell iManager 2.7.5," on page 255	This section is new.

E.14.3   What's New or Changed for Novell Cluster Services

Location	Change
Section 2.4, "What's New (September 2012 Patches)," on page 36	This section is new.

E.15   August 31, 2012

Updates were made to the following section. The changes are explained below.

- Section E.15.1, "Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes," on page 297

E.15.1   Configuring and Managing Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 12.13, "Changing the Name Space for a Clustered NSS Volume," on page 200	This section is new.

E.16   August 6, 2012

Updates were made to the following section. The changes are explained below.

* Section E.16.1, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 298

E.16.1   Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
"Renaming a Pool for a Pool Cluster Resource" on page 199	Because renaming involves changing information for the Pool object in eDirectory, you must ensure that the shared pool is active on a cluster node that has its NCP server object in the same context as the Pool object of the pool you are going to rename.
	Added details for the procedure.
Step 4c in "Extending the Size of a LUN" on page 204	Added information to flush and rebuild the maps.
Step 4e in "Extending the Size of a LUN" on page 204	The evms_activate command was moved to follow the verification of multipath maps.
	Corrected the command:
	evms_activate
Step 2d in "Creating and Sharing a New LUN Device" on page 208	Corrected the command:
	evms_activate

E.17   June 27, 2012

Updates were made to the following sections. The changes are explained below.

* Section E.17.1, "Configuring and Managing Cluster Resources," on page 298
* Section E.17.2, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 299

E.17.1   Configuring and Managing Cluster Resources

Location	Change
Section 10.1.2, "Using Parameter Values with Spaces in a Cluster Script," on page 148	This section is new.

E.17.2 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
"Expanding the Size of a Clustered Pool" on page 202	Corrected the procedures for an online pool expansion.

E.18 May 31, 2012

Updates were made to the following section. The changes are explained below.

- Section E.18.1, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 299

E.18.1 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
"NCP Server for Linux" on page 176	NCP Server does not support cross-protocol locks across a cluster migration or failover of the resource.
"Adding NFS Export for a Clustered Pool Resource" on page 193	This section is new.

E.19 April 19, 2012

Updates were made to the following section. The change is explained below.

- Section E.19.1, "Configuring Cluster Policies and Priorities," on page 299
- Section E.19.2, "Files for Novell Cluster Services," on page 300
- Section E.19.3, "Managing Clusters," on page 300
- Section E.19.4, "Planning for Novell Cluster Services," on page 300

E.19.1 Configuring Cluster Policies and Priorities

Location	Change
Section 8.4, "Configuring Cluster Event Email Notification," on page 109	Because both Postfix and the GWIA default to using port 25, you must configure the GWIA to bind exclusively to its resource's secondary IP address in order to avoid a port conflict between Postfix and the GWIA.

E.19.2   Files for Novell Cluster Services

Location	Change
Table B-1 on page 273	References to the hang-check timer were removed. It was obsoleted in OES 2 SP2 Patches. The NCS tolerance is now used to determine whether NCS has been starved of CPU. You can use directive `kernel.panic` in file `/etc/sysctl.conf` file to control the behaviors after NCS has determined to kill the node.

E.19.3   Managing Clusters

Location	Change
Table 9-2, "Cluster Resource States," on page 128	Updated the Alert entry to identify alert types: Start alert, Failover alert, and Failback alert.

E.19.4   Planning for Novell Cluster Services

Location	Change
Section 4.9.3, "Modifying the Polling Interval, No Path Retry, and Failback Settings in the multipath.conf File," on page 74	Added information about the polling_interval.

E.20   January 2012 Scheduled Maintenance

Updates were made to the following sections. The changes are explained below.

- Section E.20.1, "Configuring and Managing Cluster Resources," on page 300
- Section E.20.2, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 301
- Section E.20.3, "Planning for Novell Cluster Services," on page 301
- Section E.20.4, "Upgrading OES 2 Linux Clusters," on page 301
- Section E.20.5, "What's New or Changed for Novell Cluster Services," on page 301

E.20.1   Configuring and Managing Cluster Resources

Location	Change
"Configuring Preferred Nodes for a Resource" on page 160	Changes are not allowed for the Preferred Nodes list for the `Master_IP_Address_Resource`.

E.20.2 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
"NCP Server for Linux" on page 176	NCP Server does not support cross-protocol locks across a cluster migration or failover of the resource.
"Cluster-Enabling an Existing NSS Pool and Its Volumes" on page 184	When you cluster-enable an existing pool, you can automatically deactivate the pool on the master node, and online the resource on a preferred node.

E.20.3 Planning for Novell Cluster Services

Location	Change
"NCP Server for Linux" on page 67	NCP Server does not support cross-protocol locks across a cluster migration or failover of the resource.

E.20.4 Upgrading OES 2 Linux Clusters

Location	Change
"Requirements for Upgrading Clusters" on page 97 "Support for SLES 10 SP4" on page 99	Novell Cluster Services added support for OES 2 SP3 services and file systems on the SUSE Linux Enterprise Server (SLES) 10 SP4 operating system. You can upgrade to SLES 10 SP4 by using the `move-to-sles10-sp4 patch` in the SLES patch channel. After a cluster upgrade, all nodes in the cluster must be running on the same SLES 10 release with the latest patches applied.

E.20.5 What's New or Changed for Novell Cluster Services

Location	Change
"What's New (January 2012 Patches)" on page 36	This section is new.

E.21 December 15, 2011

Updates were made to the following sections. The changes are explained below.

- Section E.21.1, "Configuring and Managing Cluster Resources," on page 302
- Section E.21.2, "Managing Clusters," on page 302
- Section E.21.3, "Security Considerations," on page 302
- Section E.21.4, "Troubleshooting Novell Cluster Services," on page 302

E.21.1 Configuring and Managing Cluster Resources

Location	Change
"Failure Rate" on page 156	The resource log file messages are in the `/var/opt/novell/log/ncs` directory.
"Monitoring Services That Are Critical to Clustering" on page 158	The following line was corrected to use `rcnamcd`: `exit_on_error rcnamcd status`

E.21.2 Managing Clusters

Location	Change
"Deleting a Non-Mirrored Cluster SBD Partition" on page 142	Removed a duplicate step.

E.21.3 Security Considerations

Location	Change
"Log Files" on page 262	This section was updated for completeness. The resource log file messages are in the `/var/opt/novell/log/ncs` directory.

E.21.4 Troubleshooting Novell Cluster Services

Location	Change
"Why did a resource not migrate to the node I specified?" on page 253	This section is new.

E.22 December 6, 2011

Updates were made to the following section. The changes are explained below.

◆ Section E.22.1, "Planning for Novell Cluster Services," on page 303

E.22.1 Planning for Novell Cluster Services

Location	Change
"Modifying the Polling Interval, No Path Retry, and Failback Settings in the multipath.conf File" on page 74	This section was updated for clarity. Added a recommendation to use failback "manual".

E.23 November 2011 Scheduled Maintenance

Updates were made to the following sections. The changes are explained below.

◆ Section E.23.1, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 303
◆ Section E.23.2, "Configuring and Managing Cluster Resources for Shared Linux POSIX Volumes," on page 303
◆ Section E.23.3, "Managing Clusters," on page 304
◆ Section E.23.4, "What's New or Changed for Novell Cluster Services," on page 304

E.23.1 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
"Configuring a Monitor Script for the Shared NSS Pool" on page 191	The CIFS `monitor` command is available in the September 2011 Scheduled Maintenance patch. It is the default command used in the monitor script for new NSS pool cluster resources that have CIFS enabled as an advertising protocol.

E.23.2 Configuring and Managing Cluster Resources for Shared Linux POSIX Volumes

Location	Change
"Expanding EVMS Volumes on Shared Disks" on page 234	The following steps are new: Ensure that the target disk contains only unpartitioned free space. Re-mount the volume to be expanded.

E.23.3 Managing Clusters

Location	Change
"Prerequisites for Creating an SBD Partition" on page 137	This section was expanded for clarity. The path for the SBD partition is: `/dev/evms/.nodes/<cluster_name>.sbd`
"Creating a Non-Mirrored Cluster SBD Partition" on page 138	Ensure that the device you want to use has been initialized and marked as *Shareable for Clustering* as described in Section 9.14.1, "Prerequisites for Creating an SBD Partition," on page 137.
"Using SBDUTIL to Create a Mirrored Cluster SBD Partition" on page 140 "Using EVMSGUI to Create a Mirrored Cluster SBD Partition" on page 141	Ensure that the devices you want to use have been initialized and marked as *Shareable for Clustering* as described in Section 9.14.1, "Prerequisites for Creating an SBD Partition," on page 137.

E.23.4 What's New or Changed for Novell Cluster Services

Location	Change
"What's New (November 2011 Patches)" on page 37	This section is new.

E.24 August 2011 Scheduled Maintenance

The OES 2 SP3 August 2011 Maintenance Patch provides support for OES 2 SP3 on the SUSE Linux Enterprise 10 SP4 operating system. Links were updated to point to the latest SLES 10 documentation. In addition, updates were made to the following section. The changes are explained below.

- Section E.24.1, "What's New or Changed for Novell Cluster Services," on page 304

E.24.1 What's New or Changed for Novell Cluster Services

Location	Change
"What's New (August 2011 Patches)" on page 37	This section is new.

E.25 July 22, 2011

Updates were made to the following section. The changes are explained below.

◆ Section E.25.1, "Managing Clusters," on page 305

E.25.1 Managing Clusters

Location	Change
"Starting and Stopping Novell Cluster Services" on page 125	If you need to restart adminfs, you must stop Novell Cluster Services before you stop adminfs, or you can reboot the server.

E.26 June 24, 2011

Updates were made to the following sections. The changes are explained below.

◆ Section E.26.1, "Configuring and Managing Cluster Resources," on page 305
◆ Section E.26.2, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 305
◆ Section E.26.3, "Managing Clusters," on page 306
◆ Section E.26.4, "Planning for Novell Cluster Services," on page 306
◆ Section E.26.5, "Upgrading OES 1 Linux Clusters to OES 2 Linux," on page 306

E.26.1 Configuring and Managing Cluster Resources

Location	Change
"Synchronizing Locally Modified Resourse Templates with eDirectory" on page 152	This section is new.

E.26.2 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
"Cluster-Enabling an Existing NSS Pool and Its Volumes" on page 184	If you plan to add CIFS and AFP as advertising protocols when cluster-enabling an existing pool, the Novell CIFS and Novell AFP software must be installed and running before you begin the enablement process. Their daemons must also be running before bringing the resource online.

E.26.3 Managing Clusters

Location	Change
"Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster" on page 116	If you list multiple LDAP servers, order them based on what is local to the cluster then the closest physical read/write replica.

E.26.4 Planning for Novell Cluster Services

Location	Change
"LDAP Server List" on page 61	This section is new.

E.26.5 Upgrading OES 1 Linux Clusters to OES 2 Linux

The step to use the `cluster convert` command was removed. This is not necessary when upgrading between Linux versions.

E.27 May 2011 Scheduled Maintenance

Updates were made to the following sections. The changes are explained below.

- Section E.27.1, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 306
- Section E.27.2, "Console Commands for Novell Cluster Resources," on page 306
- Section E.27.3, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 307
- Section E.27.4, "What's New or Changed for Novell Cluster Services," on page 307

E.27.1 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
"Naming Conventions" on page 175	This section is new. It is not a new behavior. It was previously undocumented.

E.27.2 Console Commands for Novell Cluster Resources

Location	Change
"MIGRATE {resource} {node}" on page 265 "ONLINE {resource}{node name}" on page 265	For the `cluster online` and `cluster migrate` commands, the node that you specify must be running in the cluster and also be in the resource's *Assigned Nodes* list.

E.27.3 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
"SBD Partitions" on page 70	An SBD must be created before you attempt to create file system cluster resources.

E.27.4 What's New or Changed for Novell Cluster Services

Location	Change
"What's New (May 2011 Patches)" on page 38	This section is new.

E.28 April 2011 Scheduled Maintenance

Updates were made to the following sections. The changes are explained below.

- Section E.28.1, "Configuring and Managing Cluster Resources," on page 307
- Section E.28.2, "What's New or Changed in Novell Cluster Services," on page 307

E.28.1 Configuring and Managing Cluster Resources

Location	Change
Section 10.1.6, "Number of Resources," on page 149	This section is new.

E.28.2 What's New or Changed in Novell Cluster Services

Location	Change
"What's New (April 2011 Patches)" on page 38	This section is new.

E.29 March 18, 2011

Updates were made to the following sections. The changes are explained below.

- Section E.29.1, "Configuring and Managing Cluster Resources," on page 308
- Section E.29.2, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 308
- Section E.29.3, "Console Commands for Novell Cluster Services," on page 308
- Section E.29.4, "Installing and Configuring Novell Cluster Services on OES 2 Linux," on page 308

E.29.1   Configuring and Managing Cluster Resources

Location	Change
Step 6 on page 154 Step 6 on page 154 "Timeout Value" on page 156	The timeout values for load, unload, and monitoring scripts are applied only when the resource is migrated to another node. They are not used during resource online/offline procedures.

E.29.2   Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 12.7, "Configuring a Monitor Script for the Shared NSS Pool," on page 191	Added a sample monitoring script.
"Renaming a Pool for a Pool Cluster Resource" on page 199	Added a step to use `cluster rename` to apply a customized Pool Resource name.
"Adding a Volume to a Clustered Pool" on page 200	This section is new.
"Expanding the Size of a Clustered Pool" on page 202	This section is new.

E.29.3   Console Commands for Novell Cluster Services

Location	Change
"RESTART [seconds]" on page 266	The cluster leave process begins immediately. Specify a value of 60 or more seconds as the time to wait before the cluster join begins.

E.29.4   Installing and Configuring Novell Cluster Services on OES 2 Linux

Location	Change
Section 5.7, "Installing or Updating the Clusters Plug-in for iManager," on page 92	This section was relocated from Section 4.5.7, "Storage-Related Plug-Ins for iManager," on page 64. Step 4 is new.

E.29.5 Managing Clusters

Location	Change
"Deleting a Cluster Node from a Cluster" on page 135	Modified for accuracy. Reconfiguring the node for a different cluster is not supported in YaST at this time.

E.29.6 Planning for a Cluster

Location	Change
"Cluster OU Partitioning and Replication" on page 51	If you do not want to put a replica of eDirectory on the node, you must configure one or multiple LDAP servers for the node to use. The LDAP servers must have a master replica or a Read/Write replica of eDirectory. For information about how to modify the LDAP server list that is used by a cluster, see Section 8.9.1, "Changing the Administrator Credentials or LDAP Server IP Addresses for a Cluster," on page 116.

E.30 December 2010 (OES 2 SP3)

Updates were made to the following sections. The changes are explained below.

- Section E.30.1, "Configuring and Managing Cluster Resources," on page 310
- Section E.30.2, "Configuring and Managing Cluster Resources for Shared Linux POSIX Volumes," on page 310
- Section E.30.3, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 310
- Section E.30.4, "Configuring Cluster Policies and Priorities," on page 311
- Section E.30.5, "Console Commands for Novell Cluster Services," on page 311
- Section E.30.6, "Converting NetWare 6.5 Clusters to OES 2 Linux," on page 311
- Section E.30.7, "Installing and Configuring Novell Cluster Services on OES 2 Linux," on page 311
- Section E.30.8, "Managing Clusters," on page 311
- Section E.30.9, "Overview of Novell Cluster Services," on page 312
- Section E.30.10, "Planning for a Cluster," on page 312
- Section E.30.11, "Planning for Novell Cluster Services," on page 312
- Section E.30.12, "Upgrading OES 2 Linux Clusters," on page 312
- Section E.30.13, "Upgrading OES 1 Linux Clusters to OES 2 Linux Clusters," on page 313
- Section E.30.14, "What's New or Changed for Novell Cluster Services," on page 313

E.30.1 Configuring and Managing Cluster Resources

Location	Change
Section 10.1.4, "Script Length Limits," on page 148	This section is new. A script can contain up to 3200 bytes of information.
Step 8 in Section 10.6.2, "Configuring Resource Monitoring," on page 157	This section is new. The *Reboot the hosting node (without synchronizing or unmounting the disks)* option is new for OES 2 SP3.
"Configuring Resource Mutual Exclusion Groups" on page 162	This feature is new for OES 2 SP3.
Step 7 in Section 10.8, "Configuring Preferred Nodes for a Resource," on page 160	Beginning in OES 2 SP3, an edit option provides a text editor where you can set the preferred failover order of the nodes assigned to a resource.
Step 5 in Section 10.9, "Configuring Resource Priorities for Load Order," on page 161	Beginning in OES 2 SP3, an edit option provides a text editor where you can set the load order priorities for resources in the cluster.
"Renaming a Cluster Resource" on page 166	This section is new.

E.30.2 Configuring and Managing Cluster Resources for Shared Linux POSIX Volumes

Location	Change
"Creating a Virtual Server Object for the Linux POSIX Volume Cluster Resource" on page 228	This section was modified to correct errors in the procedure.

E.30.3 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
"Cluster-Enabling an Existing NSS Pool and Its Volumes" on page 184	The pool to be cluster enabled must be activated on the master node in the cluster before you begin to cluster enable the pool. This allows the pool and volume information to be automatically added to the pool cluster resource load, unload, and monitor scripts. The volume entry will be automatically removed from the /etc/fstab so that the pool cluster resource load script controls when the volume is mounted. The procedure has been modified with the related details.

E.30.4   Configuring Cluster Policies and Priorities

Location	Change
"Configuring Quorum Membership and Timeout Properties" on page 107	If you change the values, you must restart Novell Cluster Services to activate the updated settings.

E.30.5   Console Commands for Novell Cluster Services

Location	Change
"CONVERT preview [resource]" on page 264	Beginning in OES 2 SP3, you can preview all resources before you complete the cluster conversion by not specifying a resource.
"DOWN" on page 264	Beginning in OES 2 SP3, you are prompted to confirm the cluster down.
"RENAME <old_resource_name> <new_resource_name>" on page 265	The cluster rename option is new for OES 2 SP3.
"RESTART [seconds]" on page 266	Beginning in OES 2 SP3, you are prompted to confirm the cluster restart. Valid values: 60 seconds or greater. The default is 60 seconds.

E.30.6   Converting NetWare 6.5 Clusters to OES 2 Linux

Information in this chapter was moved to the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*.

E.30.7   Installing and Configuring Novell Cluster Services on OES 2 Linux

Location	Change
"OES Common Proxy User" on page 57	This section is new. This OES Common Proxy User option is new in OES 2 SP3.
"Patching Novell Cluster Services" on page 94	This section is new.

E.30.8   Managing Clusters

Location	Change
"Generating a Cluster Configuration Report" on page 129	The Monitoring script is included in the report for each resource. The Mutually Exclusive Resource Groups information is included in the report.

Location	Change
"Removing All Nodes at the Same Time" on page 132	This section is new.
"Deleting a Non-Mirrored Cluster SBD Partition" on page 142	Stop Novell Cluster Services before you delete the SBD partition, then restart it after a new SBD partition exists.
"Deleting a Mirrored Cluster SBD Partition" on page 142	Stop Novell Cluster Services before you delete the software RAID that contains the SBD partition, then restart it after a new SBD partition exists.
"Removing a Segment from a Mirrored Cluster SBD Partition" on page 143	Enter cluster maintenance mode instead of bringing the cluster down.

E.30.9 Overview of Novell Cluster Services

Location	Change
"Terminology" on page 23	This section is new.

E.30.10 Planning for a Cluster

Location	Change
"Planning for a Cluster" on page 47	This section is new.

E.30.11 Planning for Novell Cluster Services

This chapter describes the installation requirements for Novell Cluster Services section. This information was previously located in the Installing Novell Cluster Services section.

In addition, the following changes were made to the requirements information:

Location	Change
"CASA" on page 66	This section is new.

E.30.12 Upgrading OES 2 Linux Clusters

Location	Change
Section 6.1, "Requirements for Upgrading Clusters," on page 97	Novell Cluster Services supports mixed-nodes running 32-bit and 64-bit OES 2 Linux during a rolling cluster upgrade.

Location	Change
Section 6.2, "Upgrading OES 2 Clusters (Rolling Cluster Upgrade)," on page 97	Stop Novell Cluster Services during the upgrade.
Section 6.5, "Upgrade Issues for OES 2 SP3," on page 99	This section is new.

E.30.13 Upgrading OES 1 Linux Clusters to OES 2 Linux Clusters

Location	Change
Section 6.2, "Upgrading OES 2 Clusters (Rolling Cluster Upgrade)," on page 97	Stop Novell Cluster Services during the upgrade.

E.30.14 What's New or Changed for Novell Cluster Services

Location	Change
"What's New (OES 2 SP3)" on page 39	This section is new.

E.31 August 2, 2010

Updates were made to the following sections. The changes are explained below.

- Section E.31.1, "Configuring Cluster Resources for Shared Linux POSIX Volumes," on page 313
- Section E.31.2, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 314
- Section E.31.3, "Managing Clusters," on page 314

E.31.1 Configuring Cluster Resources for Shared Linux POSIX Volumes

Location	Change
"Configuring a Load Script for the Linux POSIX Volume Cluster Resource" on page 223	Step 3 on page 224 describes how to use mount options in the load script.
"Creating a Virtual Server Object for the Linux POSIX Volume Cluster Resource" on page 228	This section was modified to clarify the NCS:NCP Server object is created under the Cluster object for the cluster where you created the shared Linux POSIX volume.

E.31.2 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
"Novell Storage Services" on page 174 "Guidelines for Using Pool Cluster Resources" on page 177	The Server, Pool, Volume, Cluster Resource, and Cluster objects are recommended to be in the same context (such as ou=ncs,o=novell). If they are in different contexts, you might need to cluster migrate the pool cluster resource back to the node where the pool was created.

E.31.3 Managing Clusters

Location	Change
"OES 2 SP2 with Patches and Later" on page 134	To allow an automatic reboot, set the `kernel.panic` token to value to `1` in the `/etc/sysctl.conf` file.

E.32 June 2010 Scheduled Maintenance

Updates were made to the following sections. The changes are explained below.

E.32.1 Configuring and Managing Cluster Resources

Location	Change
"Naming Conventions for Cluster Resources" on page 148	This section is new.

E.32.2 Configuring Cluster Policies and Priorities

Location	Change
"Configuring Cluster Event Email Notification" on page 109	The subject line provides information about the cluster name, resource name, action taken, and node name.

E.32.3 Console Commands for Novell Cluster Services

Location	Change
"SCAN FOR NEW DEVICES" on page 266	You can use the rescan-scsi-bus.sh script to scan for the new devices on Linux without rebooting. For information, see "Scanning for New Devices without Rebooting" (http://www.suse.com/documentation/sles10/stor_admin/data/scandev.html) in the *SLES 10 SP4: Storage Administration Guide* (http://www.suse.com/documentation/sles10/stor_admin/data/bookinfo.html).

E.32.4 Installing and Configuring Novell Cluster Services on OES 2 Linux

Location	Change
"SBD Partitions" on page 70	When you mark a device as *Shareable for Clustering,* the share information is added to the disk and consumes about 4 MB. This space is used in addition to the space needed for the SBD. For example, if you create a 1 GB (1024 MB) device for the SBD, 4 MB are used for the share information, leaving 1020 MB as available free space for the SBD partition.

E.32.5 Managing Clusters

Location	Change
"Generating a Cluster Configuration Report" on page 129	The Monitoring script is included in the report for each resource.
"Generating a Cluster Configuration Report" on page 129	Step 6 provides instructions for how to save the report to a file.
"Preventing a Cluster Node Reboot after a Node Shutdown" on page 134	Novell Cluster Services has been modified to handle panic situations for a node according to the Linux kernel panic setting on the node.
"Creating a Non-Mirrored Cluster SBD Partition" on page 138	In Step 1 on page 137, added the -s option to the command and example. Step 5 on page 138 is required only if Cluster Services was initially installed without an SBD partition or mirrored SBD partitions. However, it does no harm to verify that the *NCS: Shared Disk Flag* attribute is enabled.

Location	Change
"Using SBDUTIL to Create a Mirrored Cluster SBD Partition" on page 140	In Step 4 on page 140, added the -s option to the command and example. Step 5 on page 140 is required only if Cluster Services was initially installed without an SBD partition or mirrored SBD partitions. However, it does no harm to verify that the *NCS: Shared Disk Flag* attribute is enabled.

E.32.6 Troubleshooting Novell Cluster Services

Location	Change
"Diagnosing a Poison Pill Condition" on page 252	This section is new.
"Running supportconfig Concurrently on Multiple Nodes Causes Cluster Problems" on page 259	This section is new.

E.32.7 Upgrading OES 2 Linux Clusters

Location	Change
"Updating the iFolder Resource Template" on page 99	You must run a script one of the nodes (master node is preferred) in the cluster after upgrading to SP2. The template changes take effect without rebooting the node or restarting the cluster.

E.32.8 What's New or Changed for Novell Cluster Services

Location	Change
"What's New (June 2010 Patches)" on page 42	This section is new.

E.33 March 15, 2010

Updates were made to the following sections. The changes are explained below.

E.33.1 Configuring and Managing Cluster Resources

Location	Change
Section 10.1.5, "Planning Cluster Maintenance," on page 149	This section is new.
Section 10.2.2, "Creating a Resource Template," on page 150	Cluster Services marks the load, unload, or monitor process as failed right after the defined timeout expires, but it must wait for the process to conclude before it can start other resource operations.
Section 10.4, "Configuring a Load Script for a Cluster Resource," on page 153	
Section 10.5, "Configuring an Unload Script for a Cluster Resource," on page 154	
Section 10.6.2, "Configuring Resource Monitoring," on page 157	

E.33.2 Configuring Cluster Resources for Shared Linux POSIX Volumes

Location	Change
Section 13.4.3, "Configuring a Load Script for the Linux POSIX Volume Cluster Resource," on page 223	Cluster Services marks the load, unload, or monitor process as failed right after the defined timeout expires, but it must wait for the process to conclude before it can start other resource operations.
Section 13.4.4, "Configuring an Unload Script for a Linux POSIX Volume Cluster Resource," on page 224	
Section 13.4.5, "Enabling Monitoring and Configuring a Monitor Script for a Linux POSIX Volume Cluster Resource," on page 225	

E.33.3 Configuring Novell Cluster Services in a Xen Virtualization Environment

Location	Change
"Viewing or Modifying the Monitor Script" on page 244	Cluster Services marks the monitor process as failed right after the defined timeout expires, but it must wait for the process to conclude before it can start other resource operations.

E.33.4 Troubleshooting Novell Cluster Services

Location	Change
"File Location Information" on page 252	This section is new.

E.34 February 19, 2010

Updates were made to the following sections. The changes are explained below.

- Section E.34.1, "Installing and Configuring Novell Cluster Services on OES 2 Linux," on page 318
- Section E.34.2, "Managing Clusters," on page 318

E.34.1 Installing and Configuring Novell Cluster Services on OES 2 Linux

Location	Change
"Novell iManager 2.7.4" on page 63	The OES 2 SP2 Linux release contains a Clusters plug-in that is required when using Novell Business Continuity Clustering 1.2.1 for OES 2 SP2 Linux. For information about BCC 1.2.1, see *BCC 1.2.1: Administration Guide for OES 2 SP2 Linux* (http://www.novell.com/documentation/bcc/bcc121_admin_lx/data/bookinfo.html).

E.34.2 Managing Clusters

Location	Change
Section 9.15, "Customizing Cluster Services Management," on page 144	The references to /_admin should be /admin on Linux..

E.35   February 10, 2010

Updates were made to the following section. The changes are explained below.

◆ Section E.35.1, "Configuring Cluster Resources for Linux POSIX Volumes," on page 319

E.35.1   Configuring Cluster Resources for Linux POSIX Volumes

Location	Change
Section 13.6, "Creating a Virtual Server Object for the Linux POSIX Volume Cluster Resource," on page 228	This section is new.

E.36   January 29, 2010 (Maintenance Patch)

Updates were made to the following sections. The changes are explained below.

◆ Section E.36.1, "Installing and Configuring Novell Cluster Services on OES 2 Linux," on page 319
◆ Section E.36.2, "Managing Clusters," on page 320
◆ Section E.36.3, "Troubleshooting Novell Cluster Services," on page 320
◆ Section E.36.4, "What's New," on page 320

E.36.1   Installing and Configuring Novell Cluster Services on OES 2 Linux

Location	Change
"Novell eDirectory 8.8.6" on page 59	You can install an eDirectory master replica or replica in the cluster, but it is not required to do so for Novell Cluster Services. (eDirectory must be installed somewhere in the same tree.)
"SBD Partitions" on page 70	If you plan to mirror the SBD, you must also initialize the partition you want to use as the mirrored SBD, and mark its device as shareable for clustering.
"Configuring a New Cluster" on page 88	For each SBD partition, you must have at least 20 MB of free space on a device that has been initialized and marked as shareable for clustering.
"Adding a Node to an Existing Cluster" on page 90	Step 1 on page 90 is new. If you have a shared disk system attached to your cluster servers, an SBD partition is required and must be created before you configure the second node in the cluster.

E.36.2 Managing Clusters

Location	Change
"Generating a Cluster Configuration Report" on page 129	.This section is new.

E.36.3 Troubleshooting Novell Cluster Services

Location	Change
Section 15.4, "Diagnosing Cluster Problems," on page 252	Cluster configuration. An HTML report created by the *Clusters > Cluster Manager > Run Report* option in the Clusters plug-in to iManager.
"Cluster Resource Is Stuck in "NDS Sync" or "eDirectory Sync" State" on page 258	.This section is new.

E.36.4 What's New

Location	Change
"What's New (January 2010 Patches)" on page 42	.This section is new.

E.37 January 20, 2010

Updates were made to the following section. The changes are explained below.

E.37.1 What's New or Changed for Novell Cluster Services

Location	Change
"Cluster Restart Is No Longer Required in a Rolling Cluster Upgrade" on page 44	.This section is new.

E.38 January 4, 2010

Updates were made to the following sections. The changes are explained below.

- Section E.38.1, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 321

E.38.1 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
"Creating a Cluster-Enabled Pool and Volume with iManager" on page 180	In Step 3g on page 180, if you choose to cluster-enable the pool later, you must create a volume on the shared pool before you attempt to cluster-enable the pool.
"Creating a Cluster-Enabled Pool and Volume with NSSMU" on page 182	In Step 4g on page 183, if you choose to cluster-enable the pool later, you must create a volume on the shared pool before you attempt to cluster-enable the pool.
"Cluster-Enabling an Existing NSS Pool and Its Volumes" on page 184	The pool must contain at least one volume before you attempt to cluster-enable the existing pool.

E.39 December 15, 2009

Updates were made to the following section. The changes are explained below.

- Section E.39.1, "Installing and Configuring Novell Cluster Services on OES 2 Linux," on page 321

E.39.1 Installing and Configuring Novell Cluster Services on OES 2 Linux

Location	Change
Section 5.2.2, "Extending the Schema," on page 79	Added Step 3 to create the `/var/opt/novell/install` directory..

E.40 December 10, 2009

Updates were made to the following sections. The changes are explained below.

- Section E.40.1, "Installing and Configuring Novell Cluster Services on OES 2 Linux," on page 322
- Section E.40.2, "Managing Clusters," on page 322
- Section E.40.3, "Troubleshooting Novell Cluster Services," on page 322

E.40.1   Installing and Configuring Novell Cluster Services on OES 2 Linux

Location	Change
Section 5.2.2, "Extending the Schema," on page 79	The procedure has been modified to create a configuration file and use it with the `/opt/novell/ncs/install/ncs_install.py` script to extend the schema..

E.40.2   Managing Clusters

Location	Change
Section 9.14, "Creating or Deleting Cluster SBD Partitions," on page 136	The procedures in this section were reorganized and modified for clarity.
Section 9.14.7, "Removing a Segment from a Mirrored Cluster SBD Partition," on page 143	This section is new.

E.40.3   Troubleshooting Novell Cluster Services

Location	Change
Section 15.21, "Is there a way to uninstall Novell Cluster Services from a server?," on page 258	This section is new.

E.41   November 2009 (OES 2 SP2)

Updates were made to the following sections. The changes are explained below.

- Section E.41.1, "Configuring Novell Cluster Services in a Xen Host Environment," on page 323
- Section E.41.2, "Console Commands for Novell Cluster Services," on page 323
- Section E.41.3, "Installing Novell Cluster Services on OES 2 Linux," on page 323
- Section E.41.4, "Managing Clusters," on page 323
- Section E.41.5, "Troubleshooting Novell Cluster Services," on page 324
- Section E.41.6, "Upgrading OES 2 Linux Clusters," on page 324
- Section E.41.7, "What's New," on page 324

E.41.1 Configuring Novell Cluster Services in a Xen Host Environment

Location	Change
Section 14.1, "Prerequisites for Xen Host Server Environments," on page 237	This section is new.

E.41.2 Console Commands for Novell Cluster Services

Location	Change
Section A.4, "SBD Utility," on page 269	This section is new.

E.41.3 Installing Novell Cluster Services on OES 2 Linux

Location	Change
"Modifying the Retry Setting for the modprobe.conf.local File" on page 73	For OES 2 SP2 and later, or if you have installed the latest kernel and qla-driver, use the following setting in the `/etc/modprobe.conf.local` file: `options qla2xxx qlport_down_retry=2`
"No Path Retry" on page 74	For no_path_retry, a value of `fail` has the same meaning as a value of 0, and it is more easily understood. `no_path_retry fail`

E.41.4 Managing Clusters

Location	Change
Section 9.14.3, "Creating a Non-Mirrored Cluster SBD Partition," on page 138	Use the EVMSGUI (or EVMSN or EVMS) tool to check the names of the devices if needed, and only use the base (leaf) names with the -d option.
"Using SBDUTIL to Create a Mirrored Cluster SBD Partition" on page 140	This section is new.

E.41.5 Troubleshooting Novell Cluster Services

Location	Change
Section 15.12, "Error 603 - smdr.novell is not registered with SLP for a new cluster resource," on page 255	This section is new.

E.41.6 Upgrading OES 2 Linux Clusters

Location	Change
Section 6.4, "Upgrade Issues for OES 2 SP2," on page 99	This section is new.

E.41.7 What's New

Location	Change
Section 2.13, "What's New (OES 2 SP2)," on page 42	This section is new.

E.42 July 30, 2009

Updates were made to the following sections. The changes are explained below.

- Section E.42.1, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 324
- Section E.42.2, "Console Commands for Novell Cluster Services," on page 325
- Section E.42.3, "Installing Novell Cluster Services on OES 2 Linux," on page 325
- Section E.42.4, "Managing Clusters," on page 325
- Section E.42.5, "Security Considerations," on page 325
- Section E.42.6, "Troubleshooting Novell Cluster Services," on page 325
- Section E.42.7, "What's New or Changed for Novell Cluster Services," on page 326

E.42.1 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 12.2, "Guidelines for Using Pool Cluster Resources," on page 177	Before you rename a cluster-enabled pool, ensure that you offline the pool resource, activate the pool by using iManager or NSSMU, rename the pool, then online the pool resource.

Location	Change
"Mirroring and Cluster-Enabling Shared NSS Pools and Volumes" on page 194	This section was updated to correct errors in the procedures.

E.42.2 Console Commands for Novell Cluster Services

Location	Change
"Business Continuity Clustering Commands" on page 268	This section is new.

E.42.3 Installing Novell Cluster Services on OES 2 Linux

Location	Change
"Cluster Installation Administrator" on page 55	Added information about how to set the rights that the user needs for installing Novell Cluster Services.
"Web Browser" on page 66	Added information about character encoding settings using Unicode (UTF-8).
"Multipath I/O Configuration Requirements" on page 73	This section is new.

E.42.4 Managing Clusters

Location	Change
Section 9.13, "Deleting a Cluster Node from a Cluster," on page 135	Corrected a typographical error. Step 4 on page 136 should read "remove its NCS attributes", not NCP attributes.

E.42.5 Security Considerations

This section is new.

E.42.6 Troubleshooting Novell Cluster Services

Location	Change
Section 15.4, "Diagnosing Cluster Problems," on page 252	This section is new.

Location	Change
Section 15.16, "Could Not Delete This Resource Data_Server (Error 499)," on page 257	This section is new.

E.42.7 What's New or Changed for Novell Cluster Services

Location	Change
Section 2.14.4, "Attribute NCS: GIPC Config Is No Longer Maintained," on page 45	This section is new.

E.43 June 22, 2009

Updates were made to the following section. The changes are explained below.

◆ Section E.43.1, "Managing Clusters," on page 326

E.43.1 Managing Clusters

Location	Change
Section 9.13, "Deleting a Cluster Node from a Cluster," on page 135	Step 4 on page 136 is new.

E.44 June 5, 2009

Updates were made to the following sections. The changes are explained below.

◆ Section E.44.1, "Configuring Cluster Policies and Priorities," on page 326
◆ Section E.44.2, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 327
◆ Section E.44.3, "Console Commands for Novell Cluster Services," on page 327
◆ Section E.44.4, "Installing Novell Cluster Services on OES 2 Linux," on page 327
◆ Section E.44.5, "Managing Clusters," on page 328

E.44.1 Configuring Cluster Policies and Priorities

Location	Change
Section 8.3.2, "Tolerance," on page 108	Added examples for the Tolerance setting.

E.44.2 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 12.4, "Cluster-Enabling an Existing NSS Pool and Its Volumes," on page 184	Comment out (or remove) the volume's entry in the `/etc/fstab` file. The load and unload scripts that are created when you cluster-enable the pool will be responsible for mounting and dismounting the volume after the pool is cluster enabled.
Section 12.5, "Configuring a Load Script for the Shared NSS Pool," on page 189	Added information about using the `/opt=ns=<long|unix|mac|dos>` switch to specify the name space.

E.44.3 Console Commands for Novell Cluster Services

Location	Change
Section A.1, "Cluster Management Commands," on page 263	The `cluster scan for new devices` command is obsolete on Linux. Use Linux commands to scan for new devices.

E.44.4 Installing Novell Cluster Services on OES 2 Linux

Location	Change
Section 4.7, "Shared Disk Configuration Requirements," on page 70	The following statement was added for clarification: **IMPORTANT:** The cluster SBD partition is not required unless you have shared storage.
Section 5.4.1, "Installing Novell Cluster Services during a OES 2 Linux Installation," on page 81	Added clarification that creating an SBD is necessary only if the cluster has a shared disk system.
Section 5.4.2, "Installing Novell Cluster Services on an Existing OES 2 Linux Server," on page 83	In Step 1 on page 83, added clarification that creating an SBD is necessary only if the cluster has a shared disk system.
Section 5.5.2, "Opening the Novell Open Enterprise Server Configuration Page," on page 86	Added clarification that creating an SBD is necessary only if the cluster has a shared disk system.

E.44.5 Managing Clusters

Location	Change
Section 9.14.3, "Creating a Non-Mirrored Cluster SBD Partition," on page 138	Added clarification that creating an SBD is necessary only if the cluster has a shared disk system. Added an example for the sbdutil utility.
Section 9.14.4, "Creating a Mirrored Cluster SBD Partition," on page 140	Added clarification that creating an SBD is necessary only if the cluster has a shared disk system.

E.45 May 6, 2009

Updates were made to the following sections. The changes are explained below.

- Section E.45.1, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 328
- Section E.45.2, "Configuring and Managing Cluster Resources," on page 328
- Section E.45.3, "Installing Novell Cluster Services on OES 2 Linux," on page 329
- Section E.45.4, "Troubleshooting Novell Cluster Services," on page 329
- Section E.45.5, "Upgrading OES 2 Linux Servers," on page 329

E.45.1 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 12.8, "Adding Advertising Protocols for NSS Pool Cluster Resources," on page 191	Step 9a was modified to clarify that you can specify a different name for the Novell CIFS Server only when you create the share. Use the CIFS iManager plug-in to manage the name of existing CIFS shares.

E.45.2 Configuring and Managing Cluster Resources

Location	Change
Section 10.2.2, "Creating a Resource Template," on page 150	The following sentence in Step 5 was modified for clarification: If the script does not complete within the specified time, the resource becomes comatose when migrating to another node.
Section 10.5, "Configuring an Unload Script for a Cluster Resource," on page 154	The following sentence in Step 6 was modified for clarification: If the script does not complete within the specified time, the resource becomes comatose when migrating to another node.
Section 10.6.4, "Monitoring Services That Are Critical to Clustering," on page 158	This section is new.

Location	Change
Section 10.7, "Configuring the Start, Failover, and Failback Modes for Cluster Resources," on page 159	The following statement was added for clarification. There is no behavior change. **IMPORTANT:** Cluster Services works with NCP user connections so that the user data sessions are resumed after failover. However, non-NCP users might experience service interruption and need to reconnect to the server after the failover. Client applications using server based storage must be restarted even with NCP unless the applications are NCP reconnect aware.

E.45.3 Installing Novell Cluster Services on OES 2 Linux

Location	Change
Section 5.2, "Extending the eDirectory Schema to Add Cluster Objects," on page 78	Removed references to `ncpserver.ldif`. This file is a by-product of the extension and is not part of the setup.

E.45.4 Troubleshooting Novell Cluster Services

Location	Change
Section 15.20, "Where Is the 'Prevent Cascading Failovers' Option?," on page 258	This section is new.

E.45.5 Upgrading OES 2 Linux Servers

Location	Change
Section 6.2, "Upgrading OES 2 Clusters (Rolling Cluster Upgrade)," on page 97	You can also add OES 2 SP1 Linux servers (or later versions) to the OES 2 Linux cluster, and remove the old OES 2 Linux servers from the cluster. Ensure that any services that are available only in OES 2 SP1 (such as Novell CIFS or Novell AFP) are set up with preferred nodes for failover that are running OES 2 SP1. Several guidelines are provided about the mixed-version cluster.

E.46 March 3, 2009

Updates were made to the following sections. The changes are explained below.

- Section E.46.1, "Configuring and Managing Cluster Resources," on page 330
- Section E.46.2, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 330
- Section E.46.3, "Managing Clusters," on page 330
- Section E.46.4, "What's Next," on page 330

E.46.1   Configuring and Managing Cluster Resources

Location	Change
Section 10.14, "Deleting Cluster Resources," on page 167	The procedure was modified for clarification and accuracy.

E.46.2   Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 12.4, "Cluster-Enabling an Existing NSS Pool and Its Volumes," on page 184	Before enabling clustering for a pool and its volumes, ensure that they have Storage objects in the eDirectory tree.

E.46.3   Managing Clusters

Location	Change
Step 5 in Section 9.13, "Deleting a Cluster Node from a Cluster," on page 135	If you are converting a cluster from NetWare to Linux, you must restart the cluster instead so that clstrlib.ko is reloaded: `rcnovell-ncs restart`

E.46.4   What's Next

Location	Change
Section 2.14.3, "Behavior Change for Adding a Node," on page 44	**IMPORTANT:** A Novell Cluster Services patch is available in the patch channel and on the Novell Downloads Web site (http://www.novell.com/downloads) that allows you to add a new node seamlessly again on a Linux server. For cluster conversions, a cluster restart is still necessary after all NetWare nodes have been removed from the cluster.
Section 2.14.3, "Behavior Change for Adding a Node," on page 44	If you are converting a cluster from NetWare to Linux, you must restart the cluster instead so that clstrlib.ko is reloaded: `rcnovell-ncs restart`

E.47   February 3, 2009

Updates were made to the following sections. The changes are explained below.

E.47.1 Configuring and Managing Cluster Resources

Location	Change
Step 7 in Section 10.6, "Enabling Monitoring and Configuring the Monitor Script," on page 155	Clarified the definition of *Maximum Local Failures*.

E.47.2 Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Step 8b in Section 12.8, "Adding Advertising Protocols for NSS Pool Cluster Resources," on page 191	You can download the cifsPool.py script file as a patch Novell Downloads Web site (http://download.novell.com/Download?buildid=sA5Tdb_U2CE~).

E.47.3 Installing Novell Cluster Services on OES 2 Linux

Location	Change
"Novell iManager 2.7.4" on page 63	**NOTE:** A February 3, 2009 update to the Clusters plug-in for OES 2 SP1 Linux is available that can be used to manage the Novell Business Continuity Clustering (BCC) 1.1 SP2 for NetWare 6.5 SP8 and Novell Business Continuity Clustering 1.2 for OES 2 SP1 Linux. You can download the update from the Novell Downloads Web site (http://download.novell.com/).

E.48   January 13, 2009

Updates were made to correct errata.

E.49   December 2008 (OES 2 SP1)

Updates were made to the following sections. The changes are explained below.

E.49.1   Comparison of Clustering OES 2 Services for Linux and NetWare

This section is new.

E.49.2   Comparison of Novell Cluster Services for Linux and NetWare

This section is new.

E.49.3   Configuring and Managing Cluster Resources

Location	Change
Section 10.2, "Using Cluster Resource Templates," on page 149	This section was modified for clarity.
Section 10.14, "Deleting Cluster Resources," on page 167	Added procedures for deleting cluster resources on the master node and the non-master node.

E.49.4   Configuring Cluster Resources for Shared Linux POSIX Volumes

Location	Change
Section 13.10, "Known Issues for Working with Cluster Resources for Linux POSIX Volumes," on page 235	This section is new.

E.49.5   Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 12.3, "Creating Cluster-Enabled Pools and Volumes," on page 178	Added information about using Novell CIFS for Linux and Novell AFP for Linux when cluster-enabling pools on OES 2 SP1 Linux and later.
Section 12.2.2, "Guidelines for Shared Volumes," on page 178	Added information about using encrypted NSS volumes on a shared pool.

Location	Change
Section 12.4, "Cluster-Enabling an Existing NSS Pool and Its Volumes," on page 184	Added information about using Novell CIFS for Linux and Novell AFP for Linux when cluster-enabling pools on OES 2 SP1 Linux and later.
Section 12.8, "Adding Advertising Protocols for NSS Pool Cluster Resources," on page 191	This section is new.
"Mirroring the Partition of an Existing NSS Pool" on page 198	This section is new.
Section 12.14, "Changing an Assigned Volume ID," on page 201	This section is new.

E.49.6 Configuring Novell Cluster Services in a Virtualization Environment

Location	Change
Chapter 14, "Configuring Novell Cluster Services in a Xen Virtualization Environment," on page 237	**IMPORTANT:** All templates except Xen and XenLive are valid in guest servers (DomU) in the virtualization environment. Only the Xen and XenLive templates can be used in the OES 2 Linux (Xen) host environment (that is, in Dom0, but not in DomU).
Section 14.2, "Virtual Machines as Cluster Resources," on page 238	An overview of OCFS2 is available in "Oracle Cluster File System 2" (http://www.suse.com/documentation/sles10/book_sle_reference/data/cha_ocfs2.html) in the *SUSE Linux Enterprise Server 10 SP4 Administration Guide* (http://www.suse.com/documentation/sles10/book_sle_reference/data/book_sle_reference.html). For detailed information about using OCFS2, see the *OCFS2 Project* (http://oss.oracle.com/projects/ocfs2/) on the Oracle Web site.

E.49.7 Console Commands for Novell Cluster Services

Location	Change
Section A.3, "extend_schema Command," on page 268	This section is new.

E.49.8 Installing Novell Cluster Services on OES 2 Linux

Location	Change
"Novell eDirectory 8.8.6" on page 59	If the eDirectory administrator user name or password contains special characters (such as $, #, and so on), ensure that you escape each special character by preceding it with a backslash (\) when you enter credentials.

Location	Change
"NCP Server for Linux" on page 67	This section is new.
"Novell AFP" on page 69	This section is new.
"Novell CIFS" on page 69	This section is new.
"SLP" on page 61	This section is new.
"Novell Domain Services for Windows" on page 69	This section is new.
Section 4.2, "IP Address Requirements," on page 58	Added information about the credentials needed by container administrators for installing Novell Cluster Services in trees where the schema has already been extended.
Section 5.2, "Extending the eDirectory Schema to Add Cluster Objects," on page 78	This section is new. It explains a new feature for extending the schema in a Novell eDirectory tree before clusters are installed.
Section 5.4.1, "Installing Novell Cluster Services during a OES 2 Linux Installation," on page 81	This section has been revised for OES 2 SP1.

E.49.9 Managing Clusters

Location	Change
Section 8.9, "Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster," on page 116	The procedure can also be used to modify the administrator user, or to modify the password of the existing administrator user. The procedure has been updated for clarity.

E.49.10 Overview of Novell Cluster Services

Location	Change
Section 1.1, "Why Should I Use Clusters?," on page 17	This section is new.
Section 1.4, "Clustering for High-Availability," on page 18	This scenario was relocated from Section 1.2, "Benefits of Novell Cluster Services," on page 17.
Section 1.5, "Shared Disk Scenarios," on page 21	This section was reorganized for clarity.

E.49.11　Troubleshooting Novell Cluster Services

Location	Change
Section 15.11, "Cluster Resource Goes Comatose Immediately After Migration or Failover," on page 255	This section is new.
Section 15.15, "NSS Takes Up to 10 Minutes to Load When the Server Is Rebooted (Linux)," on page 257	This section is new.

E.49.12　Upgrading OES 2 Clusters

This section is new.

E.49.13　What's New or Changed for Novell Cluster Services

Location	Change
Section 2.14, "What's New (OES 2 SP1)," on page 44	This section is new.

E.50　June 4, 2008

Updates were made to the following sections. The changes are explained below.

- Section E.50.1, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 335
- Section E.50.2, "Configuring Cluster Resources for Shared Linux POSIX Volumes," on page 336
- Section E.50.3, "Installation and Setup," on page 336
- Section E.50.4, "Managing Novell Cluster Services," on page 337

E.50.1　Configuring Cluster Resources for Shared NSS Pools and Volumes

Location	Change
Section 12.1, "Requirements for Creating Pool Cluster Resources," on page 173	This section is new.
Section 12.2, "Guidelines for Using Pool Cluster Resources," on page 177	This section is new.
Section 12.3.1, "Initializing and Sharing a Device," on page 179	This section is new.

Location	Change
Section 12.4, "Cluster-Enabling an Existing NSS Pool and Its Volumes," on page 184	Ensure that you offline the cluster resource before you attempt to delete the clustered pool or its cluster resource.
Section 12.10, "Mirroring and Cluster-Enabling Shared NSS Pools and Volumes," on page 194	This section is new.
Section 12.16, "Deleting NSS Pool Cluster Resources," on page 212	This section is new.

E.50.2 Configuring Cluster Resources for Shared Linux POSIX Volumes

Location	Change
Section 13.1, "Requirements for Creating Linux POSIX Volume Cluster Resources," on page 213	This section is new.
Section 13.3, "Creating Linux POSIX Volumes on Shared Devices," on page 216	Procedures were corrected.
Section 13.7, "Sample Scripts for a Linux POSIX Volume Cluster Resource," on page 231	Script commands were corrected.
Section 13.9, "Deleting Shared Storage," on page 235	This section is new.

E.50.3 Installation and Setup

The following topics were moved to separate sections:

- Chapter 8, "Configuring Cluster Policies and Priorities," on page 105
- Chapter 10, "Configuring and Managing Cluster Resources," on page 147
- Chapter 12, "Configuring Cluster Resources for Shared NSS Pools and Volumes," on page 173
- Chapter 13, "Configuring and Managing Cluster Resources for Shared Linux POSIX Volumes," on page 213

The following additional changes were made:

Location	Change
Chapter 4, "Planning for Novell Cluster Services," on page 55	This section is new.

Location	Change
Section 5.1, "Novell Cluster Services Licensing," on page 77	This section was updated.
Section 7.2, "Requirements and Guidelines for Upgrading Clusters from OES 1 Linux and OES 2 Linux," on page 101	This section is new.
"Planning the Conversion of Cluster Resources" in the *OES 2 SP3: Novell Cluster Services NetWare to Linux Conversion Guide*	This section is new.

E.50.4 Managing Novell Cluster Services

The following topics were moved to separate sections:

- Chapter 10, "Configuring and Managing Cluster Resources," on page 147
- Chapter 15, "Troubleshooting Novell Cluster Services," on page 251
- Appendix A, "Console Commands for Novell Cluster Services," on page 263
- Appendix B, "Files for Novell Cluster Services," on page 273

The following additional changes were made:

Location	Change
Chapter 9, "Managing Clusters," on page 125	This section was edited for clarity.
Section 9.2, "Monitoring Cluster and Resource States," on page 127	This section was edited for clarity.
Section 9.5, "Onlining and Offlining (Loading and Unloading) Cluster Resources from a Cluster Node," on page 130	This section is new.
Section 9.6, "Removing (Leaving) a Node from the Cluster," on page 132	This section is new.
Section 9.7, "Joining a Node to the Cluster," on page 132	This section is new.
Section 9.13, "Deleting a Cluster Node from a Cluster," on page 135	The procedure was edited for clarity.

E.51   May 2, 2008

Updates were made to the following sections. The changes are explained below.

E.51.1   Installation and Setup

Location	Change
Section 5.4.1, "Installing Novell Cluster Services during a OES 2 Linux Installation," on page 81	A one-node cluster can be configured without an SBD (split-brain detector). However, for adding a second node in the existing cluster (without SBD), you need additional pre-configuration of creating the SBD partition on the shared disk by using sbdutil on the first node.
Section 5.4.2, "Installing Novell Cluster Services on an Existing OES 2 Linux Server," on page 83	Corrected the procedure to use the YaST Control Center. You cannot install Novell Cluster Services by going directly to the yast2 ncs option at the command line.
	A one-node cluster can be configured without an SBD (split-brain detector). However, for adding a second node in the existing cluster (without SBD), you need additional pre-configuration of creating the SBD partition on the shared disk by using sbdutil on the first node.
"Configuring an Unload Script for a Linux POSIX Volume Cluster Resource" on page 224	In OES 2, use deactivate_evms_container instead of deport_evms_container.

E.51.2   Managing Novell Cluster Services

Location	Change
Section 9.2, "Monitoring Cluster and Resource States," on page 127	When the resource is in Start Alert state, you must clear the alert before you can offline the resource. After the resource is offline, you can online the resource.
Section 9.13, "Deleting a Cluster Node from a Cluster," on page 135	Details have been added for clarification.
Section 8.9, "Moving a Cluster, or Changing IP Addresses, LDAP Server, or Administrator Credentials for a Cluster," on page 116	This section is new.