

Linux POSIX Volume Administration Guide

Open Enterprise Server 11 SP1

December 2013

Novell.

Legal Notices

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

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

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

Copyright © 2011–2013 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

Novell, Inc.
1800 South Novell Place
Provo, UT 84606
U.S.A.
www.novell.com

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

Novell Trademarks

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

Third-Party Materials

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

Contents

About This Guide	7
1 Overview of Linux POSIX File Systems	9
1.1 Btrfs	9
1.2 Ext2	11
1.3 Ext3	11
1.4 ReiserFS	13
1.5 XFS	14
1.6 What's Next	15
2 What's New or Changed for Linux POSIX Volumes	17
2.1 What's New (OES 11 April 2013 Patches)	17
2.2 What's New (OES 11 SP1 January 2013 Patches)	17
2.3 What's New (OES 11 January 2013 Patches)	18
2.4 What's New (OES 11 SP1 November 2012 Patches)	19
2.5 What's New (OES 11 SP1 September 2012 Patches)	19
2.6 What's New (OES 11 September 2012 Patches)	20
2.7 What's New (OES 11 SP1)	20
2.8 What's New (OES 11)	22
3 Installing or Upgrading NSS Management Tools and Linux POSIX File Systems	25
3.1 NSS Management Tools	25
3.2 Linux POSIX File Systems	25
3.3 Clustered LVM	26
4 Using NLVM and NSSMU in a Virtualized Environment	27
5 Planning for Linux POSIX Volumes	29
5.1 Root User	29
5.2 Using RAID Devices for Linux POSIX Volumes	29
5.3 Naming Conventions for Linux LVM Volume Group and Logical Volume Names	30
5.3.1 NSSMU and NLVM Requirements for LVM Names	30
5.3.2 LVM Requirements for LVM Names	30
5.3.3 Novell Cluster Services Requirements for Clustered LVM Volume Names	30
5.4 Naming Conventions for NCP Volume Names	31
5.5 Linux LVM Volume Groups	31
5.6 Linux LVM Volume Group Cluster Resources	31
5.7 Using Antivirus Software with NCP Volumes	31
6 Managing Linux Volumes with NSSMU	33
6.1 NSSMU Quick Reference for Linux Volumes	34
6.2 Launching and Navigating NSSMU	35
6.3 Viewing a List of Devices	37

6.4	Viewing Device Details	37
6.5	Initializing a Device	39
6.6	Unsharing a Device	41
6.7	Viewing a List of Linux Volumes	41
6.8	Viewing Details for a Linux Volume	42
6.9	Creating a Non-LVM Linux Volume	44
6.10	Creating an LVM Logical Volume	48
6.11	Creating a Shared LVM Volume with Novell Cluster Services	52
6.12	Mounting a Linux Volume	53
6.13	Dismounting a Linux Volume	53
6.14	Renaming the Mount Point Path for a Linux Volume	53
6.15	Renaming a Non-LVM Linux Volume	55
6.16	Renaming an LVM Volume	55
6.17	Deleting a Linux Volume	56
7	Managing Linux Volumes with NLVM Commands	59
7.1	NLVM Commands Quick Reference for Linux Volumes	60
7.2	Viewing a List of Devices	62
7.3	Viewing Device Details	62
7.4	Initializing a Device	63
7.5	Unsharing a Device	65
7.6	Viewing a List of Linux Volumes	65
7.7	Creating a Linux POSIX Volume	66
7.8	Creating an LVM Logical Volume	67
7.9	Creating a Shared LVM Logical Volume with Novell Cluster Services	69
7.10	Mounting Linux Volumes	69
7.11	Dismounting Linux Volumes	71
7.12	Rescanning for Storage Objects on Known Devices	72
7.13	Deleting a Linux POSIX Volume	72
7.13.1	Deleting a Linux POSIX Volume	72
7.13.2	Deleting an LVM Volume Group and Logical Volume	73
7.14	Additional Information	74
8	Clustering LVM Volume Groups with Novell Cluster Services	75
8.1	Requirements for Creating LVM Cluster Resources	75
8.1.1	Novell Cluster Services	76
8.1.2	Linux Logical Volume Manager 2 (LVM2)	76
8.1.3	Clustered Logical Volume Manager Daemon (CLVMD)	76
8.1.4	Resource IP Address	76
8.1.5	Shared Storage Devices	76
8.1.6	All Nodes Must Be Present	77
8.1.7	Working in Mixed Node OES Clusters	77
8.1.8	NCP File Access with Novell NCP Server	77
8.1.9	SMB/CIFS File Access with Novell Samba	78
8.1.10	Linux File Access Protocols	79
8.2	Initializing a SAN Device	79
8.3	Configuring an LVM Volume Group Cluster Resource with NSS Management Tools	81
8.3.1	Sample Values	82
8.3.2	Creating an LVM Volume Group Cluster Resource with NSSMU	83
8.3.3	Creating an LVM Volume Group Cluster Resource with NLVM Commands	88
8.4	Configuring the LVM Cluster Resource Settings	91
8.5	Viewing or Modifying the LVM Resource Scripts	94
8.6	Sample LVM Resource Scripts	96

8.6.1	Sample LVM Resource Load Scripts Created by NSS Management Tools	97
8.6.2	Sample LVM Resource Unload Scripts Created by NSS Management Tools.	98
8.6.3	Sample LVM Resource Monitor Scripts Created by NSS Management Tools	99
8.7	Renaming the Mount Point Path for a Clustered LVM Volume	101
8.8	Renaming a Clustered LVM Logical Volume	102
8.9	Disabling Clustering for an LVM Volume	104
8.10	Deleting a Clustered LVM Volume Group and Logical Volume	106
8.11	Deleting a Clustered LVM Volume (Created in NSSMU or NLVM)	110
8.11.1	Deleting a Cluster-Enabled LVM Volume on the Master Node	110
8.11.2	Deleting a Cluster-Enabled LVM Volume on a Non-Master Node.	111

9 Backing Up Linux POSIX Volumes 113

A Documentation Updates 115

A.1	December 3, 2013	115
A.2	September 19, 2013.	115
A.3	July 31, 2013	115
A.3.1	Clustering LVM Volume Groups with Novell Cluster Services.	115
A.4	April 2013 Scheduled Maintenance	116
A.4.1	What's New for Linux POSIX Volumes	116
A.5	January 2013 Scheduled Maintenance	116
A.5.1	What's New for Linux POSIX Volumes	116
A.6	January 11, 2013	116
A.6.1	Planning for Linux POSIX Volumes.	117
A.7	November 2012 Scheduled Maintenance	117
A.7.1	Clustering LVM Volume Groups with Novell Cluster Services.	117
A.7.2	Managing Linux Volumes with NLVM Commands.	118
A.7.3	Managing Linux Volumes with NSSMU.	118
A.7.4	Overview of Linux POSIX File Systems.	118
A.7.5	What's New for Linux POSIX Volumes	118
A.8	September 2012 Scheduled Maintenance	119
A.8.1	What's New for Linux POSIX Volumes	119
A.9	August 28, 2012 (OES 11 SP1)	119
A.9.1	Backing Up Linux POSIX Volumes	119
A.9.2	Clustering LVM Volume Groups with Novell Cluster Services.	119
A.9.3	Installing or Upgrading NSS Management Tools and Linux POSIX File Systems.	120
A.9.4	Managing Linux POSIX Volumes with NLVM Commands.	120
A.9.5	Managing Linux POSIX Volumes with NSSMU.	120
A.9.6	Planning for Linux POSIX Volumes.	120
A.9.7	Using NLVM and NSSMU in a Virtualized Environment	120
A.9.8	What's New for Linux POSIX Volumes	121
A.10	January 18, 2012	121

About This Guide

This guide describes how to manage Linux POSIX volumes on a Novell Open Enterprise Server (OES) 11 Support Pack 1 (SP1) server by using the Novell Storage Services (NSS) Management Utility and the Novell Linux Volume Manager (NLVM) command line interface (CLI).

- ♦ Chapter 1, “Overview of Linux POSIX File Systems,” on page 9
- ♦ Chapter 2, “What’s New or Changed for Linux POSIX Volumes,” on page 17
- ♦ Chapter 3, “Installing or Upgrading NSS Management Tools and Linux POSIX File Systems,” on page 25
- ♦ Chapter 4, “Using NLVM and NSSMU in a Virtualized Environment,” on page 27
- ♦ Chapter 5, “Planning for Linux POSIX Volumes,” on page 29
- ♦ Chapter 6, “Managing Linux Volumes with NSSMU,” on page 33
- ♦ Chapter 7, “Managing Linux Volumes with NLVM Commands,” on page 59
- ♦ Chapter 8, “Clustering LVM Volume Groups with Novell Cluster Services,” on page 75
- ♦ Chapter 9, “Backing Up Linux POSIX Volumes,” on page 113
- ♦ Appendix A, “Documentation Updates,” on page 115

Audience

This guide is intended for storage and system administrators.

Feedback

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

Documentation Updates

For the most recent version of the *OES 11 SP1: Linux POSIX Volume Administration Guide*, visit the [OES 11 SP1 Web site \(http://www.novell.com/documentation/oes11/\)](http://www.novell.com/documentation/oes11/).

Additional Documentation

For information about the NLVM commands, see the [OES 11 SP1: NLVM Reference](#).

For information about managing the NSS file system, see the [OES 11 SP1: NSS File System Administration Guide for Linux](#).

For information about clustering with Novell Cluster Services, see the [OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide](#).

For information about using the file system management tools that are native to SUSE Linux Enterprise Server (SLES) 11 Service Pack 2 (SP2), see the *SLES 11 SP2 Storage Administration Guide*. (http://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html). The SLES guide describes Linux storage management technologies such as Linux Volume Manager (LVM), multipath I/O, Linux iSCSI initiators and targets, and the Internet Storage Name Service (iSNS).

1 Overview of Linux POSIX File Systems

SUSE Linux Enterprise Server (SLES) 11 Service Pack 2 (SP2) provides several native Linux POSIX file systems. In Novell Open Enterprise Server (OES) 11 SP1, the Novell Storage Services (NSS) Management Utility (NSSMU) and the Novell Linux Volume Manager (NLVM) commands are designed to be used with the Btrfs, Ext2, Ext3, ReiserFS, and XFS file systems.

This section describes key features of the supported Linux POSIX file systems, including access control, availability, scalability, and speed. You can use this information to help determine which of these file systems best fits your storage needs.

See the following sections for an overview of the Linux POSIX file systems that are supported by NSSMU and NLVM:

- ♦ [Section 1.1, “Btrfs,” on page 9](#)
- ♦ [Section 1.2, “Ext2,” on page 11](#)
- ♦ [Section 1.3, “Ext3,” on page 11](#)
- ♦ [Section 1.4, “ReiserFS,” on page 13](#)
- ♦ [Section 1.5, “XFS,” on page 14](#)
- ♦ [Section 1.6, “What’s Next,” on page 15](#)

The NSSMU and NLVM tools also support the NSS file system. In addition, the Storage plug-in for Novell iManager is available to manage NSS storage objects, including devices, pools, volumes, partitions, and software RAID devices. For information, see the [OES 11 SP1: NSS File System Administration Guide for Linux](#).

The NSSMU and NLVM tools also support the NSS file system on clustered NSS pools and Linux POSIX file systems on clustered Linux Logical Volume Manager (LVM) volume groups on OES 11 SP1 nodes in your Novell Cluster Services clusters. In addition, the Clusters plug-in for Novell iManager is available to manage the cluster and cluster resources for the pools and volume groups. For information, see the [OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide](#).

For a side-by-side comparison of the technical features of the file systems, see “File System Support” on the [SUSE Linux Enterprise Server Technical Information Web page \(http://www.suse.com/products/server/technical-information/#FileSystem\)](http://www.suse.com/products/server/technical-information/#FileSystem).

1.1 Btrfs

The Btrfs file system is a copy-on-write (COW) file system based on COW-friendly B-trees. It provides fault tolerance, repair, and easy management features, such as the following:

- ♦ Writable snapshots that allow you to easily roll back your system if needed after applying updates, or to back up files.
- ♦ Multiple device support that allows you to grow or shrink the file system.
- ♦ Compression to efficiently use storage space.

- Different RAID levels for metadata and user data.
- Different checksums for metadata and user data to improve error detection.
- Integration with Linux Logical Volume Manager (LVM) storage objects.
- Integration with the YaST 2 Partitioner and AutoYaST on SUSE Linux.
- Offline, in-place migration from existing Ext2, Ext3, and Ext4 file systems.

For more information, see “Btrfs” (http://www.suse.com/documentation/sles11/stor_admin/data/sec_filesystems_major.html#bwk8gda) in the *SUSE Linux Enterprise Server 11 SP2 Storage Administration Guide* (http://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html)

Access Control

Btrfs uses the POSIX access control model.

Availability

Btrfs is a logging-style file system. Instead of journaling the block changes, it writes them in a new location, then links the change in. Until the last write, the new changes are not committed.

Btrfs provides writable snapshots with the SUSE Snapper infrastructure that allow you to easily roll back your system if needed after applying updates, or to back up files. Snapper allows you to create and delete snapshots, and to compare snapshots and revert the differences between them. For information about Snapper and its integration in ZYpp (`snapper-zypp-plugin`) and YaST 2 (`yast2-snapper`), see the “Snapshots/Rollback with Snapper” (http://www.suse.com/documentation/sles11/book_sle_admin/data/cha_snapper.html) in the *SUSE Linux Enterprise Server 11 SP2 Administration Guide* (http://www.suse.com/documentation/sles11/book_sle_admin/data/book_sle_admin.html).

IMPORTANT: Because Btrfs is capable of storing snapshots of the file system, it is advisable to reserve twice the amount of disk space than the standard storage proposal. This is done automatically by the YaST2 Partitioner in the Btrfs storage proposal for the root file system.

Scalability

Btrfs is designed to handle very large storage volumes, up to 16 EiB (exbibytes). Its maximum file size is 16 EiB. It provides both online and offline expand and shrink capability.

Speed

Btrfs is designed to read and write fast. It provides performance features such as fine-grained B-tree locking, running checksum operations to their own background threads, and using a cache to reduce IO requirements when cleaning up old transactions.

Btrfs Administration

The Novell Storage Services (NSS) Management Utility (NSSMU) and Novell Linux Volume Manager (NLVM) support the creation, listing, and deletion of Btrfs file systems. The `btrfsprogs` package is required to be installed.

Btrfs is integrated in the YaST2 Partitioner and AutoYaST. It is available during the installation to allow you to set up a solution for the root file system. You can use the YaST2 Partitioner after the install to view and manage Btrfs volumes.

Btrfs administration tools are provided in the `btrfsprogs` package. For information about using Btrfs commands, see the `btrfs(8)`, `btrfsck(8)`, `mkfs.btrfs(8)`, and `btrfsctl(8)` man pages. For information about Btrfs features, see the [Btrfs wiki \(http://btrfs.wiki.kernel.org\)](http://btrfs.wiki.kernel.org).

1.2 Ext2

The Ext2 file system is the predecessor to Ext3. Ext2 is not journaled. As one of the earliest file systems on Linux, Ext2 has been heavily tested and improved over the years. This might be the reason why people often refer to it as “rock-solid”.

Access Control

Ext2 uses the POSIX access control model.

Availability

After a system outage when the file system cannot be cleanly unmounted, `e2fsck` starts to analyze the file system data. Metadata is brought into a consistent state and pending files or data blocks are written to a designated directory (called `lost+found`). In contrast to journaling file systems, `e2fsck` analyzes the entire file system and not just the recently modified bits of metadata. This takes significantly longer than checking the log data of a journaling file system. Depending on file system size, this procedure can take half an hour or more. Therefore, it is not desirable to choose Ext2 for any server that needs high availability.

Scalability

Ext2 does not scale well to large volumes or to a great number of files.

Speed

Because Ext2 does not maintain a journal and uses significantly less memory, it is sometimes faster than other file systems.

Easy Upgradability to Ext3

Because Ext3 is based on the Ext2 code and shares its on-disk format as well as its metadata format, upgrades from Ext2 to Ext3 are very easy.

1.3 Ext3

The Ext3 file system is a journaled file system that has the greatest use in Linux today. It is the default file system in SUSE Linux Enterprise Server 11 distributions. It is quite robust and quick. Ext3 is based on the Ext2 code and shares its on-disk format as well as its metadata format. The reliability and solidity of Ext2 is elegantly combined in Ext3 with the advantages of a journaling file system.

Access Control

Ext3 uses the POSIX extended access control model.

Availability

Ext3 journaling is designed to take care of both metadata and data integrity and keep them in a consistent state. The amount of time to recover an Ext3 file system after an unclean system shutdown does not depend on the size of the file system or the number of files; rather, it depends on the size of the journal used to maintain consistency. The default journal size takes about a second to recover, depending on the speed of the hardware.

Integrity and performance can be customized by specifying the journaling mode as `journal`, `ordered`, or `writeback`. Enabling Ext3 in the `data=journal` mode offers maximum security (data integrity), but can slow down the system because both metadata and data are journaled. Enabling Ext3 in the `data=ordered` mode ensures both data and metadata integrity, but it uses journaling only for metadata. The file system driver collects all data blocks that correspond to one metadata update. These data blocks are written to disk before the metadata is updated. As a result, consistency is achieved for metadata and data without sacrificing performance. A third option to use is `data=writeback`, which allows data to be written to the main file system after its metadata has been committed to the journal. This option is often considered the best in performance. It can, however, allow old data to reappear in files after a crash and recovery even though the internal file system integrity is maintained.

IMPORTANT: Ext3 uses the `data=ordered` option as the default.

Scalability

Ext3 does not scale well to large volumes or to a great number of files. The recently added Ext3 h-trees feature significantly improves its scalability over Ext2. However, it is still not as scalable as some of the other file systems in Linux. With h-trees, Ext3 scales similarly to NTFS. Without h-trees, Ext3 handles about 5,000 files in a directory.

Speed

Ext3's journaling optimizes hard drive head motion, which typically provides a higher throughput than Ext2. The journaling mode you use determines the performance achieved.

Easy and Highly Reliable Upgrades from Ext2

The `tune2fs` utility makes it possible to easily upgrade an Ext2 file system to Ext3 without reformatting.

Downgrading from Ext3 to Ext2 is also easy. Just perform a clean dismount of the Ext3 file system and remount it as an Ext2 file system.

Default Ext3 Inode Size

To allow space for extended attributes and ACLs for a file on Ext3 file systems, the default inode size for Ext3 was increased from 128 bytes on SLES 10 to 256 bytes on SLES 11. The maximum number of files that can be stored for an Ext3 file system on SLES 11 is reduced by half as compared to Ext3 file systems created on SLES 10 because each inode consumes twice the space. As a result, you need to use larger partitions to hold the same number of files on SLES 11. The inode setting cannot be changed for an existing Ext3 file system unless the file system gets extended.

If you do not use extended attributes or ACLs on your ext3 file systems, you can set the inode size to 128 bytes by using any of the following methods:

- ♦ **Modifying the default for all new Ext3 files:** Modify the `/etc/mke2fs.conf` file to set the default `inode_ratio` to 8192. The value applies to all new Ext3 file systems.
- ♦ **At the command line:** Pass the bytes-per-inode option (`-i 8192`) to the `mkfs.ext3` command when you create a new Ext3 file system.
- ♦ **During installation with YaST:** Pass the *Bytes per inode* option when you create a new Ext3 file system during the installation. In the YaST Partitioner on the *Edit Partition* page under *Formatting Options*, select *Format partition > Ext3*, then click *Options*. In the File system options dialog box, select 8192 from the *Bytes per inode* drop-down list, then click *OK*.
- ♦ **During installation with autoyast:** In an autoyast profile, you can use the `fs_options` tag:

```
<fs_options>
  <opt_bytes_per_inode>
    <option_str>-I</option_str>
    <option_value>8192</option_value>
  </opt_bytes_per_inode>
</fs_options>
```

For information, see [SLES11 ext3 partitions can only store 50% of the files that can be stored on SLES10 \[Technical Information Document 7009075\]](http://www.novell.com/support/kb/doc.php?id=7009075) (<http://www.novell.com/support/kb/doc.php?id=7009075>).

1.4 ReiserFS

The Reiser File System (ReiserFS) is a journaling file system that is designed to improve the scalability and performance over the Ext2 and Ext3 file systems. By comparison, ReiserFS provides better disk space utilization, better disk access performance, faster crash recovery, and reliability through data journaling.

Access Control

ReiserFS uses the POSIX extended access control model.

Availability

ReiserFS uses ordered journaling by default. The file system driver collects all data blocks that correspond to one metadata update. These data blocks are written to disk before the metadata is updated. As a result, consistency is achieved for metadata and data without sacrificing performance. The commit policy depends on the journal size but is based on the number of blocks to commit. Using a journal to keep track of recent metadata changes makes a file system check a matter of seconds, even for huge file systems.

Scalability

ReiserFS provides the best performance and scalability when there are many files and the files are small. It scales and performs extremely well on Linux, out-scaling Ext3 with h-trees. In addition, ReiserFS is designed to very efficiently use disk space. As a result, it is a good choice on Linux where there are many small files in the file system. Because collaboration (email) and many Web serving applications have many small files, ReiserFS is best suited for these types of workloads.

Speed

For small files, file data and inode information are often stored next to each other. They can be read with a single disk I/O operation, meaning that only one access to disk is required to retrieve all the information needed.

1.5 XFS

The XFS file system is a highly scalable, high-performance 64-bit journaling file system that is designed to meet extreme computing challenges. XFS is very good at manipulating large files and performs well on high-end hardware.

Access Control

XFS uses the POSIX extended access control model.

Availability

XFS supports metadata journaling, which facilitates quicker crash recovery. The journal can be stored within the data section (an internal log) of the file system, or on a separate device to minimize disk contention. On XFS, the journal tracks high-level operations that are being performed. Journal updates are performed asynchronously to avoid incurring a performance penalty. If a server crash occurs, XFS retains file system consistency by redoing operations that were incomplete prior to the crash. Recovery occurs at mount time. The recovery speed is independent of the size of the file system. If recently modified data was not flushed to disk before a system crash, XFS zeros any unwritten data blocks on reboot in order to avoid any possible security issues that might be caused by residual data.

Scalability

XFS originated at Silicon Graphics, Inc. (SGI) for Irix and was designed specifically for large files and large volume scalability. It scales to petabyte volumes to handle extremely large file systems (up to 16 exabytes), files (8 exabytes), and directory structures (tens of millions of entries). Video and multimedia files are best handled by this file system.

XFS uses allocation groups that can be independently addressed by the file system. This allows concurrent file system access to groups by multiprocessor systems.

Speed

Before writing the data to the file system, XFS reserves (preallocates) the free space needed for a file. Thus, file system fragmentation is greatly reduced. Performance is increased because the contents of a file are not distributed throughout the file system.

Free space and inodes are handled by B+ trees inside the allocation groups. The use of B+ trees greatly contributes to XFS's performance and scalability.

XFS uses delayed allocation to reduce the number of writes it needs to make. It holds a pending transaction in RAM and reserves the appropriate amount of space, but does not write it immediately. Some short-lived temporary data might become obsolete by the time XFS decides where actually to save it, so it does not need to be written. In this way, XFS increases write performance and reduces file system fragmentation. Because delayed allocation results in less frequent write events than in other file systems, it is likely that data loss after a crash during a write is more severe.

Proprietary Versions

SGI offers a closed source cluster parallel version of XFS called cXFS, which uses an asymmetrical model. The unique feature of cXFS is that its slave nodes can run on UNIX, Linux, and Windows, making it a cross-platform file system. Its master node must run on SGI hardware.

1.6 What's Next

To use NSSMU to create and manage Linux POSIX volumes and LVM volumes, see [Chapter 6, "Managing Linux Volumes with NSSMU,"](#) on page 33.

To use NLVM commands to create and manage Linux POSIX volumes and LVM logical volumes, see [Chapter 7, "Managing Linux Volumes with NLVM Commands,"](#) on page 59.

To use NSSMU or NLVM to cluster-enable LVM volume groups and logical volumes, see [Chapter 8, "Clustering LVM Volume Groups with Novell Cluster Services,"](#) on page 75.

2 What's New or Changed for Linux POSIX Volumes

This section describes the changes made to the management of Linux POSIX volumes with the Novell Linux Volume Manager (NLVM) and the Novell Storage Services Management Utility (NSSMU) since the Novell Open Enterprise Server (OES) 11 release.

- ♦ [Section 2.1, “What’s New \(OES 11 April 2013 Patches\),” on page 17](#)
- ♦ [Section 2.2, “What’s New \(OES 11 SP1 January 2013 Patches\),” on page 17](#)
- ♦ [Section 2.3, “What’s New \(OES 11 January 2013 Patches\),” on page 18](#)
- ♦ [Section 2.4, “What’s New \(OES 11 SP1 November 2012 Patches\),” on page 19](#)
- ♦ [Section 2.5, “What’s New \(OES 11 SP1 September 2012 Patches\),” on page 19](#)
- ♦ [Section 2.6, “What’s New \(OES 11 September 2012 Patches\),” on page 20](#)
- ♦ [Section 2.7, “What’s New \(OES 11 SP1\),” on page 20](#)
- ♦ [Section 2.8, “What’s New \(OES 11\),” on page 22](#)

2.1 What’s New (OES 11 April 2013 Patches)

Upgrade to eDirectory 8.8.7

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

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

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

2.2 What’s New (OES 11 SP1 January 2013 Patches)

Upgrade to Novell iManager 2.7.6

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

Novell iManager 2.7.6 provides the following enhancements:

- ♦ Microsoft Internet Explorer 10 certification in the desktop user interface view on Windows 8 (excluding Windows 8 RT) and Windows Server 2012.

- ♦ Apple Safari 6.0 certification on Mac OS X Mountain Lion (version 10.8).
- ♦ iManager Workstation certification on Windows 8 Enterprise Edition (32-bit and 64-bit).
- ♦ Manager 2.7.6 support for Tomcat 7.0.32. and Java 1.7.0_04 versions.

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

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

Novell Client Support for Windows 8 and Server 2012

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

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

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

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

NCP Server

NCP Server has been modified to automatically log dismounts of NSS volumes and NCP volumes.

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

Upgrade to Novell iManager 2.7.6

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

Novell iManager 2.7.6 provides the following enhancements:

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

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

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

Novell Client Support for Windows 8 and Server 2012

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

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

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

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

New Novell Cluster Services Plug-in for iManager 2.7.5 and Later

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

NCP Server

NCP Server has been modified to automatically log dismounts of NSS volumes and NCP volumes.

2.4 What’s New (OES 11 SP1 November 2012 Patches)

In addition to bug fixes, the following changes and enhancements were made in the November 2012 Scheduled Maintenance for OES 11 SP1:

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

2.5 What’s New (OES 11 SP1 September 2012 Patches)

In addition to bug fixes, the following changes and enhancements were made in the September 2012 Scheduled Maintenance for OES 11 SP1:

- ♦ **Creating Linux POSIX Volumes:** Linux POSIX file systems should not be created on NSS software RAID. The NSSMU and NLVM management tools have been modified to block NSS software RAID from being presented as device options when you create a Linux POSIX volume.

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

In addition to bug fixes, the following changes and enhancements were made in the September 2012 Scheduled Maintenance for OES 11:

- ♦ **Creating Linux POSIX Volumes:** Linux POSIX file systems should not be created on NSS software RAIDs. The NSSMU and NLVM management tools have been modified to block NSS software RAIDs from being presented as device options when you create a Linux POSIX volume.

2.7 What's New (OES 11 SP1)

The Linux POSIX file systems support OES 11 SP1 running on 64-bit SUSE Linux Enterprise Server (SLES) 11 SP2. In addition to bug fixes, the following changes and enhancements are available:

NSSMU

The NSS Management Utility (NSSMU) supports OES 11 SP1 file systems running on 64-bit SUSE Linux Enterprise Server (SLES) 11 SP2. In addition to bug fixes, the following changes and enhancements for Linux POSIX volume management were made to NSSMU in the OES 11 SP1 release:

- ♦ **Create Linux Volume:**
 - ♦ The *NCP Enable* option can be used to enable the Linux POSIX file system on the volume to be accessed with the NetWare Core Protocol (NCP).
 - ♦ Support was added for the Btrfs file system. The Btrfs file system is now available as an alternative to other supported file systems (Ext2, Ext3, Reiserfs, and XFS) when you create a Linux volume. The *btrfsprogs* package is required. You can use NSSMU or NLVM to manage the Btrfs volumes as you would manage other Linux data volumes. Btrfs volumes can also be configured for NCP file access and for clustering with Novell Cluster Services.
- ♦ **Confirmation Prompt:** You are automatically prompted to confirm actions for certain commands that destroy data, such as when you initialize a device or delete Linux POSIX volumes.
- ♦ **File System Types:** When you create a Linux POSIX volumes, the supported file system types are now listed in alphabetical order: btrfs, ext2, ext3, reiserfs, and xfs.
- ♦ **Linux Volume Details:** On the Linux Volumes page, a volume's information includes the *NCP Enabled* parameter and status (*Yes* or *No*).
- ♦ **Mount Point Path:** The default mount point path for Linux volumes is `/usr/novell/<volume_name>`.
- ♦ **Partitions:** On the NSSMU Partitions page, the Partitions list panel has been widened to accommodate long partition names. If you widen the NSSMU display screen, it widens the Partitions list panel and allows you to see the partitions' full names.

NLVM Commands

The Novell Linux Volume Manager (NLVM) commands support OES 11 SP1 file systems running on 64-bit SUSE Linux Enterprise Server (SLES) 11 SP2. In addition to bug fixes, the following changes and enhancements for Linux POSIX volume management were made in the OES 11 SP1 release:

- ♦ **Sizes:** The following enhancements are available for the `size` option:
 - ♦ The default multiplier for sizes was changed from M (megabytes) to G (gigabytes). If no multiplier is provided, the value is assumed to be in gigabytes. For example:

```
size=20    (20 gigabytes is used)
```
 - ♦ Use the T multiplier to specify sizes in terabytes (TB).
 - ♦ All sizes can be entered as whole numbers or with fractional parts such as 200.45G and 3.98T.
- ♦ **Multiple Devices:** You can specify multiple devices for the following commands:
 - ♦ **Init:** You can specify multiple devices to be initialized at the same time. Each device is formatted with the same specified partitioning scheme and share state.
 - ♦ **Share:** You can specify multiple devices to be shared at the same time.
 - ♦ **Unshare:** You can specify multiple devices to be unshared at the same time.
- ♦ **No Prompt:** The `--no-prompt` NLVM option can be used with the `nlvm delete linux volume` command and the `nlvm init` command to prevent a confirmation message from being displayed.
- ♦ **Confirmation Prompt:** You are automatically prompted to confirm actions to initialize a device or to delete a Linux POSIX volume.
- ♦ **Create Linux Volume:**
 - ♦ Support was added for the Btrfs file system. The Btrfs file system is now available as an alternative to other supported file systems (Ext2, Ext3, ReiserFs, and XFS) when you create a Linux volume. The `btrfsprogs` package is required. You can use NSSMU or NLVM to manage the Btrfs volumes as you would manage other Linux data volumes. Btrfs volumes can also be configured for NCP file access and for clustering with Novell Cluster Services
 - ♦ The `ncp` option for the `nlvm create linux volume` command can be used to enable the Linux POSIX file system on the volume to be accessed with the NetWare Core Protocol (NCP).
- ♦ **List Linux Volumes:** The volume information displays NCP if the Linux volume is enabled for NCP.

Btrfs File System

NSSMU, NLVM Linux commands, NCP Server, and Novell Cluster Services were modified to support the Btrfs file system for Linux volumes. The Btrfs file system is now available as an alternative to other supported file systems (Ext2, Ext3, Reiserfs, and XFS) when you create a Linux volume. The `btrfsprogs` package is required. You can use NSSMU or NLVM to manage the Btrfs volumes as you would manage other Linux data volumes. Btrfs volumes can also be configured for NCP file access and for clustering with Novell Cluster Services.

SUSE Linux Enterprise Server (SLES) 11 Service Pack 2 (SP2) also supports the Btrfs file system in the YaST2 Partitioner. You can use the Partitioner during the system installation to configure a Btrfs solution for the system device. YaST2 automatically installs the `btrfsprogs` package if you use Btrfs for the root (/) file system. Bootloader support for `/boot` on Btrfs is not available in SLES 11 SP2; it is planned to be available beginning in SUSE Linux Enterprise 12.

For information about the features of Btrfs, see “Btrfs” (http://www.novell.com/documentation/oes11/stor_posixvol_lx/data/bwk8gda.html) in the *OES 11 SP1: Linux POSIX Volume Administration Guide* (http://www.novell.com/documentation/oes11/stor_posixvol_lx/data/bookinfo.html).

NCP Volumes and Cluster Resources

NSSMU and NLVM were modified to support the creation of NCP volumes on Linux POSIX file systems. You can NCP enable Linux POSIX volumes, an LVM volumes, and clustered LVM volumes as you create them. When you enable NCP for a Linux volume, an NCP volume of the same name (using all capital letters) is automatically created. An NCP share is created at the root of the Linux volume.

For clustered LVM volumes, the NCP Virtual Server object is automatically created and the NCP mount, dismount, and status commands are added to the scripts for the LVM volume group cluster resource. The Novell Cluster Services Generic File System template is used to specify the default cluster policies and scripts for the cluster resource. After you create the resource, you can use the Clusters plug-in for Novell iManager to set up preferred nodes for the resource and to manage the resource.

Linux Volume Support in SLES 11 SP2

For information about Linux volume support changes in SUSE Linux Enterprise Server 11 SP1, see “What’s New in SLES 11 SP2” (https://www.suse.com/documentation/sles11/stor_admin/data/bwkb7my.html) in the *SUSE Linux Enterprise Server 11 SP2 Storage Administration Guide* (https://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html).

2.8 What’s New (OES 11)

Creating Linux Volumes with NSSMU and NLVM Commands

For Novell Open Enterprise Server (OES) 11 servers, the Novell Storage Services Management Utility (NSSMU) and the Novell Linux Volume Manager (NLVM) commands allow you to create, mount, and delete the following Linux POSIX storage objects:

- ♦ Linux POSIX volumes (non-LVM volumes)
- ♦ Linux Logical Volume Manager 2 (LVM2) volume groups and LVM logical volumes
- ♦ Shared Linux Clustered LVM (cLVM) volume groups and logical volumes that are cluster-enabled with Novell Cluster Services

Linux Volume Support in SLES 11 SP1

For information about Linux volume support changes in SUSE Linux Enterprise Server 11 SP1, see “What’s New in SLES 11 SP1” (https://www.suse.com/documentation/sles11/stor_admin/data/bndgyod.html) in the *SUSE Linux Enterprise Server 11 SP2 Storage Administration Guide* (https://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html).

Linux Volume Support in SLES 11

For information about Linux volume support changes in SUSE Linux Enterprise Server 11, see “What’s New in SLES 11” (https://www.suse.com/documentation/sles11/stor_admin/data/bndgyoe.html) in the *SUSE Linux Enterprise Server 11 SP2 Storage Administration Guide* (https://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html).

3 Installing or Upgrading NSS Management Tools and Linux POSIX File Systems

This section describes the installation and upgrade requirements for using Novell Storage Services (NSS) management tools to manage Linux POSIX volumes on your Novell Open Enterprise Server (OES) 11 SP1 server.

- [Section 3.1, “NSS Management Tools,” on page 25](#)
- [Section 3.2, “Linux POSIX File Systems,” on page 25](#)
- [Section 3.3, “Clustered LVM,” on page 26](#)

3.1 NSS Management Tools

Novell Open Enterprise Server 11 SP1 is required for using the NSS Management Utility (NSSMU) and the Novell Linux Volume Manager (NLVM) to manage Linux POSIX volumes. For general information about installing, upgrading, and patching your OES 11 (or later) server, see the [OES 11 SP1: Installation Guide](#).

During the OES installation, select the *Novell Storage Services* pattern from the list of *OES Services*. This automatically installs NSSMU and NLVM.

Updates for the NSS management tools are delivered in the OES 11 (or later) patch channels.

3.2 Linux POSIX File Systems

The software packages for most of the Linux POSIX file systems (such as Ext2, Ext3, ReiserFS, and XFS) are installed and upgraded by default whenever you install or upgrade the SUSE Linux Enterprise Server (SLES) 11 SP 1 (or later) operating system on your OES 11 (or later) server. No action is required.

The Btrfs file system is available on SLES 11 SP2 or later. Using the Btrfs file system requires the `btrfsprogs` package. The package is installed automatically if you select Btrfs as the file system for the system device when you install the SLES 11 SP2 (or later) operating system. Otherwise, the `btrfsprogs` package must be installed manually by using the YaST 2 Software Installer. Thereafter, the package is updated by default whenever you patch or upgrade your system.

Updates for the Linux POSIX file systems and the `btrfsprogs` package are delivered via the SLES 11 SP2 (or later) patch channel.

3.3 Clustered LVM

Clustered LVM (CLVM) requires the Linux kernel 2.6.32.45-0.3 or later. You can get the latest kernel version by using the SLES 11 SP2 update channel.

4 Using NLVM and NSSMU in a Virtualized Environment

The Novell Linux Volume Manager (NLVM) utility and Novell Storage Services Management Utility (NSSMU) run in a virtualized environment just as they do on a physical server running Novell Open Enterprise Server 11 and later, and require no special configuration or other changes.

- ♦ For information on setting up virtualized OES 11 SP1, see “[Installing, Upgrading, or Updating OES on a VM](#)” in the *OES 11 SP1: Installation Guide*.
- ♦ To get started with Xen virtualization, see the [Virtualization with Xen](http://www.suse.com/documentation/sles11/book_xen/data/book_xen.html) documentation (http://www.suse.com/documentation/sles11/book_xen/data/book_xen.html).
- ♦ To get started with KVM virtualization, see the [Virtualization with KVM](http://www.suse.com/documentation/sles11/book_kvm/data/book_kvm.html) documentation (http://www.suse.com/documentation/sles11/book_kvm/data/book_kvm.html).

5 Planning for Linux POSIX Volumes

Consider the requirements and caveats in this section when planning to use the Novell Storage Services (NSS) Management Utility (NSSMU) or Novell Linux Volume Manager (NLVM) command line commands to manage Linux POSIX volumes on your Novell Open Enterprise Server (OES) 11 SP1 servers.

- ♦ [Section 5.1, “Root User,” on page 29](#)
- ♦ [Section 5.2, “Using RAID Devices for Linux POSIX Volumes,” on page 29](#)
- ♦ [Section 5.3, “Naming Conventions for Linux LVM Volume Group and Logical Volume Names,” on page 30](#)
- ♦ [Section 5.4, “Naming Conventions for NCP Volume Names,” on page 31](#)
- ♦ [Section 5.5, “Linux LVM Volume Groups,” on page 31](#)
- ♦ [Section 5.6, “Linux LVM Volume Group Cluster Resources,” on page 31](#)
- ♦ [Section 5.7, “Using Antivirus Software with NCP Volumes,” on page 31](#)

5.1 Root User

The Linux system root user privileges are required to use NSSMU and the NLVM commands.

5.2 Using RAID Devices for Linux POSIX Volumes

You can use a hardware RAID device to achieve device fault tolerance for Linux POSIX volumes. Novell Storage Services (NSS) management tools can see and use any available free space on a hardware RAID device.

We recommend that you do not use Linux software RAIDs (such as MD RAIDs and Device Mapper RAIDs) for devices that you plan to use for storage objects that are managed by NSS management tools. The Novell Linux Volume Manager (NLVM) utility and the NSS Management Utility (NSSMU) list Linux software RAID devices that you have created by using Linux tools. Beginning with Linux Kernel 3.0 in OES 11 SP1, NLVM and NSSMU can see these devices, initialize them, and allow you to create storage objects on them. However, this capability has not yet been fully tested.

IMPORTANT: In OES 11, a server hang or crash can occur if you attempt to use a Linux software RAID when you create storage objects that are managed by NSS management tools.

Linux software RAID devices do not support concurrent activation on multiple nodes, and are not cluster aware. They cannot be used for shared disk storage objects such as cLVM volume groups and the Novell Cluster Services SBD (split-brain detector) partitions.

NSS software RAID devices are supported for use with NSS pools and Novell Cluster Services SBD (split-brain-detector) partitions. Linux POSIX file systems are not supported on NSS software RAID devices.

IMPORTANT: Do not create linux partitions (or any non-Novell type partition) on an NSS software RAID device. Doing so causes all pool creations on that RAID device to fail.

5.3 Naming Conventions for Linux LVM Volume Group and Logical Volume Names

Consider the following conventions for naming Linux Logical Volume Manager (LVM) volume groups and logical volumes:

- [Section 5.3.1, “NSSMU and NLVM Requirements for LVM Names,” on page 30](#)
- [Section 5.3.2, “LVM Requirements for LVM Names,” on page 30](#)
- [Section 5.3.3, “Novell Cluster Services Requirements for Clustered LVM Volume Names,” on page 30](#)

5.3.1 NSSMU and NLVM Requirements for LVM Names

NLVM and NSSMU require that Linux LVM volume group names and logical volume names be unique from any volume, device, pool, RAID, and other Device Mapper name. The LVM group name is limited to 128 characters. The LVM logical volume name is limited to 64 characters.

If you enable NCP for the volume, see [Section 5.4, “Naming Conventions for NCP Volume Names,” on page 31](#).

5.3.2 LVM Requirements for LVM Names

LVM2 allows volume group names and logical volume names to contain characters A to Z, a to z, 0 to 9, underscore (_), hyphen (-), dot (.), and plus (+). The names cannot begin with a hyphen.

Reserved names and character strings that are used internally by LVM cannot be used as volume group names or logical volume names. A volume group cannot be called anything that exists in `/dev/` at the time of creation. It cannot be named `'.'` (a single dot) or `'..'` (double dot).

A logical volume cannot be named the following reserved words:

```
. (a single dot)
.. (double dot)
snapshot
pvmove
```

The logical volume name also cannot contain the following strings:

```
_mlog
_mimage
```

5.3.3 Novell Cluster Services Requirements for Clustered LVM Volume Names

In a Novell Cluster Services cluster, the names of clustered LVM volume groups and logical volumes must be unique across all nodes in the cluster.

5.4 Naming Conventions for NCP Volume Names

NCP volume names can be up to 14 alphanumeric characters, using uppercase letters A through Z and numbers 0 through 9. Underscores (`_`) are allowed.

If you NCP enable a Linux volume as you create it with NSSMU or the `nlvm create linux volume` command, the NCP volume name uses the specified Linux volume name, but all letters are capitalized. NCP treats the Linux volume name as case insensitive. Ensure that the specified Linux volume name does not exceed 14 characters, does not use special characters, and is unique across all nodes in the cluster for both Linux and NCP.

5.5 Linux LVM Volume Groups

NLVM uses the Linux Logical Volume Manager to create volume groups. LVM requires that the devices you use to create a volume group are already initialized and contain no partitions. LVM uses the entire device for the volume group.

5.6 Linux LVM Volume Group Cluster Resources

Novell Cluster Services 2.1, NLVM, and NSSMU use the Clustered Logical Volume Manager (CLVM) to manage LVM volume group cluster resources. CLVM requires the Linux kernel 2.6.32.45-0.3 or later. You can get the latest kernel version by using the SLES 11 SP2 update channel.

When you create clustered LVM volume groups on shared storage, all of the nodes in the cluster must have shared physical access to the devices that you want to use to create the volume group. A quorum of nodes must be present in the cluster. The volume group cluster resource is brought online on only one node at a time.

LVM requires that the devices you use to create a volume group are already initialized and contain no partitions. In a cluster, a device should be physically attached to all nodes in a cluster. The device must not be marked as Shareable for Clustering because that adds a 4 KB partition on the device to store the shared state. LVM uses the entire device for the volume group.

5.7 Using Antivirus Software with NCP Volumes

For information about using antivirus software with NCP volumes, see “[McAfee Antivirus Requires Additional Configuration](#)” in the *OES 11 SP1: Planning and Implementation Guide*.

6 Managing Linux Volumes with NSSMU

The Novell Storage Services (NSS) Management Utility (NSSMU) for Novell Open Enterprise Server (OES) 11 SP1 allows you to create and manage Linux POSIX file systems. This section describes how to use NSSMU to create and manage Linux POSIX volumes on your OES server.

NSSMU also allows you to create and manage NSS pools, volumes, software RAIDs, and pool snapshots. For information, see the *OES 11 SP1: NSS File System Administration Guide for Linux*.

- ♦ [Section 6.1, “NSSMU Quick Reference for Linux Volumes,” on page 34](#)
- ♦ [Section 6.2, “Launching and Navigating NSSMU,” on page 35](#)
- ♦ [Section 6.3, “Viewing a List of Devices,” on page 37](#)
- ♦ [Section 6.4, “Viewing Device Details,” on page 37](#)
- ♦ [Section 6.5, “Initializing a Device,” on page 39](#)
- ♦ [Section 6.6, “Unsharing a Device,” on page 41](#)
- ♦ [Section 6.7, “Viewing a List of Linux Volumes,” on page 41](#)
- ♦ [Section 6.8, “Viewing Details for a Linux Volume,” on page 42](#)
- ♦ [Section 6.9, “Creating a Non-LVM Linux Volume,” on page 44](#)
- ♦ [Section 6.10, “Creating an LVM Logical Volume,” on page 48](#)
- ♦ [Section 6.11, “Creating a Shared LVM Volume with Novell Cluster Services,” on page 52](#)
- ♦ [Section 6.12, “Mounting a Linux Volume,” on page 53](#)
- ♦ [Section 6.13, “Dismounting a Linux Volume,” on page 53](#)
- ♦ [Section 6.14, “Renaming the Mount Point Path for a Linux Volume,” on page 53](#)
- ♦ [Section 6.15, “Renaming a Non-LVM Linux Volume,” on page 55](#)
- ♦ [Section 6.16, “Renaming an LVM Volume,” on page 55](#)
- ♦ [Section 6.17, “Deleting a Linux Volume,” on page 56](#)

6.1 NSSMU Quick Reference for Linux Volumes

Use this section as a quick reference to manage devices, view partitions, and create and manage Linux volumes on your OES 11 SP1 server. To launch NSSMU, log in to the server as the root user, open a terminal console, then enter `nssmu`.

IMPORTANT: For information about the NSS file system management features of NSSMU, see the “[NSS Management Utility \(NSSMU\) Quick Reference](#)” in the *OES 11 SP1: NSS File System Administration Guide for Linux*.

Table 6-1 NSSMU Management Options for Linux Volumes

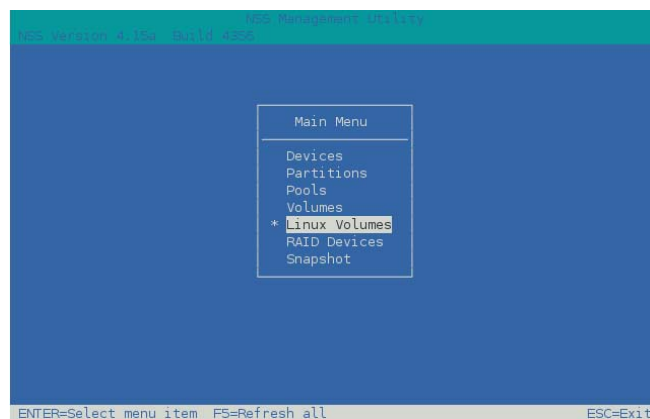
NSSMU Management Options	Description
<i>Devices</i>	Use the Devices page to initialize and maintain physical storage devices and NSS software RAID devices available to this server. Linux volumes cannot use NSS software RAID devices. NSSMU does not recognize native Linux software RAID devices.
F3 = Initialize device (Do not initialize your system device.)	
F5 = Refresh display	Initializing the selected device erases its partition table, effectively destroying all of its data. A device that has not been initialized reports a partition type of <i>Uninitialized</i> .
F6 = Share (shareable/not shareable for clustering)	
Space = Select/Unselect	When you initialize a device, you can select the DOS or the GUID Partition Table (GPT) partitioning scheme for a given device.
F8 = More	The DOS partitioning scheme supports devices up to 2 TB (terabytes) in size. It allows up to four partitions on a device.
Enter = Show partitions (does not report the 4K partition created for Shareable for Clustering)	The GPT partitioning scheme supports device sizes up to 2E64 sectors (that is, up to 8388608 petabytes (PB) based on the 512-byte sector size). It allows up to 128 partitions per disk. Each of its disks partitions is a logical device that is identified by a unique 128-bit (16-byte) GUID.
Esc = Previous menu	Use Shareable for Clustering only for devices you plan to use for a Novell Cluster Services SBD (split brain detector) partition or for a cluster enabled NSS pool. Clustered Linux LVM volumes require a device that is initialized, contains no partitions, and is not shared (Shareable for Clustering is set to <i>No</i>). The entire device is used for the volume group. Clustered LVM controls the share state of the device, not the NSS Shareable for Clustering setting.
<i>Partitions</i>	Use this option to display details about partitions. All types of partitions are displayed, including those for Linux POSIX file systems, NSS file systems, and NSS software RAID.
Ins = Create an NSS partition (disabled)	
Del = Delete an NSS partition	If you widen the NSSMU display screen, it widens the Partitions list panel and allows you to see the partitions' full names.
F3 = Mirror an NSS pool partition	
F5 = Refresh details of the partition	The Create option is disabled. Partitions are automatically created as you define Linux volumes, NSS pools, or NSS software RAID.
F6 = Label	
Enter = Show volumes	Do not use the delete and mirror options for Linux POSIX file systems. Partitions used for Linux volumes are destroyed automatically when you delete the volume from the Volumes page.
Esc = Previous menu	

NSSMU Management Options	Description
<i>Linux Volumes</i>	Use the Linux Volumes page to create one of the following types of Linux volumes:
Ins = Create	♦ Native Linux POSIX volumes (non-LVM)
Del = Delete	♦ LVM volume group and logical volume
F3 = Rename	♦ Clustered LVM volume group, logical volume, Novell Cluster Services cluster resource, and NCP virtual server (if NCP is enabled)
F5 = Refresh details of the Linux volumes	♦ NCP volume. Enabling NCP for a Linux volume automatically creates an NCP volume of the same name. The NCP share is created at the root of the volume. For a clustered volume, an NCP Virtual Server object is created for the cluster resource.
F7 = Mount or Dismount	
Esc = Previous menu	
	NSSMU supports Btrfs (requires <code>btrfsprogs</code>), Ext2, Ext3, ReiserFS, and XFS file systems.
	You can mount or dismount the Linux volume.
	You can delete a Linux volume. The volume must be dismounted (or offline for clustered volumes) before you can delete it.
	You can use Rename to modify the mount point for the volume.

6.2 Launching and Navigating NSSMU

- 1 Log in to the server as the `root` user, then open a terminal console.
- 2 Launch NSSMU by entering

`nssmu`



- 3 The NSSMU Main Menu page offers the following options:

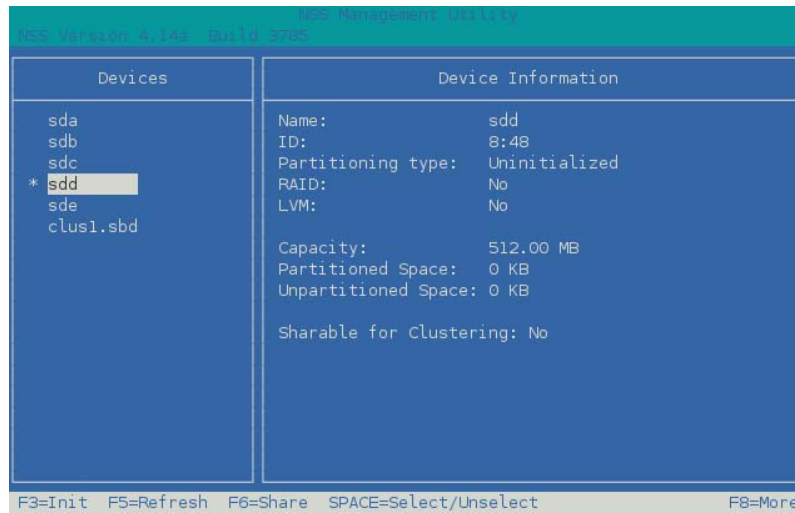
NSSMU Main Menu	Description
Devices	<p>View information about devices.</p> <p>Initialize a device in MSDOS or GPT format.</p> <p>Enable or disable a device as <i>Shareable for Clustering</i>. This is used for devices that you plan to use as the Split Brain Detector (SBD) for a Novell Cluster Services cluster or for NSS pools.</p>
Partitions	<p>View a list of partitions on a device, including Linux and LVM partitions.</p> <p>If you widen the NSSMU display screen, it widens the Partitions list panel and allows you to see the partitions' full names.</p>
Pools	<p>Create and manage NSS pools.</p> <p>Create and cluster-enable NSS pools on devices marked as Shareable for Clustering.</p>
Volumes	Create and manage NSS volumes.
Linux Volumes	<p>Create and manage Linux POSIX file systems, such as Btrfs (requires <i>btrfsprogs</i>), Ext2, Ext3, Reiser, and XFS.</p> <p>You can create the following volume types:</p> <ul style="list-style-type: none"> ♦ Native Linux POSIX volumes (non-LVM) ♦ LVM volume group and logical volume ♦ Clustered LVM volume group, logical volume, Novell Cluster Services cluster resource, and NCP virtual server (if NCP is enabled) ♦ NCP volume. Enabling NCP for a Linux volume automatically creates an NCP volume of the same name. The NCP share is created at the root of the volume. For a clustered volume, an NCP Virtual Server object is created for the cluster resource.
RAID Devices	Create and manage NSS software RAID devices.
Snapshots	Create and manage NSS pool snapshots.

- 4 Select a menu option by using the Up-arrow and Down-arrow keys, then press Enter.
- 5 When you are done working in a functional area, press Esc to return to the Main Menu page.
- 6 From the Main Menu page, press Esc to exit NSSMU.

6.3 Viewing a List of Devices

You can use NSSMU to view a list of the physical storage devices and NSS software RAID devices available to this server. Linux volumes cannot use NSS software RAID devices. NSSMU does not recognize native Linux software RAID devices.

- 1 From the NSSMU menu, select *Devices*, then press Enter.
- 2 Under *Devices*, view the device node names of the available devices.



- 3 Press the Up-arrow and Down-arrow keys on your keyboard to select a device and view information about it. For information, see [Section 6.4, “Viewing Device Details,” on page 37](#).
- 4 Press Esc twice to exit NSSMU.

6.4 Viewing Device Details

NSSMU allows you to view the following information about each device:


Parameter	Description
Name	Specifies the Linux device node name in the <code>/dev</code> directory. Examples: sda, sdb, sdc, sdd, and so on
ID	Specifies the major and minor number assigned to a block device on the server. Example: 8:0, 8:16, 8:32, 8:48

Parameter	Description
Partitioning type	<p>Specifies if the DOS or GPT partitioning scheme is used to format the device. A device that is not partitioned is reported as <i>Uninitialized</i>.</p> <p>A device that contains a non-mirrored Novell Cluster Services split-brain detector (SBD) partition is typically formatted as MS-DOS. A device that contains a mirrored SBD partition has a partition type of <i>NA</i> (not applicable) because it is a RAID 1 device.</p> <p>Values:</p> <ul style="list-style-type: none"> ♦ Uninitialized: The device is not partitioned. ♦ DOS: The device is partitioned in MS-DOS format. ♦ GPT: The device is partitioned in GUID Partition Table format. ♦ CSM: The device contains a legacy Cluster Segment Manager partition that was created by using the Enterprise Volume Management System (EVMS) on an OES 2 cluster node, and imported for use on OES 11 (or later) cluster nodes. ♦ LVM2: The entire device is an LVM2 volume group that contains an LVM2 logical volume. For example, the device is used for a clustered LVM2 volume group and logical volume. ♦ NA: The device is an NSS or SBD RAID 1 mirrored device that does not have a partitioning scheme.
RAID	<p>Specifies whether the device is an NSS software RAID device.</p> <p>Values: Yes or No</p>
LVM	<p>Specifies whether the device contains an LVM volume group.</p> <p>Values: Yes or No</p>
Capacity	Specifies the total amount of space available on the device.
Partitioned space	Specifies the amount of space in use on the device.
Unpartitioned space	Specifies the amount of free unpartitioned space available on the device. For MSDOS devices, free space includes unallocated space in an extended DOS partition.
Shareable for clustering	<p>The device can be used for shared NetWare partitions. Use with Novell Cluster Services SBD partitions and cluster-enabled NSS pools. The status is stored on a 4 KB partition on the device.</p> <p>Values: Yes or No (default)</p>

To view device details:

- 1 From the NSSMU menu, select *Devices*, then press Enter.
- 2 Select the device that you want to view by using the Up-arrow and Down-arrow keys.

- 3 View details about a selected device under *Device Information*.



- 4 Press Esc twice to exit NSSMU.

6.5 Initializing a Device

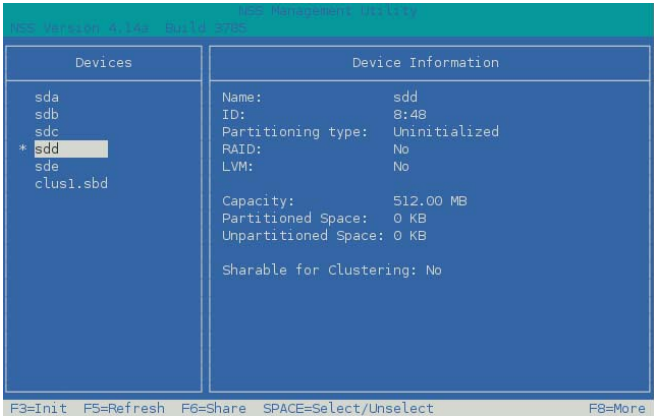
You must initialize a device to set up its device format before you can create volumes on it. You can also initialize a device to wipe its current structure and reconfigure it. Devices that you want to use for Linux POSIX volumes should be in an unshared state whether they are local or SAN devices.

IMPORTANT: NLVM and NSSMU do not support using Linux software RAID devices and NSS software RAID devices with Linux POSIX file systems. You can use a hardware RAID device to achieve device fault tolerance for Linux POSIX volumes.

Initializing a device formats it with an MSDOS or a GPT partitioning scheme. MSDOS supports devices up to 2 TB in size. GPT supports devices of any size. The default partitioning scheme is MSDOS for devices less than 2 TB, and GPT for devices greater than or equal to 2 TB. NSSMU automatically selects the default value in the menu. If you do not want to use the default, you can select the other option. If you specify MSDOS for a device that has a size greater than or equal to 2 TB, the device's free space is limited to 2 TB; the remainder of the device is unusable space.

WARNING: Initializing a device removes all partitions and data from the device. Do not initialize the device that contains the operating system.


- 1 From the NSSMU menu, select *Devices*, then press Enter.
- 2 Select the device that you want to initialize by using the Up-arrow and Down-arrow keys.
Information about a selected device is displayed under *Device Information*.
If a device has never been initialized, its *Partitioning type* is reported as *Uninitialized*.



- 3 Press F3 to initialize the device.
- 4 Read the warning message, then press y (Yes) to continue, or press n (No) to cancel.



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



- 6 View the *Device Information* to confirm that the partitioning type is applied and the unpartitioned space shows free space.

6.6 Unsharing a Device

The *Shareable for Clustering* option on the NSSMU Devices page is intended for marking devices as shareable that you plan to use for cluster-enabled NSS pools or for Novell Cluster Services SBD (split brain detector) partitions.

For Linux volumes, you use devices that are not marked as shareable for clustering. When you cluster-enable an LVM volume group for Novell Cluster Services, clustered LVM controls the share state of the device, not the *Shareable for Clustering* setting. If a device was previously used for an NSS pool or SBD partition and is marked as *Shareable for Clustering*, it contains a small 4 KB partition on it to store the state. Before you can create a clustered LVM volume group on the device, you must change the device's *Shareable for Clustering* state from *Yes* to *No*.

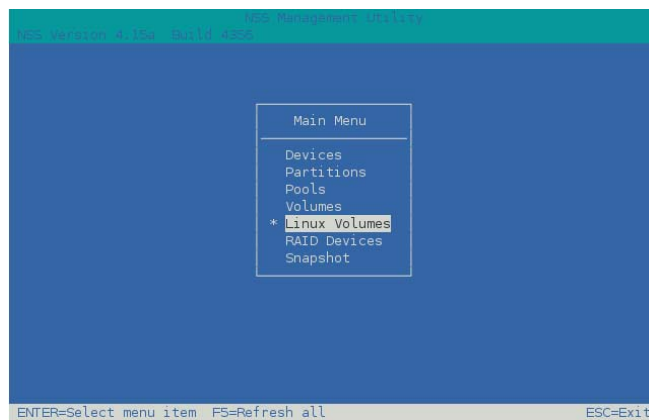
To unshare a device:

- 1 From the NSSMU menu, select *Devices*, then press Enter.
- 2 Select the device that you want to unshare by using the Up-arrow and Down-arrow keys.
- 3 Press F6 to change the share state of the device from *Yes* to *No*.
- 4 Press Esc twice to exit NSSMU.

6.7 Viewing a List of Linux Volumes

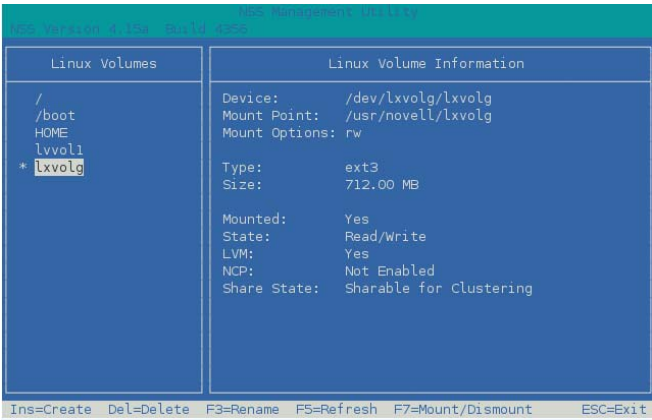
You can use the NSSMU Linux Volumes page to view information about Linux POSIX and LVM logical volumes that are mounted on the system.

- 1 From the NSSMU menu, select *Linux Volumes*, then press Enter.



- 2 Under *Linux Volumes*, view the volume name of the available volumes.

Linux POSIX volumes show the volume name in all caps, such as HOME. Linux LVM logical volumes show the volume name as the name you gave the volume when you created it, such as lvvol1 or lxvolg.



- 3 Use the Up arrow and Down arrow to select a volume and view details about it.
For information about the fields, see [Section 6.8, “Viewing Details for a Linux Volume,”](#) on [page 42](#).
- 4 Press Esc twice to exit NSSMU.

6.8 Viewing Details for a Linux Volume

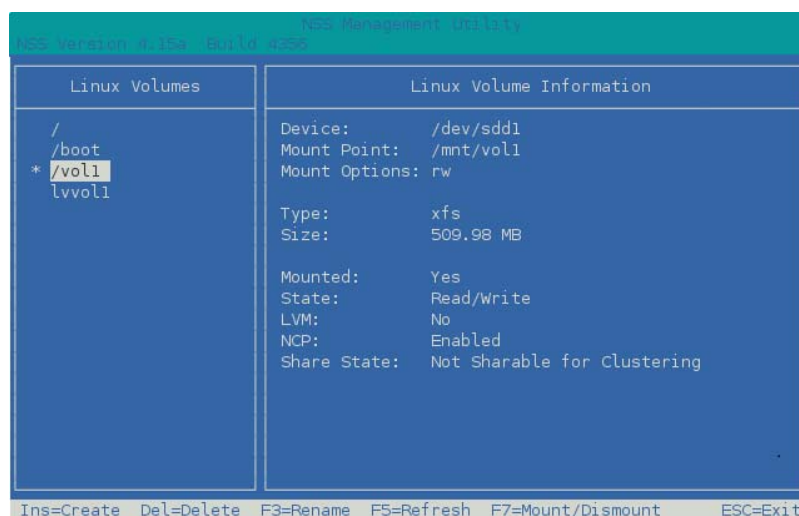
NSSMU allows you to view the following information about each Linux volume:

Parameter	Description
Device	Specifies the full device node path. Example for a Linux POSIX volume: /dev/sda2 Example for an LVM volume: /dev/lvvol1/lvvol1
Mount Point	Specifies the path on the root file system where this volume is mounted. Examples: /mnt/vol1 /home /media/ext3/ext3-1 /mnt/lvvol1
Mount options	Specifies the mount options that are applied whenever this volume is automatically mounted after a reboot. Example: rw
Type	Specifies the file system type. Examples: btrfs (requires btrfsprogs), ext2, ext3, reiserfs, xfs
Size	Specifies the amount of space reserved for this volume. Example: 30.58 GB

Parameter	Description
Mounted	Specifies whether the volume is mounted or unmounted. Value: Yes or No
State	Specifies the availability for the file system. Example: Read/Write
LVM	Specifies whether the specified volume is an LVM logical volume. Value: Yes or No
NCP enabled	Specifies whether the volume is enabled for NCP (NetWare Core Protocol). The NCP volume name is the same as the Linux volume name, and the NCP share is created at the root of the volume. If the volume is cluster-enabled, this information is not reported when the cluster resource is in an offline or comatose state. Values: Yes or No
Share State	Specifies whether the volume is cluster enabled for Novell Cluster Services. Values: Shareable for Clustering or Not Shareable for Clustering

- 1 From the NSSMU main menu, select *Linux Volumes*, then press Enter.
- 2 Select a volume from the list of volumes to view details about it.

For example, the following screen shows details about a native Linux POSIX volume with an XFS file system:



The following screen shows details about an unshared LVM logical volume with an Ext3 file system:

```

NSS Management Utility
NSS Version 4.15a Build 4339

Linux Volumes                                Linux Volume Information
-----
/
/boot
/vol1
* lvvol1

Device:      /dev/lvvol1/lvvol1
Mount Point: /mnt/lvvol1
Mount Options: rw

Type:       ext3
Size:       508.00 MB

Mounted:    Yes
State:      Read/Write
LVM:        Yes
NCP:        Enabled
Share State: Not Sharable for Clustering

Ins=Create Del=Delete F3=Rename F5=Refresh F7=Mount/Dismount ESC=Exit

```

The following screen shows details about a clustered LVM logical volume with an Ext3 file system. The Linux Clustered Logical Volume Manager (CLVM) controls and reports the volume group's Share Status as *Shareable for Clustering*.

```

NSS Management Utility
NSS Version 4.15a Build 4339

Linux Volumes                                Linux Volume Information
-----
/
/boot
* vol44

Device:      /dev/vol44/vol44
Mount Point: /mnt/vol44
Mount Options: rw

Type:       ext3
Size:       99.98 GB

Mounted:    Yes
State:      Read/Write
LVM:        Yes
NCP:        Not Enabled
Share State: Shareable for Clustering

Ins=Create Del=Delete F3=Rename F5=Refresh F7=Mount/Dismount ESC=Exit

```

- 3 Press Esc twice to exit NSSMU.

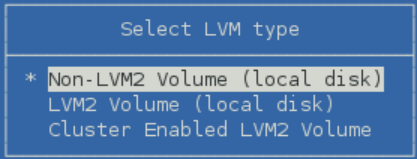
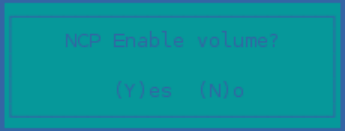
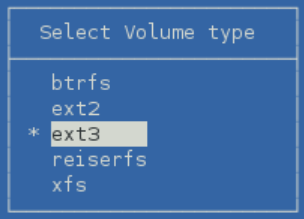
6.9 Creating a Non-LVM Linux Volume

You can use NSSMU to create a native Linux POSIX volume on a device. You must have free unpartitioned space available on a device. The device should not be shareable for clustering.

- 1 From the NSSMU main menu, select *Devices*, then press Enter.
- 2 View the devices to determine which device you want to use for the volume and the amount of space available on it.
- 3 If the device has not been initialized, or if you need to reformat the device, initialize the device.
For information, see [Section 6.5, "Initializing a Device," on page 39](#).

WARNING: Do not initialize a device that contains data you want to keep.

- 4 From the NSSMU main menu, select *Linux Volumes*, then press Enter.
- 5 Press Insert to begin creating a new Linux volume.
- 6 When you are prompted, specify the following parameters for the volume:

Parameter	Description
LVM Type	<p>Select <i>Non-LVM2 Volume</i> from the following volume types to create a native Linux POSIX volume, then press Enter.</p> <ul style="list-style-type: none">♦ Non-LVM2 Volume (local disk)♦ LVM2 Volume (local disk)♦ Cluster Enabled LVM2 Volume 
NCP	<p>Enables the Linux POSIX file system on the volume to be accessed with the NetWare Control Protocol (NCP). It creates an NCP volume of the same name, and creates an NCP share at the root of the volume.</p> 
File System Type	<p>Select one of the following file system options to make on the volume, then press Enter.</p> <ul style="list-style-type: none">♦ btrfs (requires btrfsprogs)♦ ext2♦ ext3♦ reiserfs♦ xfs 

Parameter	Description
Mount Point	<p>Specify the full Linux path where you want to mount the device, then press Enter. For example, <code>/mnt/vol1</code>.</p> <pre>Enter full mountpoint path: /home</pre> <p>The last directory is assumed to be the name you want to give the volume. The name must be a unique volume name on the server.</p> <p>NOTE: If you enable NCP for the volume, the final folder name in the path must comply with the limitations for NCP volume names described in Section 5.4, “Naming Conventions for NCP Volume Names,” on page 31. If you use lowercase letters for the volume name, they are automatically changed to uppercase for the NCP volume name.</p>
Make Options	<p>(Optional) Specify the make options to use for your selected file system, then press Enter.</p> <pre>Enter any make options:</pre> <p>For information, see the <code>mkfs(8)</code>, <code>mkfs.btrfs(8)</code>, <code>mkfs.ext2(8)</code>, <code>mkfs.ext3(8)</code>, <code>mkfs.reiser(8)</code>, and <code>mkfs.xfs(8)</code> man pages.</p> <p>If a specified make option is invalid for the specified file system, the file system creation fails.</p>
Mount Options	<p>Specify the default set of options to use when mounting the volume after a reboot, then press Enter.</p> <pre>Enter any mount options: rw</pre> <p>The Read/Write (<code>rw</code>) mount option is specified by default.</p> <p>For information about mount options that are available for each file system, see the <code>mount(8)</code> man page.</p>

Parameter	Description
Devices	<p>Use the Up-arrow and Down-arrow keys to highlight a device from the list, then press Insert or Enter to select the device.</p> <p>IMPORTANT: NSSMU does not support using Linux software RAID devices and NSS software RAID devices with Linux POSIX file systems. You can use a hardware RAID device to achieve device fault tolerance for Linux POSIX volumes.</p>

Device	Free Size	Shared	Selected
			Total 0 KB
Device	Free Size	Shared	
sda	1655.13 MB	No	
sdb	398.98 MB	No	
sdc	101.98 MB	No	
* sdd	511.95 MB	No	
sde	101.98 MB	No	

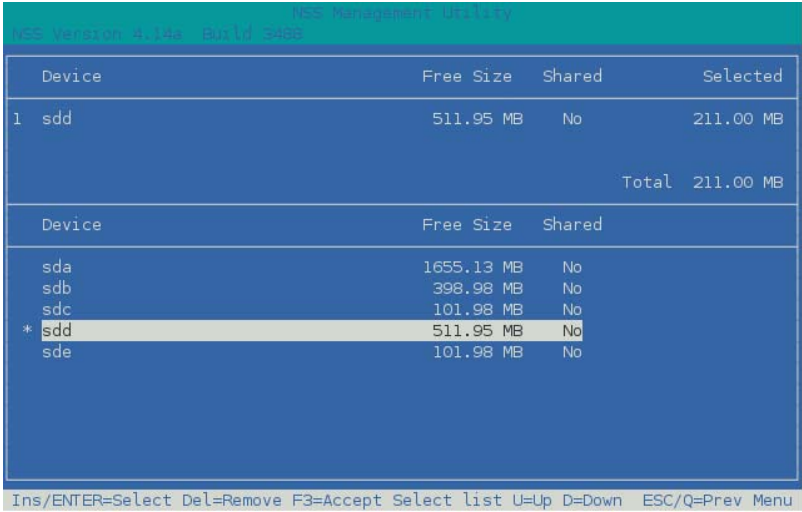
Ins/ENTER=Select Del=Remove F3=Accept Select list U=Up D=Down ESC/Q=Prev Menu

Partition Size	<p>For non-LVM volumes and LVM2 volumes, specify the maximum amount of space in MB to use from the selected device, then press Enter.</p> <p>You can use part or all of the available space on the selected device.</p> <p>For cluster-enabled LVM2 volumes, this option is not available. The entire device is automatically used by the LVM volume group.</p>
----------------	---

Device	Free Size	Shared	Selected
			Total 0 KB
Device	Free Size	Shared	
sda		No	
sdb		No	
sdc	101.98 MB	No	
* sdd	511.95 MB	No	
sde	101.98 MB	No	

Enter size in MB for this segment ESC=Prev Menu

Parameter	Description
Confirm	Select F3 to accept and create the Linux POSIX volume.



```

NSS Management Utility
NSS Version 4.14a Build 3433

Device      Free Size  Shared  Selected
-----
1 sdd       511.95 MB  No      211.00 MB
Total      211.00 MB

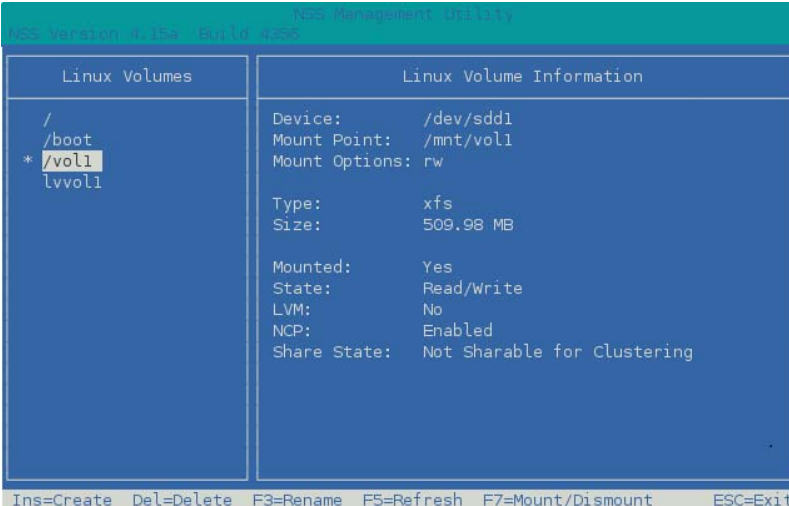
Device      Free Size  Shared
-----
sda         1655.13 MB No
sdb         398.98 MB No
sdc         101.98 MB No
* sdd       511.95 MB No
sde         101.98 MB No

Ins/ENTER=Select Del=Remove F3=Accept Select list U=Up D=Down ESC/Q=Prev Menu

```

7 View the details of the newly created volume.

For information, see [Section 6.8, “Viewing Details for a Linux Volume,”](#) on page 42.



Linux Volumes	Linux Volume Information
<pre> / /boot * /vol1 lvvol1 </pre>	<pre> Device: /dev/sdd1 Mount Point: /mnt/vol1 Mount Options: rw Type: xfs Size: 509.98 MB Mounted: Yes State: Read/Write LVM: No NCP: Enabled Share State: Not Sharable for Clustering </pre>

Ins=Create Del=Delete F3=Rename F5=Refresh F7=Mount/Dismount ESC=Exit

8 Press Esc twice to exit NSSMU.

6.10 Creating an LVM Logical Volume

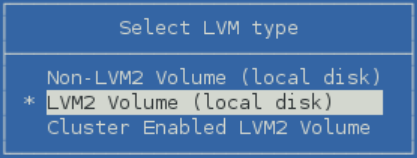
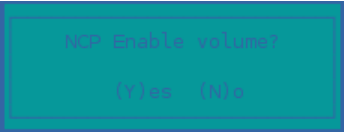

You can use NSSMU to create a Linux LVM volume group and logical volume on a device, make a file system on the volume, and mount the volume. You must have free unpartitioned space available on a device. The device should not be shareable for clustering.

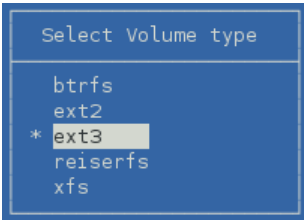



- 1 From the NSSMU main menu, select *Devices*, then press Enter.
- 2 View the devices to determine which device you want to use for the volume and the amount of space available on it.

- 3 If the device has not been initialized, or if you need to reformat the device, initialize the device. For information, see [Section 6.5, “Initializing a Device,” on page 39](#).

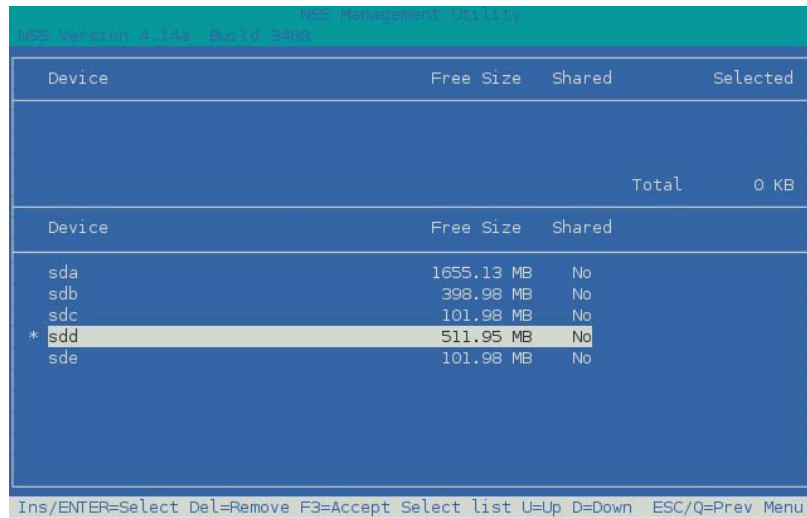
WARNING: Do not initialize a device that contains data you want to keep.

- 4 From the NSSMU main menu, select *Linux Volumes*, then press Enter.
- 5 Press Insert to begin creating a new Linux LVM logical volume.
- 6 When you are prompted, specify the following parameters for the volume:

Parameter	Description
LVM Type	<p>Select <i>LVM2 Volume</i> from the following volume types to create an LVM volume group and volume, then press Enter.</p> <ul style="list-style-type: none">♦ Non-LVM2 Volume (local disk)♦ LVM2 Volume (local disk)♦ Cluster Enabled LVM2 Volume 
NCP	<p>Enables the Linux POSIX file system on the volume to be accessed with the NetWare Control Protocol (NCP). It creates an NCP volume of the same name, and creates an NCP share at the root of the volume.</p> 
Volume name	<p>Type a name to use for the LVM volume group and logical volume (such as <code>lvvol1</code>), then press Enter.</p>  <p>NOTE: If you enable NCP for the volume, the LVM volume name must comply with the limitations for NCP volume names described in Section 5.4, “Naming Conventions for NCP Volume Names,” on page 31. If you use lowercase letters for the volume name, they are automatically changed to uppercase for the NCP volume name.</p>


Parameter	Description
File System Type	<p>Select one of the following file system options to make on the volume, then press Enter:</p> <ul style="list-style-type: none"> ♦ btrfs (requires btrfsprogs) ♦ ext2 ♦ ext3 ♦ reiserfs ♦ xfs 
Mount Point	<p>Specify the full Linux path where you want to mount the device, then press Enter. The default path is <code>/usr/novell/<lv_volume_name></code>.</p>  <p>For example, type <code>/mnt/lvol1</code>, then press Enter.</p> <p>The mount path's final directory can be the same or different than the name used for the logical volume.</p>
Make Options	<p>(Optional) Specify the make options to use for your selected file system, then press Enter.</p>  <p>For information, see the <code>mkfs(8)</code>, <code>mkfs.btrfs(8)</code>, <code>mkfs.ext2(8)</code>, <code>mkfs.ext3(8)</code>, <code>mkfs.reiser(8)</code>, and <code>mkfs.xfs(8)</code> man pages.</p> <p>If a specified make option is invalid for the specified file system, the file system creation fails.</p>
Mount Options	<p>Specify the default set of options to use when mounting the volume after a reboot, then press Enter.</p>  <p>The Read/Write (<code>rw</code>) mount option is specified by default.</p> <p>For information about mount options that are available for each file system, see the <code>mount(8)</code> man page.</p>

Parameter	Description
Devices	Use the Up-arrow and Down-arrow keys to highlight a device from the list, then press Insert or Enter to select the device.



Partition Size	Specify the maximum amount of space in MB to use from the selected device, then press Enter.
----------------	--

You can use part or all of the available space on the selected device.



Parameter	Description
Confirm	Select F3 to accept and create the Linux POSIX volume.

```

NSS Management Utility
NSS Version 4.14a Build 3489

Device      Free Size  Shared  Selected
1 sdd        511.95 MB  No      211.00 MB
Total      211.00 MB

Device      Free Size  Shared
sda         1655.13 MB  No
sdb         398.98 MB  No
sdc         101.98 MB  No
* sdd        511.95 MB  No
sde         101.98 MB  No

Ins/ENTER=Select Del=Remove F3=Accept Select list U=Up D=Down ESC/Q=Prev Menu

```

- 7 View the details of the newly created volume.
- For information, see [Section 6.8, “Viewing Details for a Linux Volume,”](#) on page 42.

```

NSS Management Utility
NSS Version 4.13a Build 4359

Linux Volumes      Linux Volume Information

/
/boot
/voll
* lvvoll1

Device:      /dev/lvvol1/lvvol1
Mount Point: /mnt/lvvol1
Mount Options: rw

Type:        ext3
Size:        508.00 MB

Mounted:     Yes
State:       Read/Write
LVM:         Yes
NCP:         Enabled
Share State: Not Sharable for Clustering

Ins=Create Del=Delete F3=Rename F5=Refresh F7=Mount/Dismount ESC=Exit

```

- 8 Press Esc twice to exit NSSMU.

6.11 Creating a Shared LVM Volume with Novell Cluster Services

For information about creating a clustered LVM volume group cluster resource with NSSMU, see [Section 8.3.2, “Creating an LVM Volume Group Cluster Resource with NSSMU,”](#) on page 83 in Chapter 8, “Clustering LVM Volume Groups with Novell Cluster Services,” on page 75.

6.12 Mounting a Linux Volume

Use the *Mount/Dismount* option on the NSSMU Linux Volumes page to mount an unmounted Linux volume.

IMPORTANT: If an LVM logical volume is clustered, you must use the `cluster online` command to allow the load script to mount the volume. You can also use the *Clusters > Cluster Manager* page in Novell iManager to online a resource.

- 1 From the NSSMU menu, select *Linux Volumes*, then press Enter.
- 2 In the *Linux Volumes* list, select the volume.
- 3 In the volume's details, view the volume's current mount status.
The *Mounted* status is *No* if the volume is not mounted.
- 4 Press F7 to mount the volume.
- 5 Select the volume, then view its details to verify that the *Mounted* status changes from *No* to *Yes*.
- 6 Press Esc twice to exit NSSMU.

6.13 Dismounting a Linux Volume

Use the *Mount/Dismount* option on the NSSMU Linux Volumes page to dismount a mounted Linux volume.

IMPORTANT: If an LVM logical volume is clustered, you must use the Novell Cluster Services `cluster offline` command to allow the unload script to dismount the volume. You can also use the *Clusters > Cluster Manager* page in Novell iManager to offline a resource.

- 1 From the NSSMU menu, select *Linux Volumes*, then press Enter.
- 2 In the *Linux Volumes* list, select the volume.
- 3 In the volume's details, view the volume's current mount status.
The *Mounted* status is *Yes* if the volume is mounted.
- 4 Press F7 to dismount the volume.
- 5 Select the volume, then view its details to verify that the *Mounted* state changes from *Yes* to *No*.
- 6 Press Esc twice to exit NSSMU.

6.14 Renaming the Mount Point Path for a Linux Volume

The *Rename* option on the NSSMU Linux Volumes page allows you to modify the mount point of a volume. You cannot rename the mount point path for the root (/) volume or other system volumes such as /boot.

An LVM volume name does not depend on the mount point path. For an LVM volume, modifying the mount point path in any way does not effect the LVM volume group name and logical volume name.

Renaming the mount point path for a non-LVM volume might also cause the volume name to change, depending on the change you make to the path. A native Linux POSIX volume assumes its name from the final directory in the mount point path. If you change the name of the final directory, the new directory name also becomes the new name of the volume.

The following examples demonstrate how renaming the mount point path for a non-LVM volume can affect the volume name:


Volume Information	New Mount Point and Change	New Volume Name
Volume Type: Non-LVM Mount Point: /projects/volp1 Volume Name: /volp1	Change the path to the final directory. Use the same name for the final directory. /old/projects/volp1	The volume name does not change.
Volume Type: Non-LVM Mount Point: /home/janedoe Volume Name: /janedoe	Change the final directory name. The preceding directories can be the same or different. /home/janebuck or /users/janebuck	The new name for the final directory becomes the new volume name: /janebuck The final directory name must be unique as a volume name on the server.

The volume is dismounted in order to apply the new mount point path to the volume. The volume is not automatically mounted afterwards.

WARNING: To avoid potential data corruption, you should dismount the volume gracefully before you rename it.

Use the procedure in this section to rename the mount point path for a non-clustered Linux volume. For information about modifying the mount point path for a clustered LVM volume, see [Section 8.7, “Renaming the Mount Point Path for a Clustered LVM Volume,”](#) on page 101.

- 1 From the NSSMU main menu, select *Linux Volumes*, then press Enter.
- 2 In the *Linux Volumes* list, browse each volume to verify that the mount point you want to use is not already in use by another volume.
- 3 Select the volume of interest.
- 4 If the volume is mounted, press F7 to dismount the volume.
- 5 Select the volume again, verify that it is not mounted, then press F3 to rename the mount point path.
- 6 Read the caution message, then press y (Yes) to continue, or press n (No), or press Esc to cancel the task.



- 7 In the *New Mount Point* field, specify the full Linux path of the new mount point that you want to use, then press Enter.

The original mount point path is automatically populated in the *New Mount Point* field. Modify it as needed.



```
New mount point: /home/bob2
```

- 8 Select the volume, then view its details to verify that the mount point has changed.
If the final directory was changed in the mount point path of a non-LVM volume, notice that the volume name has changed. The volume is not automatically mounted.
- 9 Select the volume, then press F7 to mount the volume.
- 10 Select the volume, then view its details to verify that the *Mounted* status is *Yes*.

6.15 Renaming a Non-LVM Linux Volume

The *Rename* option on the NSSMU Linux Volumes page allows you to modify the mount point of a volume. NSSMU uses the name of the last directory in a non-LVM mount point path as the volume name for the non-LVM volume. For example, if you create a non-LVM volume with the mount point path `/usr/novell/lxvol1`, the last directory name is `lxvol1`, and the non-LVM Linux volume becomes `lxvol1`.

Renaming the last directory in a mount point path for a non-LVM volume effectively renames the volume. You should gracefully dismount the volume before using the *Rename* option in NSSMU. For information about using the *Rename* option, see [Section 6.14, “Renaming the Mount Point Path for a Linux Volume,” on page 53](#). After you rename the volume successfully, you can mount the volume.

6.16 Renaming an LVM Volume

NSSMU does not provide a way to modify the name of an LVM volume. You can use the Linux `lvrename` command to rename an existing LVM logical volume in a volume group from the old volume name to the new volume name. The command does not change the volume group name. You should gracefully dismount the volume before attempting to rename it.

- 1 Log in to the server as the `root` user, then open a terminal console.
- 2 Launch NSSMU and dismount the logical volume:
 - 2a At the command prompt, enter

```
nssmu
```
 - 2b Select *Linux Volumes* and press Enter.
 - 2c Select the LVM volume, then press F7 Dismount.
 - 2d Select the volume, then view its details to verify that the *Mounted* state changes from *Yes* to *No*.
 - 2e Press Esc twice to exit NSSMU.
- 3 At the command prompt, enter

```
lvrename </dev/vg_name/old_lv_name> </dev/vg_name/new_lv_name>
```

Replace `vg_name` with the name of the volume group. If you created the LVM logical volume in NSSMU, the volume group name is the same as the logical volume name. If you created the LVM logical volume with NLVM or with LVM commands, you might have used a different name for the volume group.

Replace *old_lv_name* with the current name of the LVM logical volume.

Replace *new_lv_name* with the new name of the LVM logical volume.

For example, to change the name of the logical volume on volume group *vghome* from *lv_users* to *lv_home*, enter

```
lvrename /dev/vghome/lv_users /dev/vghome/lv_home
```

- 4 (Optional) Rename the mount point path that you use for the logical volume. The mount point should still be unmounted when you perform this task.

For information, see [Section 6.14, “Renaming the Mount Point Path for a Linux Volume,”](#) on page 53.

- 5 In a text editor, modify the */etc/fstab* file with the new logical volume name. You should also specify the new mount point if you renamed it.

- 6 Launch NSSMU and mount the renamed LVM logical volume:

6a At the command prompt, enter

```
nssmu
```

6b Select *Linux Volumes* and press Enter.

6c Select the LVM volume, then press F7 Mount.

6d Select the volume, then view its details to verify that the Mounted state changes from *No* to *Yes*.

6e Press Esc twice to exit NSSMU.

6.17 Deleting a Linux Volume

You can use the *Delete* option on the NSSMU Linux Volumes page to delete an unclustered Linux volume. The volume must be unmounted.

Deleting the Linux volume does not automatically delete the NCP volume and its Volume object. You should remove the NCP volume and its Volume object before you delete the Linux volume.

WARNING: The delete process removes the partition for the volume, and destroys all of the data on the partition.

Use the procedure in this section to delete unclustered Linux volumes. For information about deleting a clustered LVM volume group and logical volume, see [Section 8.10, “Deleting a Clustered LVM Volume Group and Logical Volume,”](#) on page 106.

- 1 If you enabled the Linux volume for NCP, use Novell Remote Manager to dismount the NCP volume, then remove the NCP volume and its Volume object.

For information, see “[Removing an NCP Volume](#)” in the *OES 11 SP1: NCP Server for Linux Administration Guide*.

- 2 From the NSSMU menu, select *Linux Volumes*, then press Enter.
- 3 In the *Linux Volumes* list, select the volume.
- 4 If the volume is mounted, press F7 to dismount it.
- 5 In the *Linux Volumes* list, select the volume, then verify that the Mounted state is *No*.
- 6 Press *Delete* to delete the volume.
- 7 Press *y* (yes) to confirm, or press *n* (no), or press Esc to cancel the delete process.

- 8** In the *Linux Volumes* list, verify that the volume is no longer displayed.
- 9** Press Esc to return to the main menu, then select *Devices*.
- 10** Select the device that contained the volume, then verify that the amount of free space increased.
- 11** Press Esc twice to exit NSSMU.

7 Managing Linux Volumes with NLVM Commands

The Novell Linux Volume Manager (NLVM) command line interface can be used to create and manage Linux POSIX file systems. For information about the syntax and options for the NLVM commands used in this section, see the [OES 11 SP1: NLVM Reference](#).

- ♦ [Section 7.1, “NLVM Commands Quick Reference for Linux Volumes,” on page 60](#)
- ♦ [Section 7.2, “Viewing a List of Devices,” on page 62](#)
- ♦ [Section 7.3, “Viewing Device Details,” on page 62](#)
- ♦ [Section 7.4, “Initializing a Device,” on page 63](#)
- ♦ [Section 7.5, “Unsharing a Device,” on page 65](#)
- ♦ [Section 7.6, “Viewing a List of Linux Volumes,” on page 65](#)
- ♦ [Section 7.7, “Creating a Linux POSIX Volume,” on page 66](#)
- ♦ [Section 7.8, “Creating an LVM Logical Volume,” on page 67](#)
- ♦ [Section 7.9, “Creating a Shared LVM Logical Volume with Novell Cluster Services,” on page 69](#)
- ♦ [Section 7.10, “Mounting Linux Volumes,” on page 69](#)
- ♦ [Section 7.11, “Dismounting Linux Volumes,” on page 71](#)
- ♦ [Section 7.12, “Rescanning for Storage Objects on Known Devices,” on page 72](#)
- ♦ [Section 7.13, “Deleting a Linux POSIX Volume,” on page 72](#)
- ♦ [Section 7.14, “Additional Information,” on page 74](#)

7.1 NLVM Commands Quick Reference for Linux Volumes

Use this section as a quick reference for the NLVM commands that support Linux POSIX volumes. Enter commands in a terminal console or script as the `root` user. For details about each command, see the *OES 11 SP1: NLVM Reference*.

Task	Command
<p>Create a clustered Linux LVM volume group and logical volume.</p> <p>Specify the <code>ncp</code> option to enable NCP for the volume and create an NCP share at the root of the volume, and to automatically create the NCP Server storage object for the cluster resource.</p> <p>See also Chapter 8, “Clustering LVM Volume Groups with Novell Cluster Services,” on page 75.</p>	<pre>nlvm create linux volume type=<btrfs ext2 ext3 reiserfs xfs> device=<devicename> mp=<full_mount_point_path> [mntopt=<fs_mount_options>] lvm name=<lv_name> [group=<vg_name>] shared ip=<resource_ip_address> [ncp]</pre> <p>Example:</p> <pre>nlvm create linux volume type=xfs device=sde mp=/mnt/vol42 mntopt=rw lvm name=lvmvol42 shared ip=10.10.10.42</pre>
<p>Create a Linux LVM volume group and logical volume.</p> <p>Specify the <code>ncp</code> option to enable NCP for the volume and create an NCP share at the root of the volume.</p> <p>See also Section 7.8, “Creating an LVM Logical Volume,” on page 67.</p>	<pre>nlvm create linux volume type=<btrfs ext2 ext3 reiserfs xfs> device=<devicename> size=<value[K M G T]> mp=<full_mount_point_path> [mntopt=<fs_mount_options>] lvm name=<lv_name> [group=<vg_name>] [ncp]</pre> <p>Example:</p> <pre>nlvm create linux volume type=xfs device=sde size=511M mp=/mnt/vol42 mntopt=rw lvm name=lvmvol42 group=lvmvg42</pre>

Task	Command
<p>Create a Linux POSIX volume.</p> <p>Specify the ncp option to enable NCP for the volume and create an NCP share at the root of the volume.</p> <p>See also Section 7.7, “Creating a Linux POSIX Volume,” on page 66.</p>	<pre>nlvm create linux volume type=<btrfs ext2 ext3 reiserfs xfs> device=<device_name anydisk> size=<value[K M G T] max> mp=<full_mount_point_path> [mntopt=<fs_mount_options>] [ncp]</pre> <p>Example:</p> <pre>nlvm create linux volume type=ext3 device=sda size=20G mp=/mnt/vol1 mntopt=rw</pre>
<p>Delete a Linux volume</p> <p>See also Section 7.13.1, “Deleting a Linux POSIX Volume,” on page 72.</p> <p>See also Section 7.13.2, “Deleting an LVM Volume Group and Logical Volume,” on page 73.</p> <p>See also Section 8.10, “Deleting a Clustered LVM Volume Group and Logical Volume,” on page 106.</p>	<pre>nlvm delete linux volume <volume_name></pre> <p>Example: Linux POSIX volume</p> <pre>nlvm delete linux volume /vol1</pre> <p>Example: Linux LVM logical volume</p> <pre>nlvm delete linux volume lvvol1</pre>
<p>Initialize a device</p> <p>See also Section 7.4, “Initializing a Device,” on page 63.</p>	<pre>nlvm [--force] init <device_name> [format=<gpt msdos>] unshared</pre> <p>Example:</p> <pre>nlvm init sde format=msdos unshared</pre>
<p>Rescan for storage objects on known devices</p> <p>See also Section 7.12, “Rescanning for Storage Objects on Known Devices,” on page 72.</p>	<pre>nlvm [-m] rescan</pre>
<p>Unshare a device</p> <p>See also Section 7.5, “Unsharing a Device,” on page 65.</p>	<pre>nlvm unshare <device></pre> <p>Example:</p> <pre>nlvm unshare sde</pre>
<p>View a list of active devices</p> <p>See also Section 7.2, “Viewing a List of Devices,” on page 62.</p>	<pre>nlvm list devices [more] [exclude=<raid nonraid shared nonshared lvm nonlvm>]</pre>
<p>View a list of Linux volumes</p> <p>See also Section 7.6, “Viewing a List of Linux Volumes,” on page 65.</p>	<pre>nlvm list linux volumes</pre>
<p>View details for a device</p> <p>See also Section 7.3, “Viewing Device Details,” on page 62.</p>	<pre>nlvm list <device></pre> <p>Example:</p> <pre>nlvm list device sde</pre>

7.2 Viewing a List of Devices

You can use the `nlvm list devices` command to view information about each of the active devices on the system. For command usage information, see “[List Devices](#)” in the [OES 11 SP1: NLVM Reference](#).

The command returns the following information:

Parameter	Description	Sample Values
Device name	Displays the node name of the device in the <code>/dev</code> directory.	<code>sde</code>
Size	Displays the total amount of space on the device.	11.0GB
Free	Displays the amount of free unpartitioned space on the device.	0KB 199.50GB
Format	Displays the partitioning scheme used to format the device. MSDOS supports devices up to 2 TB in size. GPT supports devices of any size.	MSDOS GPT
Shared	Displays whether the Shareable for Clustering setting is enabled or disabled. The Shareable for Clustering setting should be enabled for devices used in clustered NSS pools and for SBD partitions.	Yes, enabled No, disabled
RAID	Displays whether the device is used as a segment in an NSS software RAID device.	Yes or No

If the `more` option is specified, the information returned for each device is the same as for the `nlvm list device` command.

- 1 Log in to the server as the `root` user, then open a terminal console.
- 2 View a list of active devices by entering

```
nlvm list devices [more] [exclude=<raid|nonraid|shared|nonshared|lvm|nonlvm>]
```

For example:

```
avalon:~/Desktop # nlvm list devices
sda size=11.00GB free=1.30GB format=MSDOS shared=No RAID=No
sdb size=102.00MB free=0KB format=MSDOS shared=Yes RAID=No
sdc size=102.00MB free=0KB format=MSDOS shared=Yes RAID=No
sdd size=512.00MB free=0KB format=None shared=No RAID=No
sde size=512.00MB free=0KB format=None shared=No RAID=No
clus1.sbd size=99.57MB free=0KB format=None shared=Yes RAID=Yes
```

7.3 Viewing Device Details

You can use the `nlvm list device <device_name>` command to view details about a specified device on the system. For command usage information, see “[List Device](#)” in the [OES 11 SP1: NLVM Reference](#).

The command returns the following information:

Device name

Device major:minor
Size of device in MB, GB, or TB
Free space remaining on the device in KB, MB, GB, or TB
Partitioning format – MSDOS, GPT, CSM, LVM (meaning Clustered Linux LVM volume), None
Geometry heads:sectors per track
Shared – (Yes or No) Whether this device is marked as shared
RAID – (Yes or No) Whether this is an NSS software RAID device

For RAID devices, it provides the following information:

Type – 0, 1, or 5 Software RAID type
Segs – Number of segments that this RAID should have
Missing – Segment number (if any) that is missing in the RAID
Stripe – RAID stripe size in bytes (typically kilobytes)
Enbl – (Yes or No) Whether the RAID is enabled on this node
Sync – (1 or 0) Whether the RAID is in sync
% – Percent complete of remirror or restripe

For RAID segments, it provides the following information:

Segment index
Segment (partition) name
Device name of the segment
Segment size
Sync – (1 or 0) If this mirror segment is in sync
% – Percent this mirror segment is remirrored

For a device's partitions, it provides the following information:

Partition name
Partition size
Partition type
Pool name if the partition is the NSS type and it contains a pool

To view device details:

- 1 Log in to the server as the `root` user, then open a terminal console.
- 2 View details about a device by entering

```
nlvm list device <device_name>
```

For example:

```
avalon:~/Desktop # nlvm list device sde
sde (8:64) size=512.00MB(1048576) free=511.98MB(1048544) format=MSDOS h:s=255:6
3 shared=No RAID=No
```

7.4 Initializing a Device

You must initialize a device to set up its device format before you can create volumes on it. You can also initialize a device to wipe its current structure and reconfigure it. Devices that you want to use for Linux POSIX volumes should be in an unshared state whether they are local or SAN devices. For command usage information, see “[Init Device](#)” in the *OES 11 SP1: NLVM Reference*.

IMPORTANT: NLVM and NSSMU do not support using Linux software RAID devices and NSS software RAID devices with Linux POSIX file systems. You can use a hardware RAID device to achieve device fault tolerance for Linux POSIX volumes.

Initializing a device formats it with an MSDOS or a GPT partitioning scheme. MSDOS supports devices up to 2 TB in size. GPT supports devices of any size. The default is MSDOS. If the device size is greater than 2 TB and the partitioning scheme is not specified, the default partitioning scheme of MSDOS applies, and the device size is truncated to 2 TB with the remainder as unusable space.

WARNING: Initializing a device removes all partitions and data from the device. Do not initialize the device that contains the operating system.

Devices that have never been initialized have a format of `None`. Devices that are being used for a Novell Cluster Services SBD (split brain detector) partition also have a format of `None`; however, you should not use the `nlvm init` command to remove an SBD partition. For information about removing an SBD partition, see “[Creating or Deleting Cluster SBD Partitions](#)” in the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide*.

- 1 Log in to the server as the `root` user, then open a terminal console.
- 2 View a list of active devices by entering

```
nlvm list devices
```

- 3 Identify the device that you want to initialize. If a device contains data, be prepared to lose all data on the device when you initialize it.

In the following example, the `sdd` and `sde` devices are uninitialized and display a format of `None`. The `clus1.sbd` device is a mirrored RAID device that is used for the SBD partition of a cluster named `clus1`. The SBD device should not be used.

```
avalon:~/Desktop # nlvm list devices
sda size=11.00GB free=1.30GB format=MSDOS shared=No RAID=No
sdb size=102.00MB free=0KB format=MSDOS shared=Yes RAID=No
sdc size=102.00MB free=0KB format=MSDOS shared=Yes RAID=No
sdd size=512.00MB free=0KB format=None shared=No RAID=No
sde size=512.00MB free=0KB format=None shared=No RAID=No
clus1.sbd size=99.57MB free=0KB format=None shared=Yes RAID=Yes
```

- 4 Initialize the device by entering

```
nlvm [--force] [--no-prompt] init <device_name> [format=<gpt|msdos>] unshared
```

You are automatically prompted to confirm the initialize action. Respond Y (Yes) or N (No). Use the `--no-prompt` NLVM option to suppress the confirmation.

Replace `device_name` with the node name of the device to be initialized, such as `sde`. The device name must be the first option after `init`.

Specify `gpt` or `msdos` as the partitioning scheme to use when formatting the device.

The `unshared` option removes all partitions from a device. If the device was previously set as `shared`, this removes the `shared` setting from the device.

For devices that contain data, specify the `--force` option to force the initialization if the device contains the `root (/)`, `swap`, or `/boot` partition, or if the `init` command cannot delete any pools on the disk.

For example, to initialize a device with the MSDOS partitioning scheme and leave it as unshared, enter

```
nlvm init sde format=msdos unshared
```

- 5 List details about the device to verify that the device is formatted, and the amount of free space has increased.

```
nlvm list device <device_name>
```

For example, enter

```
nlvm list device sde
```

```
avalon:~/Desktop # nlvm list device sde
sde (8:64) size=512.00MB(1048576) free=511.98MB(1048544) format=MSDOS h:s=255:6
3 shared=No RAID=No
```

7.5 Unsharing a Device

The sharing state of devices in NLVM is intended for marking devices as shareable that you plan to use for cluster-enabled NSS pools or for Novell Cluster Services SBD (split brain detector) partitions.

For Linux volumes, you use devices that are not marked as shareable for clustering. When you cluster-enable an LVM volume group for Novell Cluster Services, clustered LVM controls the share state of the device, not the Shareable for Clustering setting. If a device was previously used for an NSS pool or SBD partition and reports a Shared state of *Yes*, it contains a small 4 KB partition to store the state. Before you can create a clustered LVM volume group on the device, you must change the device's *Shared* state from *Yes* to *No*.

You can use the `nlvm unshare` command to change the share state of a device. For command usage information, see [“List Linux Volumes”](#) in the *OES 11 SP1: NLVM Reference*.

- 1 Log in to the server as the `root` user, then open a terminal console.
- 2 View a list of active devices by entering

```
nlvm list devices
```

- 3 Verify the device node name and share state of the device.
- 4 Unshare the device by entering

```
nlvm unshare sdd
```

- 5 View a list of active devices by entering

```
nlvm list devices
```

- 6 Verify that the share state of the device has changed from *Yes* to *No*.

7.6 Viewing a List of Linux Volumes

You can use the `nlvm list linux volumes` command to view information about Linux POSIX and LVM logical volumes that are mounted on the system. For command usage information, see [“Unshare”](#) in the *OES 11 SP1: NLVM Reference*.

- 1 Log in to the server as the `root` user, then open a terminal console.
- 2 View a list of active Linux volumes by entering

```
nlvm list linux volumes
```

For example:

```
avalon:~/Desktop # nlvm list linux volumes
Linux Volumes:
  Name=/ Path=/dev/sda3 mountpoint=/ type=ext3 Mounted
  Name=/boot Path=/dev/sda1 mountpoint=/boot type=ext2 Mounted
  Name=/vol1 Path=/dev/sda5 mountpoint=/mnt/vol1 type=ext3 Mounted
  Name=/vol2 Path=/dev/sda6 mountpoint=/mnt/vol2 type=ext3 Mounted
```

7.7 Creating a Linux POSIX Volume

You can use the `nlvm create linux volume` command to create a native Linux POSIX volume on a device. You must have free unpartitioned space available on a device. For command usage information, see “[Create Linux Volume](#)” in the *OES 11 SP1: NLVM Reference*.

- 1 Log in to the server as the `root` user, then open a terminal console.
- 2 List the devices by entering the following command, then verify that the device that you want to use has been formatted.

```
nlvm list device <device_name>
```

Note the amount of free unpartitioned space available.

- 3 To create the volume, enter

```
nlvm create linux volume
  type=<btrfs|ext2|ext3|reiserfs|xfs>
  device=<device_name|anydisk>
  size=<value[K|M|G|T]|max>
  mp=<full_mount_point_path>
  [mntopt=<fs_mount_options>]
  [ncp]
```

Specify the type of file system to create on the volume. Valid values are `btrfs`, `ext2`, `ext3`, `reiserfs`, and `xfs`.

Replace `device_name` with the node name of the device where you want to create the volume, such as `sda` or `sdd`. You can also specify `anydisk` to use any unshared disk that has sufficient free unpartitioned space to create the volume.

Specify the size of the partition to create for the Linux volume, or specify `max` to use all of the free unpartitioned space on the specified device. All sizes are in bytes and can be specified with one of the following multipliers: K (kilobytes), M (megabytes), G (gigabytes), or T (terabytes). Multipliers are case insensitive and are multiples of 1024. If no multiplier is specified, it is assumed to be G. If `max` is entered, all of the free unpartitioned space on the device is used. The minimum allowed size is 1 megabyte.

Replace `full_mount_point_path` with the full Linux path where you want to mount the device. The last directory is assumed to be the name you want to give the volume. The name must be unique on the server. For example, `/mnt/vol2`.

NOTE: If you enable NCP for the volume, the name of the final folder in the path must comply with the limitations for NCP volume names described in [Section 5.4, “Naming Conventions for NCP Volume Names,”](#) on page 31. If you use lowercase letters for the volume name, they are automatically changed to uppercase for the NCP volume name.

Replace `fs_mount_options` with the mount options to use when mounting the volume. For a list of available options, see the `mount (8)` man page. The default `mntopt` value is `rw`.

Use the `ncp` option to enable the Linux POSIX file system on the volume to be accessed with the NetWare Control Protocol (NCP). It creates an NCP volume of the same name, and creates an NCP share at the root of the volume.

For example, enter the following command to create a 20 GB Linux POSIX volume named `vol2` with the Ext3 file system that is mounted for read and write access at `/mnt/vol1`:

```
nlvm create linux volume type=ext3 device=sda size=20G mp=/mnt/vol1 mntopt=rw
```

If the creation is successful, the response is:

```
Volume mounted at /mnt/vol2
```

- 4 Verify that the volume is mounted by entering

```
nlvm list linux volumes
```

For example:

```
avalon:~/Desktop # nlvm list linux volumes
Linux Volumes:
  Name=/ Path=/dev/sda3 mountpoint=/ type=ext3 Mounted
  Name=/boot Path=/dev/sda1 mountpoint=/boot type=ext2 Mounted
  Name=/vol1 Path=/dev/sda5 mountpoint=/mnt/vol1 type=ext3 Mounted
  Name=/vol2 Path=/dev/sda6 mountpoint=/mnt/vol2 type=ext3 Mounted
```

If NCP is enabled for the volume, an NCP value is also displayed.

7.8 Creating an LVM Logical Volume

You can use the `nlvm create linux volume` command to create a Linux LVM volume group and logical volume, make a file system on the volume, and mount the volume. LVM requires a device that is unpartitioned and unshared. For command usage information, see “[Create Linux Volume](#)” in the *OES 11 SP1: NLVM Reference*.

For conceptual information about LVM, see “[Understanding the Logical Volume Manager](#)” (http://www.suse.com/documentation/sles11/stor_admin/data/sec_yast2_system_lvm_explanation.html) in the *SUSE Linux Enterprise Server 11 SP1 Storage Administration Guide* (http://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html).

- 1 Log in to the server as the root user, then open a terminal console.
- 2 List the devices by entering the following command, then verify that the device that you want to use has been formatted.

```
nlvm list device <device_name>
```

Note the amount of free unpartitioned space available.

- 3 To create the volume, enter

```
nlvm create linux volume
  type=<btrfs|ext2|ext3|reiserfs|xfs>
  device=<devicename>
  size=<value[K|M|G|T]>
  mp=<full_mount_point_path>
  [mntopt=<fs_mount_options>]
  lvm
  name=<lv_name>
  [group=<vg_name>]
  [ncp]
```

Specify the type of file system to create on the volume. Valid values are `ext2`, `ext3`, `reiserfs`, and `xfs`.

Replace *device_name* with the node name of the device where you want to create the volume, such as `sde`.

IMPORTANT: NLVM does not support using Linux software RAID devices and NSS software RAID devices with Linux POSIX file systems. You can use a hardware RAID device to achieve device fault tolerance for Linux POSIX volumes.

Specify the size of the amount of free unpartitioned space on the device. All sizes are in bytes and can be specified with one of the following multipliers: K (kilobytes), M (megabytes), G (gigabytes), or T (terabytes). Multipliers are case insensitive and are multiples of 1024. If no multiplier is specified, it is assumed to be G.

Replace *full_mount_point_path* with the full Linux path where you want to mount the device. The final directory's name can be the same or different than the name that you assign to the logical volume.

Replace *fs_mount_options* with the mount options to use when mounting the volume. For a list of available options, see the `mount (8)` man page. The default `mntopt` value is `rw`.

Specify the `lv` option to create an LVM volume group and logical volume.

Replace *lv_name* with the name you want to give the LVM logical volume.

NOTE: If you enable NCP for the volume, the LVM volume name must comply with the limitations for NCP volume names described in [Section 5.4, "Naming Conventions for NCP Volume Names," on page 31](#). If you use lowercase letters for the volume name, they are automatically changed to uppercase for the NCP volume name.

Replace *vg_name* with the name you want to give the LVM volume group. If the group option is not used, the volume group uses the same name as the logical volume.

Use the `ncp` option to enable the Linux POSIX file system on the volume to be accessed with the NetWare Control Protocol (NCP). It creates an NCP volume of the same name, and creates an NCP share at the root of the volume.

For example, enter the following command to create an LVM volume group named `lvmvg42` and volume named `lvmvol42` with the XFS file system that is mounted for read and write access at `/mnt/vol42`:

```
nlvm create linux volume type=xfs device=sde size=511M mp=/mnt/vol42 mntopt=rw
lvm name=lvmvol42 group=lvmvg42
```

If the creation is successful, the response is:

```
Volume mounted at /mnt/vol42
```

4 Verify that the volume is mounted by entering

```
nlvm list linux volumes
```

For example:

```
avalon:~/Desktop # nlvm list linux volumes
```

```
Linux Volumes:
```

```
Name=/ Path=/dev/sda3 mountpoint=/ type=ext3 Mounted
```

```
Name=/boot Path=/dev/sda1 mountpoint=/boot type=ext2 Mounted
```

```
Name=/vol1 Path=/dev/sda5 mountpoint=/mnt/vol1 type=ext3 Mounted
```

```
Name=/vol2 Path=/dev/sda6 mountpoint=/mnt/vol2 type=ext3 Mounted
```

```
Name=lvmvol42 Path=/dev/lvmvg42/lvmvol42 mountpoint=/mnt/vol42 type=xfs Mounted
```

If NCP is enabled for the volume, an NCP value is also displayed.

- 5 View information about the volume group you created by entering the following Linux LVM command:

```
vgdisplay [vg_name]
```

For example:

```
avalon:~/Desktop # vgdisplay
--- Volume group ---
VG Name                lvmvg42
System ID
Format                 lvm2
Metadata Areas         1
Metadata Sequence No   2
VG Access               read/write
VG Status               resizable
MAX LV                 0
Cur LV                 1
Open LV                 1
Max PV                 0
Cur PV                 1
Act PV                 1
VG Size                 508.00 MB
PE Size                 4.00 MB
Total PE               127
Alloc PE / Size         127 / 508.00 MB
Free PE / Size           0 / 0
VG UUID                 A5ZE4n-4KM7-1lbh-fBrf-sNSV-n6iU-CSovSS
```

7.9 Creating a Shared LVM Logical Volume with Novell Cluster Services

For information about creating a clustered LVM volume group cluster resource with NLVM commands, see [Section 8.3.3, “Creating an LVM Volume Group Cluster Resource with NLVM Commands,”](#) on page 88 in Chapter 8, “Clustering LVM Volume Groups with Novell Cluster Services,” on page 75.

7.10 Mounting Linux Volumes

There is no NLVM command available for mounting Linux volumes. NSSMU provides a *Mount* option on the Linux Volumes page. For information, see [Section 6.12, “Mounting a Linux Volume,”](#) on page 53.

You can also use the native Linux `mount (8)` command to mount Linux volumes that you created with NSSMU or NLVM. See the `mount (8)` man page for information about usage and the mount options for your file system.

IMPORTANT: If an LVM logical volume is clustered, you must use the `cluster online` command to allow the load script to mount the volume. You can also use the *Clusters > Cluster Manager* page in Novell iManager to online a resource.

- 1 Log in to the server as the `root` user, then open a terminal console.
- 2 View a list of Linux volumes by entering the following command, then find the name, mount point, and mount status of the volume you want to mount.

```
nlvm list linux volumes
```

For example, the following Linux POSIX volume's name is `/vol1`. Its mount point is `/mnt/vol1`. It is not mounted.

```
avalon:~/Desktop # nlvm list linux volumes
```

```
Linux Volumes:
```

```
Name=/ Path=/dev/sda3 mountpoint=/ type=ext3 Mounted
```

```
Name=/boot Path=/dev/sda1 mountpoint=/boot type=ext2 Mounted
```

```
Name=/vol1 Path=/dev/sdd1 mountpoint=/mnt/vol1 type=xfs ☐
```

The following LVM logical volume's name is `lvvol1`. Its mount point is `/mnt/lvvol1`. It is not mounted. The LVM volume group is active on the server; otherwise, the volume would not be available in the list.

```
avalon:~/Desktop # nlvm list linux volumes
```

```
Linux Volumes:
```

```
Name=/ Path=/dev/sda3 mountpoint=/ type=ext3 Mounted
```

```
Name=/boot Path=/dev/sda1 mountpoint=/boot type=ext2 Mounted
```

```
Name=lvvol1 Path=/dev/lvvol1/lvvol1 mountpoint=/mnt/lvvol1 type=xfs ☐
```

- 3 Mount the volume by entering

```
mount -t <fstype> <device_name> <full_mount_point_path>
```

Replace *fstype* with the file system type of the volume.

For a Linux POSIX volume, replace *device_name* with the Linux path of the device. For an LVM volume, replace *device_name* with the full device path of the logical volume, such as `/dev/<vg_name>/<lv_name>`.

Replace *full_mount_point_path* with the volume's mount point.

For example, enter

```
mount -t xfs /dev/sdf /mnt/vol1
```

```
mount -t ext3 /dev/lvvol1/lvvol1 /mnt/lvvol1
```

7.11 Dismounting Linux Volumes

There is no NLVM command available for dismounting Linux volumes. NSSMU provides a *Dismount* option on the Linux Volumes page. For information, see [Section 6.13, “Dismounting a Linux Volume,”](#) on page 53.

You can use the native Linux `umount (8)` command to dismount Linux volumes. See the `umount (8)` man page for information about usage and the dismounting options for your file system.

IMPORTANT: If an LVM logical volume is clustered, you must use the Novell Cluster Services `cluster offline` command to allow the unload script to dismount the volume. You can also use the *Clusters > Cluster Manager* page in Novell iManager to offline a resource.

- 1 Log in to the server as the `root` user, then open a terminal console.
- 2 View a list of Linux volumes by entering the following command, then find the name, mount point, and mount status of the volume that you want to dismount.

```
nlvm list linux volumes
```

For example, the following Linux POSIX volume's name is `/vol1`. Its mount point is `/mnt/vol1`. Its mount status is *Mounted*.

```
avalon:~/Desktop # nlvm list linux volumes
Linux Volumes:
  Name=/ Path=/dev/sda3 mountpoint=/ type=ext3 Mounted
  Name=/boot Path=/dev/sda1 mountpoint=/boot type=ext2 Mounted
  Name=/vol1 Path=/dev/sdd1 mountpoint=/mnt/vol1 type=xfs Mounted
```

The following LVM logical volume's name is `lvvol1`. Its mount point is `/mnt/lvvol1`. Its mount status is *Mounted*.

```
avalon:~/Desktop # nlvm list linux volumes
Linux Volumes:
  Name=/ Path=/dev/sda3 mountpoint=/ type=ext3 Mounted
  Name=/boot Path=/dev/sda1 mountpoint=/boot type=ext2 Mounted
  Name=lvvol1 Path=/dev/lvvol1/lvvol1 mountpoint=/mnt/lvvol1 type=xfs Mounted
```

- 3 Dismount the volume by entering the native Linux `umount` command:

```
umount <full_mount_point_path>
```

Replace *full_mount_point_path* with the volume's mount point. For example, enter

```
umount /mnt/vol1
umount /mnt/lvvol1
```

- 4 Verify that the volume's status is no longer *Mounted* by entering

```
nlvm list linux volumes
```

The volume does not display a status of *Mounted*. For example:

```
avalon:~/Desktop # nlvm list linux volumes
Linux Volumes:
  Name=/ Path=/dev/sda3 mountpoint=/ type=ext3 Mounted
  Name=/boot Path=/dev/sda1 mountpoint=/boot type=ext2 Mounted
  Name=/vol1 Path=/dev/sdd1 mountpoint=/mnt/vol1 type=xfs
```

7.12 Rescanning for Storage Objects on Known Devices

You can use the `nlvm rescan` command to scan for storage objects (such as partitions, NSS pools, and NSS software RAIDs) on known devices. It creates and updates Device Mapper objects, and mounts pools as needed. Use the `-m` option if you do not want to automatically mount any unmounted pools. For command usage information, see “[Rescan](#)” in the *OES 11 SP1: NLVM Reference*.

- 1 Log in to the server as the `root` user, then open a terminal console.
- 2 Rescan the system for storage objects by entering

```
nlvm -m rescan
```

The `-m` option causes it to not automatically mount any pools on the system.

7.13 Deleting a Linux POSIX Volume

You can use the `nlvm delete linux volume` command to delete a Linux POSIX volume. The volume must be unmounted. For command usage information, see “[Delete Linux Volume](#)” in the *OES 11 SP1: NLVM Reference*.

WARNING: The delete process removes the partition for the volume, and destroys all of the data on the partition.

- ♦ [Section 7.13.1, “Deleting a Linux POSIX Volume,” on page 72](#)
- ♦ [Section 7.13.2, “Deleting an LVM Volume Group and Logical Volume,” on page 73](#)

7.13.1 Deleting a Linux POSIX Volume

- 1 Log in to the server as the `root` user, then open a terminal console.
- 2 View a list of Linux volumes by entering the following command, then find the name, mount point, and mount status of the volume that you want to delete.

```
nlvm list linux volumes
```

For example, the following Linux POSIX volume’s name is `/vol1`. Its mount point is `/mnt/vol1`. Its mount status is *Mounted*.

```
avalon:~/Desktop # nlvm list linux volumes
Linux Volumes:
  Name=/ Path=/dev/sda3 mountpoint=/ type=ext3 Mounted
  Name=/boot Path=/dev/sda1 mountpoint=/boot type=ext2 Mounted
  Name=/vol1 Path=/dev/sdd1 mountpoint=/mnt/vol1 type=xfs Mounted
```

- 3 Dismount the volume by entering the native Linux `umount` command:

```
umount <full_mount_point_path>
```

Replace `full_mount_point_path` with the volume’s mount point. For example, enter

```
umount /mnt/vol1
```

- 4 Delete the Linux POSIX volume by entering

```
nlvm [--no-prompt] delete linux volume <volume_name>
```


You are automatically prompted to confirm the delete action. Respond Y (Yes) or N (No). Use the `--no-prompt` NLVM option to suppress the confirmation.

For example, enter

```
nlvm delete linux volume /vol1
```

- 5 Verify that the volume no longer appears in the list of Linux volumes by entering

```
nlvm list linux volumes
```

The volume does not appear in the list.

```
avalon:~/Desktop # nlvm list linux volumes
Linux Volumes:
  Name=/ Path=/dev/sda3 mountpoint=/ type=ext3 Mounted
  Name=/boot Path=/dev/sda1 mountpoint=/boot type=ext2 Mounted
```

- 6 Verify that the space used by the volume's partition is now available as free unpartitioned space by entering

```
nlvm list devices
```

For example:

```
avalon:~/Desktop # nlvm list devices
sda size=11.00GB free=1.30GB format=MSDOS shared=No RAID=No
sdb size=102.00MB free=0KB format=MSDOS shared=Yes RAID=No
sdc size=102.00MB free=0KB format=MSDOS shared=Yes RAID=No
sdd size=512.00MB free=511.98MB format=MSDOS shared=No RAID=No
sde size=512.00MB free=511.98MB format=MSDOS shared=No RAID=No
clus1.sbd size=99.57MB free=0KB format=None shared=Yes RAID=Yes
```

7.13.2 Deleting an LVM Volume Group and Logical Volume

Use the procedure in this section to delete an LVM volume group and logical volume. If the LVM volume group and logical volume are clustered with Novell Cluster Services, use the instructions in [Section 8.10, "Deleting a Clustered LVM Volume Group and Logical Volume,"](#) on page 106.

- 1 Log in to the server as the root user, then open a terminal console.
- 2 View a list of Linux volumes by entering the following command, then find the name, mount point, and mount status of the volume that you want to delete.

```
nlvm list linux volumes
```

For example, the following LVM logical volume's name is `lvvol1`. Its mount point is `/mnt/lvvol1`. Its mount status is *Mounted*.

```
avalon:~/Desktop # nlvm list linux volumes
Linux Volumes:
  Name=/ Path=/dev/sda3 mountpoint=/ type=ext3 Mounted
  Name=/boot Path=/dev/sda1 mountpoint=/boot type=ext2 Mounted
  Name=lvvol1 Path=/dev/lvvol1/lvvol1 mountpoint=/mnt/lvvol1 type=xfs Mounted
```

- 3 Dismount the volume by entering the native Linux `umount` command:

```
umount <full_mount_point_path>
```

Replace *full_mount_point_path* with the volume's mount point. For example, enter

```
umount /mnt/lvvol1
```

- 4 Delete the LVM logical volume and its volume group by entering

```
nlvm delete linux volume <lv_name>
```

You are automatically prompted to confirm the delete action. Respond Y (Yes) or N (No). Use the `--no-prompt` NLVM option to suppress the confirmation.

For example, enter

```
nlvm delete linux volume lvvol1
```

- 5 Verify that the volume no longer appears in the list of Linux volumes by entering

```
nlvm list linux volumes
```

The volume does not appear in the list.

```
avalon:~/Desktop # nlvm list linux volumes
Linux Volumes:
  Name=/ Path=/dev/sda3 mountpoint=/ type=ext3 Mounted
  Name=/boot Path=/dev/sda1 mountpoint=/boot type=ext2 Mounted
```

- 6 Verify that the LVM volume group is deleted by entering the native LVM `vgdisplay` command:

```
vgdisplay
```

- 7 Verify that the space used by the volume group's partition is now available as free unpartitioned space by entering

```
nlvm list devices
```

For example:

```
avalon:~/Desktop # nlvm list devices
sda size=11.00GB free=1.30GB format=MSDOS shared=No RAID=No
sdb size=102.00MB free=0KB format=MSDOS shared=Yes RAID=No
sdc size=102.00MB free=0KB format=MSDOS shared=Yes RAID=No
sdd size=512.00MB free=511.98MB format=MSDOS shared=No RAID=No
sde size=512.00MB free=511.98MB format=MSDOS shared=No RAID=No
clus1.sbd size=99.57MB free=0KB format=None shared=Yes RAID=Yes
```

7.14 Additional Information

For information about native LVM commands and YaST 2 tools to create and manage LVM volume groups and logical volumes, see the following resources:

- ♦ “Linux LVM Management Tools” in the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide*
- ♦ “LVM Configuration” (http://www.suse.com/documentation/sles11/stor_admin/data/lvm.html) in the *SUSE Linux Enterprise Server 11 SP1 Storage Administration Guide* (http://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html)

8 Clustering LVM Volume Groups with Novell Cluster Services

After you have installed and configured Novell Cluster Services, you can create shared cluster resources for Linux Logical Volume Manager (LVM) volume groups. You create an LVM logical volume on the volume group, and add a Linux POSIX file system such as Btrfs, Ext2, Ext3, ReiserFS, or XFS.

You can enable NCP for the volume to enable the Linux volume to be accessed with the NetWare Core Protocol. An NCP volume of the same name is automatically created, and the NCP share is created at the root of the volume. It also creates the NCP Virtual Server object for cluster resource in Novell eDirectory.

This section describes how to configure the LVM volume group cluster resource, logical volume, and file system with NSSMU.

- [Section 8.1, “Requirements for Creating LVM Cluster Resources,” on page 75](#)
- [Section 8.2, “Initializing a SAN Device,” on page 79](#)
- [Section 8.3, “Configuring an LVM Volume Group Cluster Resource with NSS Management Tools,” on page 81](#)
- [Section 8.4, “Configuring the LVM Cluster Resource Settings,” on page 91](#)
- [Section 8.5, “Viewing or Modifying the LVM Resource Scripts,” on page 94](#)
- [Section 8.6, “Sample LVM Resource Scripts,” on page 96](#)
- [Section 8.7, “Renaming the Mount Point Path for a Clustered LVM Volume,” on page 101](#)
- [Section 8.8, “Renaming a Clustered LVM Logical Volume,” on page 102](#)
- [Section 8.9, “Disabling Clustering for an LVM Volume,” on page 104](#)
- [Section 8.10, “Deleting a Clustered LVM Volume Group and Logical Volume,” on page 106](#)
- [Section 8.11, “Deleting a Clustered LVM Volume \(Created in NSSMU or NLVM\),” on page 110](#)

8.1 Requirements for Creating LVM Cluster Resources

Your system must meet the requirements in this section in addition to the cluster requirements described in “[Planning for Novell Cluster Services](#)” in the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide*.

- [Section 8.1.1, “Novell Cluster Services,” on page 76](#)
- [Section 8.1.2, “Linux Logical Volume Manager 2 \(LVM2\),” on page 76](#)
- [Section 8.1.3, “Clustered Logical Volume Manager Daemon \(CLVMD\),” on page 76](#)
- [Section 8.1.4, “Resource IP Address,” on page 76](#)
- [Section 8.1.5, “Shared Storage Devices,” on page 76](#)

- [Section 8.1.6, “All Nodes Must Be Present,” on page 77](#)
- [Section 8.1.7, “Working in Mixed Node OES Clusters,” on page 77](#)
- [Section 8.1.8, “NCP File Access with Novell NCP Server,” on page 77](#)
- [Section 8.1.9, “SMB/CIFS File Access with Novell Samba,” on page 78](#)
- [Section 8.1.10, “Linux File Access Protocols,” on page 79](#)

8.1.1 Novell Cluster Services

Novell Cluster Services must be installed, configured, and running when you create and manage the shared LVM volume group and logical volume. The cluster must be active.

8.1.2 Linux Logical Volume Manager 2 (LVM2)

The Linux Logical Volume Manager (LVM) 2 software supports LVM volume groups and logical volumes. LVM2 must be installed and running on each node in the cluster. LVM2 runs automatically on OES 11 and later servers; no separate installation or setup is required.

8.1.3 Clustered Logical Volume Manager Daemon (CLVMD)

The Linux Clustered Volume Manager Daemon (CLVMD, `clvmd`) software allows you to exclusively mount a shared volume group on one node at a time in a cluster. It distributes the LVM metadata updates around a cluster. CLVM must be installed and running on each node in the cluster. CLVMD runs automatically on OES 11 and later servers; no separate installation or setup is required.

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

8.1.4 Resource IP Address

Each cluster resource requires a unique static IP address that is in the same subnet as the IP addresses that are used for the cluster and cluster nodes. The IP address is used to provide access and failover capability for the cluster-enabled volume.

8.1.5 Shared Storage Devices

The shared SAN storage device that you use for an LVM volume group cluster resource must be initialized and have no partitions on it. When the device is used in a cluster resource, LVM uses the entire device for the volume group. Ensure that you size your LUNs accordingly. Use the SAN management tools to assign the LUN to all nodes in the cluster.

IMPORTANT: If you use NSS management tools to manage devices, do not enable the Shareable for Clustering option. Doing so adds a 4 KB partition to the device, which makes it unavailable to LVM.

8.1.6 All Nodes Must Be Present

LVM requires the presence of all the nodes in the cluster to modify the metadata on shared storage. This allows LVM to get the exclusive locks it needs to perform actions on shared storage.

Before you attempt to create or modify LVM volume group cluster resources:

- ♦ All of the nodes must be joined in the cluster and running properly.
- ♦ The `clvmd` daemon must be running on all nodes.

8.1.7 Working in Mixed Node OES Clusters

LVM volume group cluster resources are not supported in mixed-node OES clusters when you upgrade from OES 2 SP3 to OES 11 (or later). Complete the upgrade before you create new cluster resources.

8.1.8 NCP File Access with Novell NCP Server

Novell NCP Server can be used to provide NCP file access to Linux POSIX file systems on OES 11 and later servers. Its NCP volumes feature can be used to provide NCP access to files on an LVM volume group cluster resource. NCP Server must be installed, configured, and running on each node in the cluster.

- ♦ [“Naming Conventions for NCP Volumes” on page 77](#)
- ♦ [“Creating an NCP Volume for a New Clustered LVM Volume” on page 77](#)
- ♦ [“Creating an NCP Volume on an Existing Clustered LVM Volume” on page 78](#)
- ♦ [“Using Antivirus Software with NCP Volumes” on page 78](#)

Naming Conventions for NCP Volumes

NCP volume names can be up to 14 alphanumeric characters, using uppercase letters A through Z and numbers 0 through 9. Underscores (`_`) are allowed.

If you NCP enable a Linux volume as you create it with NSSMU or the `nlvm create linux volume` command, the NCP volume name uses the specified Linux volume name, but all letters are capitalized. NCP treats the Linux volume name as case insensitive. Ensure that the specified Linux volume name does not exceed 14 characters, does not use special characters, and is unique across all nodes in the cluster for both Linux and NCP.

Creating an NCP Volume for a New Clustered LVM Volume

You can configure NCP file access for an LVM volume group cluster resource when you create the resource by using NSSMU or the `nlvm create linux volume` command. With the NCP option enabled, these tools automatically add commands to the resource scripts that mount, dismount, and monitor an NCP volume. The NCP volume is named the same as the LVM logical volume name, and all letters in the name are capitalized. The tools automatically create an NCP Virtual Server object for the volume group cluster resource.

Creating an NCP Volume on an Existing Clustered LVM Volume

You can create an NCP Virtual Server object for the LVM cluster resource to make the resource visible in the Novell eDirectory tree. The virtual server alone does not provide NCP file access.

You can add NCP file access support to an existing LVM cluster resource:

1. Create an NCP Virtual Server object for the LVM cluster resource.
2. Create the NCP volume at the mount point path for the Linux volume. You can also create shares at subdirectories on the volume. This creates an NCP Volume object.
3. Modify the NCP configuration file to comment out the instance for the NCP volume. This allows the LVM cluster resource to control mounts and dismounts for the NCP volume.
4. Modify the LVM cluster resource scripts to add commands that define, mount, dismount, and monitor the NCP volume.
5. Take the resource offline, then bring it online to apply the changes.

For details about setting up NCP volumes on an existing clustered Linux volume, see [“Configuring NCP Volumes with Novell Cluster Services”](#) in the *OES 11 SP1: NCP Server for Linux Administration Guide*.

Using Antivirus Software with NCP Volumes

For information about using antivirus software with NCP volumes, see [“McAfee Antivirus Requires Additional Configuration”](#) in the *OES 11 SP1: Planning and Implementation Guide*.

8.1.9 SMB/CIFS File Access with Novell Samba

Samba is an open source software suite that lets Linux and other non-Windows servers provide file and print services to clients that support the Microsoft SMB (Server Message Block) and CIFS (Common Internet File System) protocols. Novell Samba is the Linux Samba software that has been modified to work with Novell eDirectory.

Novell Samba can be used to provide SMB/CIFS access to files on an LVM volume group cluster resource for Novell eDirectory users who are enabled for Linux User Management (LUM). For information about LUM-enabling your eDirectory users, see the *OES 11 SP1: Novell Linux User Management Administration Guide*.

Novell Samba must be installed and configured on each node in the cluster. The cluster load script starts the service when you online the Samba cluster resource, and the unload script stops the service when you offline it. For information about using the Samba resource template to create a Samba cluster resource that is based on an LVM volume group, see [“Configuring Samba for LVM Volume Groups and Novell Cluster Services”](#) in the *OES 11 SP1: Novell Samba Administration Guide*.

IMPORTANT: If you enable both NCP and Novell Samba file access for users, we recommend that you enable the Cross-Protocol Lock (CPL) parameter for NCP Server. CPL helps prevent potential data corruption when files are accessed by non-NCP file access protocols and by other applications that directly access the files with POSIX APIs. CPL is enabled by default. For information, see [“Configuring Cross-Protocol File Locks for NCP Server”](#) in the *OES 11 SP1: NCP Server for Linux Administration Guide*.

8.1.10 Linux File Access Protocols

You can provide native Linux file access to files on an LVM volume group cluster resource for Novell eDirectory users who are enabled for Linux User Management (LUM). The Linux file access protocols must also be LUM enabled on each node of the cluster. For information about LUM-enabling your eDirectory users and native Linux file access protocols, see the [OES 11 SP1: Novell Linux User Management Administration Guide](#).

8.2 Initializing a SAN Device

Before you begin, you must initialize the SAN device to set up its device format. You can also initialize a device to wipe its current structure and reconfigure it.

WARNING: Initializing a device removes all partitions and data from the device. Do not initialize the device that contains the operating system.

Devices that you want to use for a clustered Linux LVM volume group should contain no partitions and be in an unshared state. When you initialize the SAN device by using NSS management tools, ensure that the Shareable for Clustering option is disabled.

IMPORTANT: NLVM and NSSMU do not support using Linux software RAID devices and NSS software RAID devices with Linux POSIX file systems. You can use a hardware RAID device to achieve device fault tolerance for Linux POSIX volumes.

Initializing a device formats it with an MSDOS or a GPT partitioning scheme. MSDOS supports devices up to 2 TB in size. GPT supports devices of any size. The default is MSDOS. If the device size is greater than 2 TB and the partitioning scheme is not specified, the default partitioning scheme of MSDOS applies, and the device size is truncated to 2 TB with the remainder as unusable space.

Devices that have never been initialized have a format of `None` (NLVM commands) or `Uninitialized` (in NSSMU). Devices that are being used for a Novell Cluster Services SBD (split brain detector) partition also have a format of `None`; however, you should not use the `nlvm init` command to remove an SBD partition. For information about removing an SBD partition, see “[Creating or Deleting Cluster SBD Partitions](#)” in the [OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide](#).

Use the procedure in this section to initialize the SAN device that you want to use for the LVM volume group. Do not mark it as shareable for clustering.

- 1 Ensure that the SAN device is attached to all of the nodes in the cluster.
- 2 Log in to the master node of the cluster as the `root` user, then open a terminal console.
- 3 Launch NSSMU:

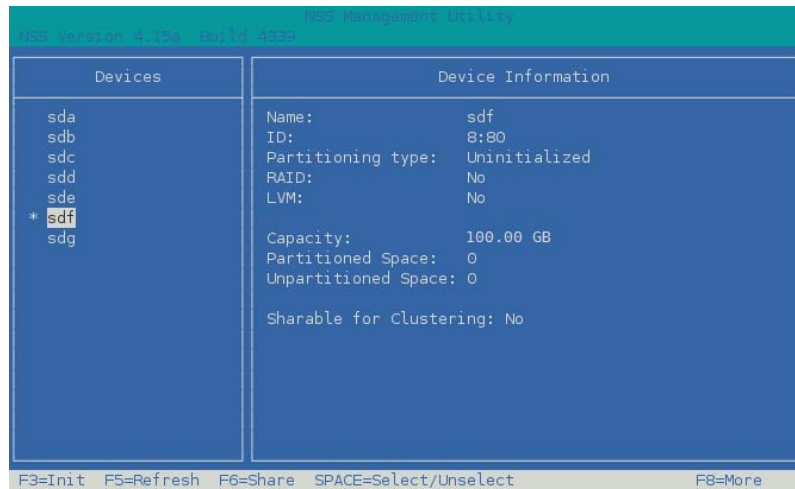
```
nssmu
```

- 4 In the NSSMU main menu, select *Devices*, then press Enter.



- 5 In the *Devices* list, select the SAN device (such as `sdf`), then view information about it.

A device that has never been initialized reports a partition type of *Uninitialized*. If the device contains partitions or data, be prepared to lose all data on the device when it is initialized. The clustered volume group requires the entire device.




- 6 Press F3 to initialize the selected device.

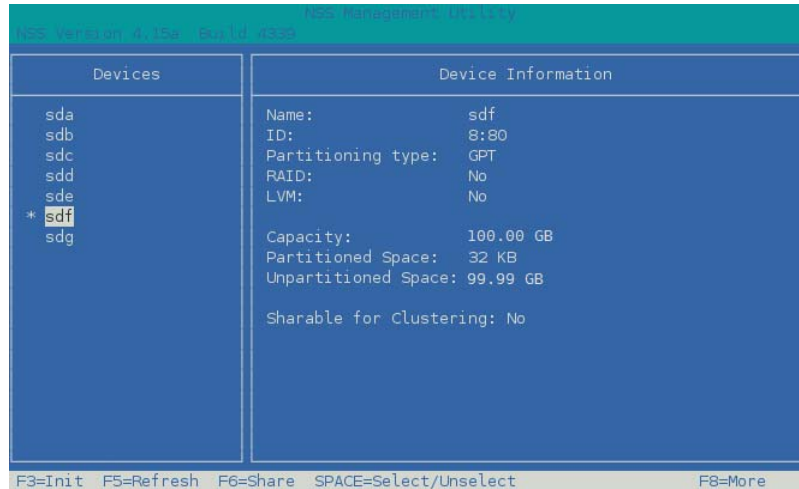
You are prompted to confirm the initialization. If you continue, any existing partitions on the device are deleted, and a new partition table is created.



- 7 Read the advisory message, then do one of the following:
- ♦ Press y (Yes) to confirm that you want to initialize the device, and continue to [Step 8](#).
 - ♦ Press n (No) (or press Esc) to cancel. Return to [Step 5](#) and choose a different device.
- 8 Specify the partitioning scheme to use as DOS or GPT, then press Enter.
- DOS supports devices up to 2 TB in size. GPT supports devices of any size.



- 9 Select the device and view its details to verify that the device is initialized and unshared (that is, *Shareable for Clustering* is set to *No*.) If *Shareable for Clustering* is set to *Yes*, press F6 to unshare the device. This is necessary to make the entire device available to Clustered LVM.



- 10 Press Esc twice to exit NSSMU.

8.3 Configuring an LVM Volume Group Cluster Resource with NSS Management Tools

This section describes how to use the NSSMU utility and NLVM commands to create a clustered LVM volume group and logical volume on a single SAN device that is assigned to all nodes in the cluster. The volume is exclusively mounted on only one node at a time. Clustered LVM manages the locks for the exclusive mount.

After you create the resource, you should view and modify the cluster resource settings and scripts before you bring the resource online. You can also add lines to its load script, unload script, and monitor script to customize the resource for other uses.

If you enable NCP when you create the volume, commands are automatically added to the resource scripts to mount, dismount, and monitor an NCP volume. The NCP volume is named the same as the LVM logical volume name, and all letters in the name are capitalized. An NCP Virtual Server object is created for the resource and commands are added to the scripts to bind or unbind it with the resource IP address.

- ♦ [Section 8.3.1, “Sample Values,” on page 82](#)
- ♦ [Section 8.3.2, “Creating an LVM Volume Group Cluster Resource with NSSMU,” on page 83](#)
- ♦ [Section 8.3.3, “Creating an LVM Volume Group Cluster Resource with NLVM Commands,” on page 88](#)

For information about adding NCP file access to an existing volume, see the following sections in the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide*:

- ♦ [“Creating a Virtual Server Object for an LVM Volume Group Cluster Resource”](#)
- ♦ [“Enabling NCP File Access for a Clustered LVM Volume”](#)

8.3.1 Sample Values

The configuration scenarios use following sample values. Ensure that you replace the sample values with information for your configuration.

Parameter	Sample Value
Device name for the shared SAN device	<code>/dev/sdf</code> The device is initialized and contains no partitions. It is not enabled as shareable for clustering.
Volume group name	<code>vol44</code> By default, NSSMU uses the logical volume name as the LVM volume group name. If you use the NLVM <code>create linux volume</code> command to create the LVM volume group cluster resource, you can specify a different name for the volume group, such as <code>vg44</code> .
Volume name	<code>vol44</code>
NCP volume name	<code>VOL44</code> The NCP volume is created only if you enable NCP as you create the clustered LVM volume by using NSSMU or the NLVM <code>create linux volume</code> command. The mount point path for the NCP volume is the mount point for the LVM volume. This means that the NCP share is created at the root of the mounted LVM volume when the LVM resource is online.
Linux POSIX file system type	<code>ext3</code> Valid values are <code>btrfs</code> (requires the <code>btrfsprogs</code> package), <code>ext2</code> , <code>ext3</code> , <code>reiserfs</code> , and <code>xfs</code> .
Make options for the file system	None (do not specify a value). Press Enter to continue. For a list of the supported file system options for the file system type you are making, see the <code>mkfs(8)</code> man page and the man page for the specific file system: <code>mkfs.btrfs(8)</code> , <code>mkfs.ext2(8)</code> , <code>mkfs.ext3(8)</code> , <code>mkfs.reiserfs(8)</code> , or <code>mkfs.xfs(8)</code> . IMPORTANT: The file system creation fails if you specify a make option that is not supported by the file system type.
Mount options	<code>rw</code> The Read/Write (<code>rw</code>) option is specified by default. For a list of available options that work with the file system type you are using, see the <code>mount(8)</code> man page.
Volume size	100 GB A 100 GB LUN is prepared in the shared storage subsystem. It is attached to the nodes in the cluster. The device must be initialized and contain no partitions. It should not be marked as Shareable for clustering. You are not prompted to enter a volume size. The clustered LVM volume group and logical volume use the entire device. When you select the device, all of the device's free available space is displayed in the <i>Free Size</i> field.

Parameter	Sample Value
Resource IP address	10.10.10.44 This is the IP address of the virtual server for the cluster resource. The address must be unique and in the same subnet as the cluster's IP address. Specify the IP address in IPv4 format.
Mount device	/dev/vol44/vol44 The mount device path format is <code>/dev/<volume_group_name>/<logical_volume_name></code> If you specify a different name for the volume group, such as vg44, by using the NLVM <code>create linux volume</code> command, the mount device path is <code>/dev/vg44/vol44</code> .
Mount point path	/mnt/vol44 You can use any valid Linux path as the mount point. The default mount point location for LVM logical volumes created by NSSMU and NLVM is <code>/usr/novell/<lx_volume_name></code> . NSSMU automatically creates the mount point path if it does not exist on this node. However, you must manually create the mount point on each of the other nodes in the cluster. If the path does not exist on a node when you fail over the resource to it, the resource goes comatose You can alternatively add the following line to the load script before the Linux file system mount command to create the path on a node if it does not exist: <code>ignore_error mkdir -p \$MOUNT_PATH</code>


8.3.2 Creating an LVM Volume Group Cluster Resource with NSSMU

This section describes how to use NSSMU to create and cluster-enable an LVM volume group. NSSMU automatically uses the same script format as the Generic File System template (`Generic_FS_Template`) to create an LVM volume group cluster resource.

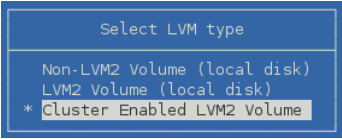
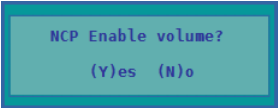

- 1 Log in to the master node of the cluster as the root user, then open a terminal console.
- 2 If you have not already done so, initialize the SAN device that you want to use for the LVM volume group.
For information, see [Section 8.2, "Initializing a SAN Device," on page 79](#).
- 3 Launch NSSMU:

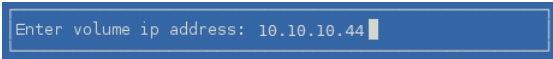
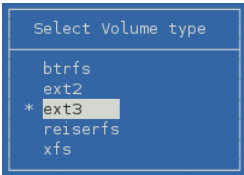
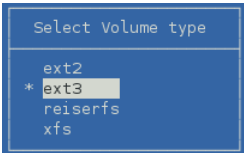

```
nssmu
```

4 In the NSSMU main menu, select *Linux Volumes*, then press Enter.



5 On the Linux Volumes page, press *Insert* to launch the volume creation wizard, then enter the following information as you are prompted for it:

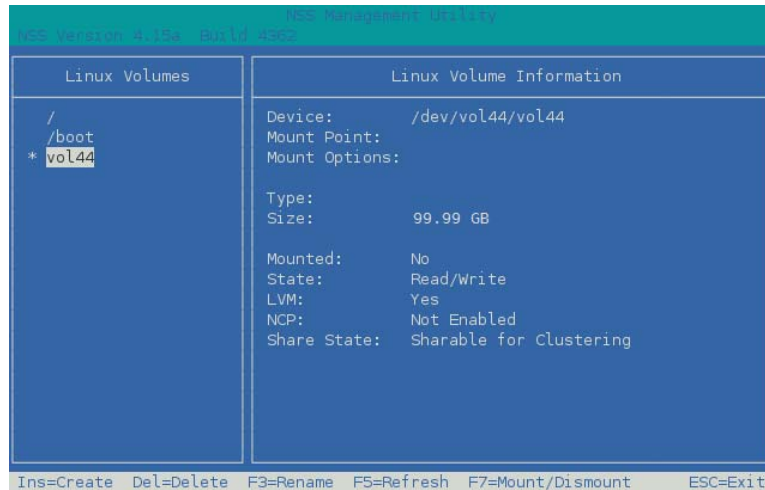
Parameter	Action
Select LVM type	Select <i>Cluster Enabled LVM2 Volume</i> , then press Enter. <div></div>
NCP Enable volume?	Specify whether to enable the LVM logical volume for NCP file access. Press y (yes) to allow NCP file access, or press n (No) to disallow NCP file access at this time. <div></div> <p>If you enable NCP, NSSMU automatically adds commands to the resource scripts that mount, dismount, and monitor an NCP volume. The NCP volume is named the same as the LVM logical volume name, and all letters in the name are capitalized. The tools automatically create an NCP Virtual Server object for the resource.</p> <p>If you do not enable NCP at this time, you can manually create a virtual server and add NCP file access later. For information, see “Creating a Virtual Server Object for an LVM Volume Group Cluster Resource” in the <i>OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide</i>.</p>
Enter volume name	Type the name for the LVM logical volume (such as <code>vol144</code>), then press Enter. <div></div> <p>The specified name is also used for the LVM volume group, which is used in the name of the LVM volume group cluster resource.</p> <p>If you enable NCP for the volume, the specified name is also applied to the NCP volume. An NCP volume name can be up to 14 alphanumeric characters. Underscores are allowed. All letters in the LVM volume name are capitalized for the NCP volume name. For example, if the LVM volume name is <code>vol144</code>, the NCP volume name is <code>VOL44</code>.</p>

Parameter	Action
<i>Enter volume IP address</i>	<p>Type the IP address to use for the LVM cluster resource in IPv4 format (such as 10.10.10.44), then press Enter.</p> 
<i>Select Volume type</i>	<p>Select one of the following the Linux POSIX file systems, then press Enter:</p> <ul style="list-style-type: none"> ♦ btrfs This option is displayed only if the <code>btrfsprogs</code> package is installed. ♦ ext2 ♦ ext3 ♦ reiserfs ♦ xfs <p>If <code>btrfsprogs</code> is installed, five file system options are displayed:</p>  <p>If <code>btrfsprogs</code> is not installed, four file system options are displayed:</p> 
<i>Enter full mount point path</i>	<p>Type the full mount point path for the LVM logical volume (such as <code>/mnt/vol44</code>), then press Enter.</p>  <p>The default path is <code>/usr/novell/<lx_volume_name></code>, such as <code>/usr/novell/vol44</code>.</p> <p>If NCP is enabled, the specified path is also used as the mount point path for the NCP volume.</p>

Parameter	Action																								
<i>Enter any make options</i>	<p>Press Enter to continue without specifying options, or specify the desired make options for the file system type you are making, then press Enter.</p> <div>Enter any make options:</div> <p>For a list of the supported file system options for the file system type you are making, see the <code>mkfs(8)</code> man page and the man page for the specific file system: <code>mkfs.btrfs(8)</code>, <code>mkfs.ext2(8)</code>, <code>mkfs.ext3(8)</code>, <code>mkfs.reiserfs(8)</code>, or <code>mkfs.xfs(8)</code>.</p> <p>IMPORTANT: The file system creation fails if you specify a make option that is not supported by the file system type.</p>																								
<i>Enter any mount options</i>	<p>Press Enter to accept the default read and write options (rw).</p> <div>Enter any mount options: rw</div> <p>The Read/Write (rw) option is specified by default. You can add other mount options, or modify this option. For a list of available options that work with the file system type you are using, see the <code>mount(8)</code> man page.</p>																								
<i>Device</i>	<p>From the list of available devices, select the device that you initialized in Section 8.2, “Initializing a SAN Device,” on page 79, such as <code>sdf</code>, then press Insert or Enter to select the device.</p> <p>The LVM volume group requires the entire device. You cannot specify the amount of space to use. The device's free available space is displayed in the Free Size field.</p> <div><div>LVM Management Utility LVM Version 4.10a Build 4302</div><table><thead><tr><th>Device</th><th>Free Size</th><th>Shared</th><th>Selected</th></tr></thead><tbody><tr><td>1 sdf</td><td>511.95 MB</td><td>No</td><td>99.99 GB</td></tr><tr><td colspan="3">Total</td><td>99.99 GB</td></tr></tbody></table> <table><thead><tr><th>Device</th><th>Free Size</th><th>Shared</th></tr></thead><tbody><tr><td>sde</td><td>511.98 MB</td><td>No</td></tr><tr><td>* sdf</td><td>99.99 GB</td><td>No</td></tr><tr><td>sdg</td><td>715.95 MB</td><td>No</td></tr></tbody></table><div>Ins/ENTER=Select Del=Remove F3=Accept Select list U=Up D=Down ESC/Q=Prev Menu</div></div>	Device	Free Size	Shared	Selected	1 sdf	511.95 MB	No	99.99 GB	Total			99.99 GB	Device	Free Size	Shared	sde	511.98 MB	No	* sdf	99.99 GB	No	sdg	715.95 MB	No
Device	Free Size	Shared	Selected																						
1 sdf	511.95 MB	No	99.99 GB																						
Total			99.99 GB																						
Device	Free Size	Shared																							
sde	511.98 MB	No																							
* sdf	99.99 GB	No																							
sdg	715.95 MB	No																							

- Press F3 to accept the setup you entered for the volume group cluster resource.
- The resource is created and brought online on the node where you created it. The resource is named `<volume_group_name>_reference`. In the example, the name of the volume group is the same as the logical volume, so the reference name is `vol144_reference`.

7 In the *Linux Volumes* list, select the newly created volume and view information about it.



Parameter	Description
Device	Specifies the full device node path for the LVM logical volume. Example: /dev/vol44/vol44
Mount Point	When the resource is online, this specifies the path on the root file system where this volume is mounted. Examples: /mnt/vol44
Mount options	When the resource is online, this specifies the mount options that are applied whenever this volume is automatically mounted after a reboot. Example: rw
Type	When the resource is online, this specifies the file system type. Examples: btrfs, ext2, ext3, reiserfs, xfs
Size	Specifies the amount of space reserved for this volume. Example: 99.58 GB
Mounted	Specifies whether the volume is mounted or unmounted. When the resource is brought online, the load script mounts the logical volume. Value: Yes or No
State	Specifies the availability for the file system. Example: Read/Write
LVM	Specifies whether the specified volume is an LVM logical volume. Value: Yes
NCP	Specifies whether the LVM logical volume is enabled for NCP (yes or no). The NCP state cannot be determined for a clustered LVM volume. You can use the Clusters plug-in for Novell iManager to determine if there are NCP commands in the load script. On the Cluster Options page, select the resource to view its properties, then click the Scripts tab.

Parameter	Description
Share State	Specifies whether the LVM logical volume is cluster enabled for a Novell Cluster Services cluster. Value: Shareable for Clustering

- 8 Press Escape twice to exit NSSMU.
- 9 Continue with [Section 8.4, “Configuring the LVM Cluster Resource Settings,”](#) on page 91.

8.3.3 Creating an LVM Volume Group Cluster Resource with NLVM Commands

This section describes how to use NLVM commands to create and cluster-enable an LVM volume group. NLVM automatically uses the same script format as the Generic File System template (`Generic_FS_Template`) to create an LVM volume group cluster resource. The NLVM command allows you to specify a group name that is different than the volume name.

- 1 Log in to the master node of the cluster as the `root` user, then open a terminal console.
- 2 If you have not already done so, initialize the SAN device that you want to use for the LVM volume group.

Ensure that the SAN device is assigned to all nodes in the cluster. The device must contain no partitions and must not be marked as shareable for clustering.

WARNING: Initializing a device removes all of the data on it.

At the command prompt, enter

```
nlvm [--force] [--no-prompt] init <device_name> format=<gpt|msdos> unshared
```

Replace *device_name* with the device node name of the device to be initialized, such as `sdf`. This must be the first command option to follow `init`.

Specify the partitioning scheme as `gpt` or `msdos`. The default is `msdos`. The MSDOS partitioning scheme supports device sizes that are less than or equal to 2 TB. If the device size is greater than 2 TB and the partitioning scheme is not specified, the default partitioning scheme of MSDOS applies, and the device size is truncated to 2 TB with the remainder as unusable space. Devices of any size can be set to use the GPT partitioning scheme.

Specify the `unshared` option to ensure that the device is not marked as shareable for clustering.

You can specify the `--force` NLVM option to force the initialization.

You are automatically prompted to confirm the initialize action. Respond `y` (Yes) or `n` (No). Use the `--no-prompt` NLVM option to suppress the confirmation.

For example, enter

```
nlvm init sdf format=gpt unshared
```


- 3 Create a clustered LVM volume group and logical volume. At the command prompt, enter the following (all on the same line, of course):

```
nlvm [nlvm_options] create linux volume
type=<btrfs|ext2|ext3|reiserfs|xfs>
device=<device_name>
[mp=</mount_path>]
[mkopt=<option1[,option2,...]>]
[mntopt=<option1[option2[]...]>]
lvm
name=<lvm_volume_name>
[group=<lvm_group_name>]
shared
ip=<IP_address_for_LVM_volgroup_cluster_resource>
[ncp]
```

For details about using this command, see “Create Linux Volume” in the *OES 11 SP1: NLVM Reference*.

Parameters and Options

type=<btrfs|ext2|ext3|reiserfs|xfs>

You must specify the type of file system to use for the volume. Btrfs requires that the btrfsprogs package is installed on all nodes. Supported file systems are btrfs, ext2, ext3, reiserfs, and xfs.

device=<device_name>

You must specify the device to use. Replace *device_name* with the device node name of the device that you want to use for the volume, such as *sdf*. The entire device is used for the LVM volume group. You cannot specify a size to use. The device must already be initialized, contain no partitions, and must not be marked as shareable for clustering.

[mp=</mount_path>]

If a mount point path is not specified, the utility assigns a default mount path of */usr/novell/<volume_name>*. Use the *mp* option to specify a custom mount point. Replace *mount_path* with the full Linux path of the mount point where the volume is to be mounted. The final directory name in the path can be the same or different than the specified volume name. If the path does not currently exist, it is automatically created on that node. You must manually create the path on other nodes in the cluster.

[mkopt=<option1[,option2,...]>]

You can use the *mkopt* option to specify the options to use when running *mkfs*. For a list of available options, see the *mkfs(8)* man page. No default option is specified. For a list of the supported file system options for the file system type you are making, see the *mkfs(8)* man page and the man page for the specific file system: *mkfs.btrfs(8)*, *mkfs.ext2(8)*, *mkfs.ext3(8)*, *mkfs.reiserfs(8)*, or *mkfs.xfs(8)*.

IMPORTANT: The file system creation fails if you specify a make option that is not supported by the file system type.

[mntopt=<option1[option2[]...]>]

You can use the *mntopt* option to specify the options to use when mounting the volume. For a list of available options, see the *mount(8)* man page. The default *mntopt* value is *rw* (Read/Write).

lvm

You must specify the *lvm* option to create an LVM volume group and logical volume. Use this option with the *name* option.

name=<lvm_volume_name>

Replace *lvm_volume_name* with a name for the LVM volume. If you do not specify the group option, this name is also used as the LVM volume group name, which is used in the cluster resource name. For LVM logical volume naming conventions, see “[Linux LVM Volume Group and Logical Volume Names](#)” in the *OES 11 SP1: NLVM Reference*.

NOTE: If you enable NCP for the volume, the LVM volume name must comply with the limitations for NCP volume names described in [Section 5.4, “Naming Conventions for NCP Volume Names,”](#) on page 31. If you use lowercase letters for the volume name, they are automatically changed to uppercase for the NCP volume name.

[group=<lvm_group_name>]

Replace *lvm_volume_group_name* with a name for the LVM volume group. The group name is also used in the cluster resource name. If you do not specify a volume group name, the group is automatically named the same as the LVM volume.

shared

You must specify the `shared` option to create a clustered LVM volume group and logical volume.

ip=<IP_address_for_LVM_volgroup_cluster_resource>

Replace *IP_address_for_LVM_volgroup_cluster_resource* with a static unique IP address to use for the LVM cluster resource. Specify the address in IPv4 format.

ncp

Specify the `ncp` option to NCP enable the LVM logical volume for NCP file access.

- ♦ If you enable NCP, NSSMU automatically adds commands to the resource scripts that mount, dismount, and monitor an NCP volume. The NCP volume is named the same as the LVM logical volume name, and all letters in the name are capitalized. The tools automatically create an NCP Virtual Server object for the resource.
- ♦ If you do not enable NCP at this time, you can manually create a virtual server and add NCP file access later. For information, see “[Creating a Virtual Server Object for an LVM Volume Group Cluster Resource](#)” in the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide*.

Example For example, at the command prompt, enter the following (all on the same line):

```
nlvm create linux volume
type=ext3
device=sdf
mp=/mnt/vol44
mntopt=rw
lvm
name=vol44
group=vg44
shared
ip=10.10.10.44
ncp
```

If the command is successful, the response is

Linux clustered volume vol44 created.

4 Verify that the cluster resource was created and brought online by entering

```
cluster status
```

The resource is named *<lv_name>_resource*. In the following example, `vol44_resource` is in the *Running* state.

```

avalon:~/Desktop # cluster status
Master_IP_Address_Resource      Running      avalon      1
lvmvol134_resource              Running      avalon      1
vol44_resource                  Running      avalon      1

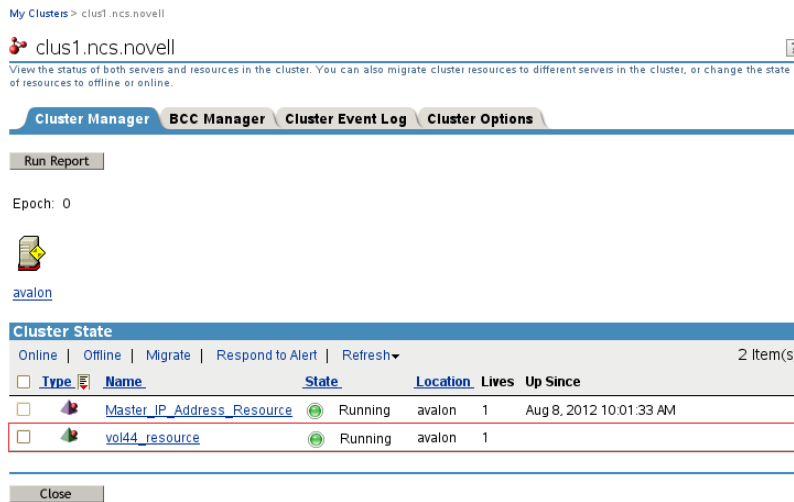
```

- 5 Continue with [Section 8.4, “Configuring the LVM Cluster Resource Settings,”](#) on page 91.

8.4 Configuring the LVM Cluster Resource Settings

Use the procedure in this section to verify that the LVM volume group cluster resource was created and is online. You can customize the resource policies, monitoring, and preferred nodes settings.

- 1 Open Novell iManager in a Web browser, then log in as a cluster administrator user.
- 2 In *Roles and Tasks*, select *Clusters > My Clusters*.
- 3 Select the cluster where you created the volume group cluster resource.
If the cluster you want to manage is not in the list, you can add it. Click *Add*, browse to locate and select the Cluster object of the cluster, then click *OK*.
- 4 In the list of resources, locate the new resource, such as vol44_resource, and notice the state of the resource. It should be online and running.



- 5 Click the resource's name link to open its Properties page.
You can alternatively go to the Cluster Options page, then click the resource's name link, or select the check box next to the resource and click *Details*.

- 6 On the Resource Policies page, view and modify the resource's Policy settings if needed. Click *Apply* if you make changes.

My Clusters > clus1.ncs.novell > vol44_resource

vol44_resource

Policies Monitoring Preferred Nodes Scripts Business Continuity

Set Start, Failover and Failback modes for the new resource. Changes other than business continuity changes made to a resource will not take affect until the resource is reloaded.

Resource Behavior	Failover Mode
<input type="checkbox"/> Resource Follows Master	<input checked="" type="radio"/> Auto
<input type="checkbox"/> Ignore Quorum	<input type="radio"/> Manual

Start Mode	Failback Mode
<input checked="" type="radio"/> Auto	<input type="radio"/> Auto
<input type="radio"/> Manual	<input checked="" type="radio"/> Disable
	<input type="radio"/> Manual

OK Cancel Apply

- 6a (Optional) Select the *Resource Follows Master* check box if you want to ensure that the resource runs only on the master node in the cluster.

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

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

The quorum default values were set when you installed Novell Cluster Services. You can change the quorum default values by accessing the properties page for the Cluster object.

Selecting this box ensures that the resource is launched immediately on any server in the Assigned Nodes list as soon as a server in the list is brought online.

- 6c By default, the Generic File System resource template sets the Start mode and Failover mode to *Auto* and the Failback Mode to *Disable*. You can change the default settings as needed.

- ♦ **Start Mode:** If the Start mode is set to *Auto*, the resource automatically loads on a designated server when the cluster is first brought up. If the Start mode is set to *Manual*, you can manually start the resource on a specific server when you want, instead of having it automatically start when servers in the cluster are brought up.
- ♦ **Failover Mode:** If the Failover mode is set to *Auto*, the resource automatically moves to the next server in the Assigned Nodes list if there is a hardware or software failure. If the Failover mode is set to *Manual*, you can intervene after a failure occurs and before the resource is started on another node.
- ♦ **Failback Mode:** If the Failback mode is set to *Disable*, the resource continues running on the node it has failed to. If the Failback mode is set to *Auto*, the resource automatically moves back to its preferred node when the preferred node is brought back online. Set the Failback mode to *Manual* to prevent the resource from moving back to its preferred node when that node is brought back online, until you are ready to allow it to happen.

7 Enable and configure monitoring for the resource, then click *Apply*.

My Clusters > clus1.ncs.novell > vol44_resource

vol44_resource

Policies Monitoring Preferred Nodes Scripts Business Continuity

To monitor the health of this resource, enable Resource monitoring. You can set the interval to poll the resource's health, and set an action if it fails to successfully load on the maximum number of local restarts. Changes other than business continuity changes made to a resource will not take affect until the resource is reloaded.

☒ Enable Resource Monitoring

Polling Interval: 1 Minutes

Failure Rate

Maximum Local Failures: 3

Time Interval: 10 Minutes
(Example: 3 Failures in 5 minutes)

Failure Action

If the Failure rate settings are reached, perform the following action:

☒ Set Resource as Comatose

☐ Migrate the Resource based on the Preferred Nodes List

☐ Reboot the Hosting Node without Syncing or Unmounting Disks

OK Cancel Apply

7a In the Properties page, select the *Monitoring* tab.

7b Select the *Enable Resource Monitoring* check box to enable resource monitoring for the selected resource.

Resource monitoring is disabled by default.

7c Specify the *Polling Interval* to control how often you want the resource monitor script for this resource to run.

You can specify the value in minutes or seconds.


7d Specify the number of failures (*Maximum Local Failures*) for the specified amount of time (*Time Interval*).

7e Specify the *Failover Action* by indicating whether you want the resource to be set to a comatose state, to migrate to another server, or to reboot the hosting node (without synchronizing or unmounting the disks) if a failure action initiates. The reboot option is normally used only for a mission-critical cluster resource that must remain available. .

- 8 Click the *Preferred Nodes* tab, assign preferred nodes for the resource by moving them from the *Unassigned* list to the *Assigned* list, then click *Apply*.

When you configure a volume group cluster resource with NSSMU or with NLVM commands, the node where you create it is automatically assigned as a preferred node for the resource.

When you bring a resource online, it is automatically loaded on the most preferred node in the list. If the node is not available, the other nodes are tried in the order that they appear in the list. You can modify the order of the nodes by clicking the Edit (pen) icon to open the list in a text editor. In the editor, click *OK* to close the editor, then click *Apply* to save your changes.



- 9 At the bottom of the page, click *OK* to close the Properties page and save your changes.
The changes do not take effect until the resource is taken offline and brought online again.
- 10 Continue with [Section 8.5, “Viewing or Modifying the LVM Resource Scripts,” on page 94](#).

8.5 Viewing or Modifying the LVM Resource Scripts

You can customize the scripts by adding lines for other products that use a shared LVM volume group resource. Compare the generic script with the templates for those products to identify what lines need to be added or modified.

- 1 In iManager, select *Clusters > My Clusters*.
- 2 Select the cluster where you created the volume group cluster resource.
- 3 Select the cluster resource’s name link to open the Properties page, then click the *Scripts* tab.
The *Scripts* tab automatically opens to the load script.
- 4 On the *Load Script* page, view or modify the load script. Click *Apply* if you make changes.
Ensure that all letters in the NCP volume name are capitalized letters A to Z. See the [“Sample LVM Resource Load Scripts Created by NSS Management Tools” on page 97](#).

Policies Monitoring Preferred Nodes Scripts Business Continuity

Load Script | Unload Script | Monitor Script

View or edit the load script for this cluster resource. Changes other than business continuity changes made to a resource will not take affect until the resource is reloaded.

Script:

```
MOUNT_FS=ext3
# define the volume group name
VOL_GROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/vol44/vol44
# define the mount point
MOUNT_POINT=/usr/novell/vol44
# define NCP server name
NCP_SERVER=clus1-vol44-SERVER
# define NCP volume name
NCP_VOLUME=VOL44

# activate the volume group
exit_on_error vgchange -a ey $VOL_GROUP_NAME
```

Timeout: 6 Minutes

OK Cancel Apply

- 5 Click the *Unload Script* link to view or modify the unload script. Click *Apply* if you make changes.

Ensure that all letters in the NCP volume name are capitalized letters A to Z. See the [“Sample LVM Resource Unload Scripts Created by NSS Management Tools”](#) on page 98.

Policies Monitoring Preferred Nodes Scripts Business Continuity

Load Script | Unload Script | Monitor Script

View or edit the unload script for this cluster resource. Changes other than business continuity changes made to a resource will not take affect until the resource is reloaded.

Script:

```
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOL_GROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/vol44/vol44
# define the mount point
MOUNT_POINT=/usr/novell/vol44
# define NCP server name
NCP_SERVER=clus1-vol44-SERVER
# define NCP volume name
NCP_VOLUME=VOL44

ignore_error ncpcon unbind --ncpservname=$NCP_SERVER --ipaddress=$RE
```

Timeout: 6 Minutes

OK Cancel Apply

- 6 Click the *Monitor Script* link to view or modify the monitor script. Click *Apply* if you make changes.

Ensure that all letters in the NCP volume name are capitalized letters A to Z. See the [“Sample LVM Resource Monitor Scripts Created by NSS Management Tools”](#) on page 99.

My Clusters > clust1.ncs.novell > vol44_resource

vol44_resource

Policies Monitoring Preferred Nodes **Scripts** Business Continuity

Load Script | Unload Script | Monitor Script

View or edit the monitor script for this cluster resource. Changes other than business continuity changes made to a resource will not take affect until the resource is reloaded.

Script

```
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOL_GROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/vol44/vol44
# define the mount point
MOUNT_POINT=/usr/novell/vol44
# define NCP server name
NCP_SERVER=clust1-vol44-SERVER
# define NCP volume name
NCP_VOLUME=VOL44

# check the IP address
ignore_error status_secondary_ipaddress $RESOURCE_IP
```

Timeout: 6 Minutes

OK Cancel Apply

7 At the bottom of the page, click **OK** to close the Properties page and save your changes.

The changes do not take effect until the resource is taken offline and brought online again.

8 If you modified the settings on any of the pages, you must take the resource offline and bring it online in order for the changes to take effect.

8a In Roles and Tasks, select *Clusters > My Clusters*.

8b On the Cluster Manager page, select the check box next to the resource, then click *Offline*.
Wait for the status to report that it is offline, then continue.

8c Select the check box next to the resource, then click *Online*.

8d Verify that the resource comes online and reports a *Running* state.

If the resource goes into a Comatose state, it is probably because you made a mistake in the lines you added or modified in the scripts. Take the resource offline, then go back to correct the scripts, and try to bring it online again.

8.6 Sample LVM Resource Scripts

The sample LVM resource scripts in this section are automatically generated for an LVM volume group cluster resource when you use NLVM or NSSMU to create a clustered LVM2 volume, as described in [Section 8.3, “Configuring an LVM Volume Group Cluster Resource with NSS Management Tools,”](#) on page 81. See [Section 8.3.1, “Sample Values,”](#) on page 82 for information about the sample values used in these scripts. Ensure that you replace sample values with those for your own system.

- ♦ [Section 8.6.1, “Sample LVM Resource Load Scripts Created by NSS Management Tools,”](#) on page 97
- ♦ [Section 8.6.2, “Sample LVM Resource Unload Scripts Created by NSS Management Tools,”](#) on page 98
- ♦ [Section 8.6.3, “Sample LVM Resource Monitor Scripts Created by NSS Management Tools,”](#) on page 99

8.6.1 Sample LVM Resource Load Scripts Created by NSS Management Tools

Compare the load scripts in this section to identify the lines that are added when you enable the LVM logical volume for NCP file access:

- ♦ [“Without NCP File Access” on page 97](#)
- ♦ [“With NCP File Access” on page 97](#)

Without NCP File Access

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name (nssmu uses volume name for group name)
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# activate the volume group
exit_on_error vgchange -a ey $VOLGROUP_NAME

# create the mount point if it does not exist on the node
ignore_error mkdir -p $MOUNT_POINT

# mount the file system
exit_on_error mount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

exit 0
```

With NCP File Access

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name (nssmu uses volume name for group name)
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# define NCP server name
NCP_SERVER=clus1-vol44-SERVER
# define NCP volume name
NCP_VOLUME=VOL44

# activate the volume group
exit_on_error vgchange -a ey $VOLGROUP_NAME

# create the mount point if it does not exist on the node
ignore_error mkdir -p $MOUNT_POINT
```

```
# mount the file system
exit_on_error mount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

# mount the NCP volume
exit_on_error ncpcon mount $NCP_VOLUME=251,PATH=$MOUNT_POINT

# bind the NCP virtual server to the resource IP address
exit_on_error ncpcon bind --ncpsservername=$NCP_SERVER --ipaddress=$RESOURCE_IP

exit 0
```

8.6.2 Sample LVM Resource Unload Scripts Created by NSS Management Tools

Compare the unload scripts in this section to identify the lines that are added when you enable the LVM logical volume for NCP file access:

- ♦ [“Without NCP File Access” on page 98](#)
- ♦ [“With NCP File Access” on page 99](#)

Without NCP File Access

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name (nssmu uses volume name for group name)
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# unmount the volume
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error umount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# deactivate the volume group
exit_on_error vgchange -a n $VOLGROUP_NAME

exit 0
```

With NCP File Access

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name (nssmu uses volume name for group name)
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# define NCP server name
NCP_SERVER=clus1-vol44-SERVER
# define NCP volume name
NCP_VOLUME=VOL44

# unbind the NCP virtual server from the resource IP address
ignore_error ncpcon unbind --ncpservername=$NCP_SERVER --ipaddress=$RESOURCE_IP

# dismount the NCP volume
ignore_error ncpcon dismount $NCP_VOLUME

# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

# unmount the volume
sleep 10 # if not using SMS for backup, please comment out this line
exit_on_error umount_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# deactivate the volume group
exit_on_error vgchange -a n $VOLGROUP_NAME

exit 0
```

8.6.3 Sample LVM Resource Monitor Scripts Created by NSS Management Tools

Compare the monitor scripts in this section to identify the lines that are added when you enable the LVM logical volume for NCP file access:

- ♦ [“Without NCP File Access” on page 100](#)
- ♦ [“With NCP File Access” on page 100](#)

Without NCP File Access

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# check the logical volume
exit_on_error status_lv $MOUNT_DEV

# test the file system
exit_on_error status_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# status the IP address
exit_on_error status_secondary_ipaddress $RESOURCE_IP

exit 0
```

With NCP File Access

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name
VOLGROUP_NAME=vol44
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/vol44
# define the mount point
MOUNT_POINT=/mnt/vol44

# define NCP server name
NCP_SERVER=clus1-vol44-SERVER
# define NCP volume name
NCP_VOLUME=VOL44

# check the LVM logical volume
exit_on_error status_lv $MOUNT_DEV

# check the NCP volume
exit_on_error ncpcon volume $NCP_VOLUME

# check the file system
exit_on_error status_fs $MOUNT_DEV $MOUNT_POINT $MOUNT_FS

# status the IP address
exit_on_error status_secondary_ipaddress $RESOURCE_IP

exit 0
```

8.7 Renaming the Mount Point Path for a Clustered LVM Volume

For a clustered LVM volume, you can rename the mount point path by modifying the mount point variable in the cluster load, unload, and monitor scripts.

- 1 Open Novell iManager in a Web browser, then log in as a cluster administrator user.
- 2 In *Roles and Tasks*, select *Clusters > My Clusters*, then select the cluster that you want to manage.
If the cluster is not in your personalized list, you can add it. Click *Add*, browse to locate and select the cluster, then click *OK*.
- 3 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.
- 4 Modify the mount point path value in the load, unload, and monitor scripts for the LVM volume group cluster resource:

- 4a 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.

- 4b Modify the load script:

- 4b1 In the load script, type the new value for the mount point in the MOUNT_POINT variable:

```
MOUNT_POINT=/media/ext3/vol44
```

- 4b2 Ensure that the following command is added above the mount command line in the load script in order to create the path on nodes if it does not exist.

```
# create the mount point if it does not exist on the node ignore_error  
mkdir -p $MOUNT_POINT
```

You can alternatively make the new path by using the `mkdir` command in a terminal console on each node. If the master node is not the most preferred node, ensure that you make the path before you bring the resource online.

- 4b3 Click *Apply*.

- 4c Click the *Unload Script* link, type the new value for the mount point in the MOUNT_POINT variable, then click *Apply*.

```
MOUNT_POINT=/media/ext3/vol44
```

- 4d Click the *Monitor Script* link, type the new value for the mount point in the MOUNT_POINT variable, then click *Apply*.

```
MOUNT_POINT=/media/ext3/vol44
```

- 4e At the bottom of the page, click *OK* to close the Properties page and save your changes.

The changes do not take effect until the resource is brought online.

- 5 Bring the resource online in order for the script changes to take effect.

- 5a In *Roles and Tasks*, select *Clusters > Cluster Manager*.


- 5b Select the check box next to the resource, then click *Online*.

- 5c Verify that the resource comes online and reports a *Running* state.

If the resource goes into a Comatose state, it is probably because you made a mistake in the lines you added or modified in the scripts. Take the resource offline, go back and correct the scripts, then try to bring it online again.

- 6 In NSSMU, verify that the new mount point is used when the clustered LVM volume resource is brought online:
- 6a Log in as the root user to the node that is hosting the resource, then start NSSMU by entering:


```
nssmu
```
 - 6b From the NSSMU main menu, select *Linux Volumes*, then press Enter.
 - 6c In the *Linux Volumes* list, select the clustered LVM volume.
 - 6d View the volume details to verify that the mount point has changed.
For example, the mount point is `/media/ext3/vol144`.



- 6e Press Esc twice to exit NSSMU.

8.8 Renaming a Clustered LVM Logical Volume

For a clustered LVM logical volume, there is no easy tool for renaming the volume. You can rename the LVM logical volume by using the `lvrename` command, then modify the volume name in the cluster load, unload, and monitor scripts. The command does not change the volume group name. The LVM cluster resource must be offline while you rename the logical volume and modify the scripts.

Using the `lvrename` command to rename the LVM logical volume does not modify other related parameters that are used by the LVM cluster resource. You can optionally also modify the following:

- ♦ Mount point path
- ♦ Resource name
- ♦ If NCP is enabled for the LVM logical volume:
 - ♦ NCP virtual server name
 - ♦ NCP volume name
 - ♦ Volume object for the NCP volume on the virtual server

Take the resource offline, modify the related parameters, modify the resource scripts with the new values, then bring the resource online.

- 1 Log in to the server as the root user, then open a terminal console.
- 2 Bring the LVM cluster resource offline.

```
cluster offline <lvm_resource_name>
```

For example, to offline the vol44_resource, enter

```
cluster offline vol44_resource
```

Wait until the resource is offline before continuing.

- 3 At the command prompt, enter

```
lvrename </dev/vg_name/old_lv_name> </dev/vg_name/new_lv_name>
```

Replace *vg_name* with the name of the volume group. If you created the LVM logical volume in NSSMU, the volume group name is the same as the logical volume name. If you created the LVM logical volume with NLVM or with LVM commands, you might have used a different name for the volume group.

Replace *old_lv_name* with the current name of the LVM logical volume.

Replace *new_lv_name* with the new name of the LVM logical volume.

For example, to change the name of the logical volume on volume group *vghome* from *lv_users* to *lv_home*, enter

```
lvrename /dev/vghome/lv_users /dev/vghome/lv_home
```

- 4 Modify the load, unload, and monitor scripts for the LVM cluster resource to use the new LVM logical volume name in the value for the MOUNT_DEV parameter.
 - 4a Open Novell iManager in a Web browser, then log in as a cluster administrator user.
 - 4b In *Roles and Tasks*, select *Clusters > My Clusters*, then select the cluster where you created the LVM cluster resource.
 - 4c 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.
 - 4d Modify the MOUNT_DEV value in the load, unload, and monitor scripts for the LVM volume group cluster resource:
 - 4d1 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.
 - 4d2 Modify the load script:
In the load script, type the new value for the logical volume in the MOUNT_DEV variable, then click *Apply*.

```
# define the device  
MOUNT_DEV=/dev/$VOLGROUP_NAME/lv_home
```
 - 4d3 Click the *Unload Script* link, type the new value for the logical volume in the MOUNT_DEV variable, then click *Apply*.

```
# define the device  
MOUNT_DEV=/dev/$VOLGROUP_NAME/lv_home
```
 - 4d4 Click the *Monitor Script* link, type the new value for the logical volume in the MOUNT_DEV variable, then click *Apply*.

```
# define the device
MOUNT_DEV=/dev/$VOLGROUP_NAME/lv_home
```

4d5 At the bottom of the page, click **OK** to close the Properties page and save your changes.

The changes do not take effect until the resource is brought online.

- 5** (Optional) Rename the mount point path that you use for the logical volume. The mount point should still be unmounted when you perform this task.

For information, see [“Renaming the Mount Point Path for a Clustered LVM Volume”](#) in the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide*.

- 6** (Optional) Rename the resource.

Ensure that the resource is offline, then enter

```
cluster rename <old_resource_name> <new_resource_name>
```

- 7** (Optional) If you created an NCP virtual server for the resource, rename the NCP virtual server name.

Use the *Directory Administration > Delete Object* task in Novell iManager to delete the old NCP virtual server (the NCS:NCP Server object) for the resource, then use `/opt/novell/ncs/bin/ncs_ncpserv.py` to create a new virtual server for the resource. For information, see [“Creating a Virtual Server Object for an LVM Volume Group Cluster Resource”](#) in the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide*. Ensure that you update the load, unload, and monitor scripts.

- 8** (Optional) If NCP file access is enabled for the LVM logical volume, rename the NCP volume and its Volume object.

8a Modify the NCP volume name for the NCP_VOLUME parameter in the load, unload, and monitor scripts for the LVM cluster resource. For information, see [“Adding NCP Volume Commands to the Resource Scripts”](#) in the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide*. Bring the resource online, but do not allow users to access the volume until the Volume object is re-created.

8b Use the *Directory Administration > Delete Object* task in Novell iManager to delete the old Volume object, then create a new Volume object that uses the newly created NCP virtual server name and new NCP volume name. For information, see [“Creating a Shared NCP Volume Object”](#) in the *OES 11 SP1: Novell Cluster Services 2.1 for Linux Administration Guide*.

- 9** Bring the resource online. At the command prompt, enter

```
cluster online <resource_name>
```

8.9 Disabling Clustering for an LVM Volume

Use the procedure in this section if you want to disable clustering for an LVM volume. Afterward, you can mount and dismount the volume only as a local LVM volume.

- 1** Delete the LVM volume group cluster resource:

1a Open Novell iManager in a Web browser, then log in as a cluster administrator user.

1b In *Roles and Tasks*, select *Clusters > My Clusters*, then select the cluster that you want to manage.

If the cluster is not in your personalized list, you can add it. Click *Add*, browse to locate and select the cluster, then click *OK*.

1c On the Cluster Manager page, select the check box next to the LVM volume group cluster resource (such as `vol144_resource`), then click *Offline*. Wait until the resource is offline to continue.

- 1d Select *Clusters > Cluster Options*.
- 1e On the Cluster Options page, select the check box next to the LVM volume group cluster resource, click *Delete*, then click *OK* to confirm.
- 1f Exit iManager.
- 2 Log in as the root user on the server where the resource was online, then open a terminal console.
- 3 Remove the shareable state from the clustered LVM volume group by entering


```
vgchange -c n <vg_name>
```

 For example:


```
vgchange -c n vol44
```
- 4 Activate the volume group on the current node by entering


```
vgchange -a ey <vg_name>
```

 For example:

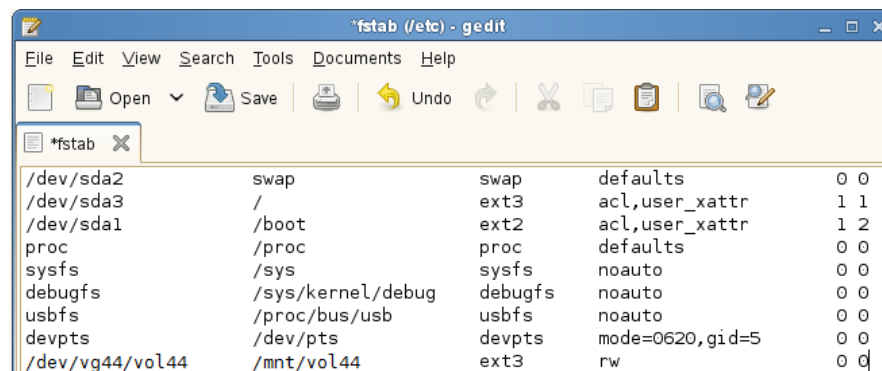

```
vgchange -a ey vol44
```
- 5 Mount the LVM logical volume by entering


```
mount -t <fstype> <device_name> <full_mount_point_path>
```

 Replace *fstype* with the file system type of the volume.
 Replace *device_name* with the full device path of the logical volume, such as `/dev/<vg_name>/<lv_name>`.
 Replace *full_mount_point_path* with the volume's mount point.
 For example, enter

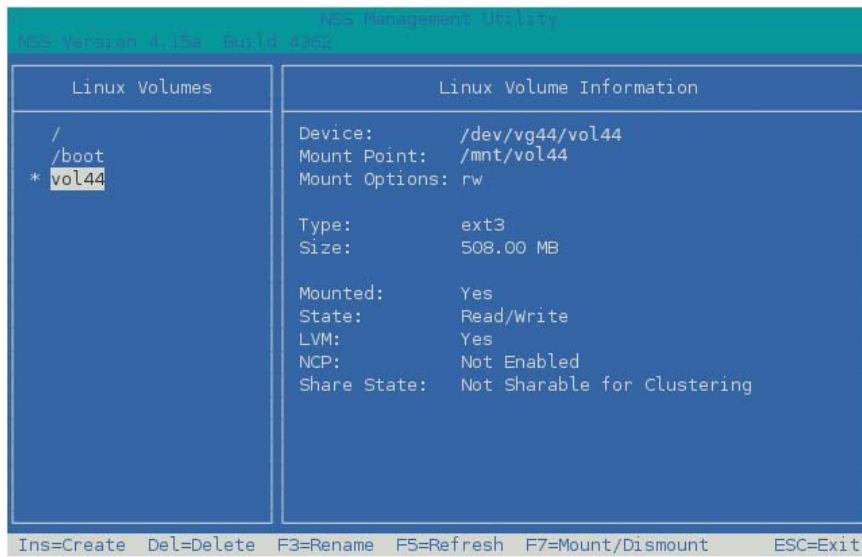

```
mount -t ext3 /dev/vol44/vol44 /media/ext3/vol44
```
- 6 In a text editor, modify the `/etc/fstab` file to specify the mount point information and file system type.
 This entry allows the volume to be mounted automatically on reboot. It also provides the automatic mount information for NSSMU to use to mount and dismount the volume locally.
 For example, complete the line for the volume's device path, such as:

```
/dev/vol44/vol44    /media/ext3/vol44    ext3    rw    0 0
```



- 7 In NSSMU, go to the Linux Volumes page, then select the volume to view its details.

The volume is no longer cluster enabled. You can mount and dismount the volume only on the current node.



8.10 Deleting a Clustered LVM Volume Group and Logical Volume

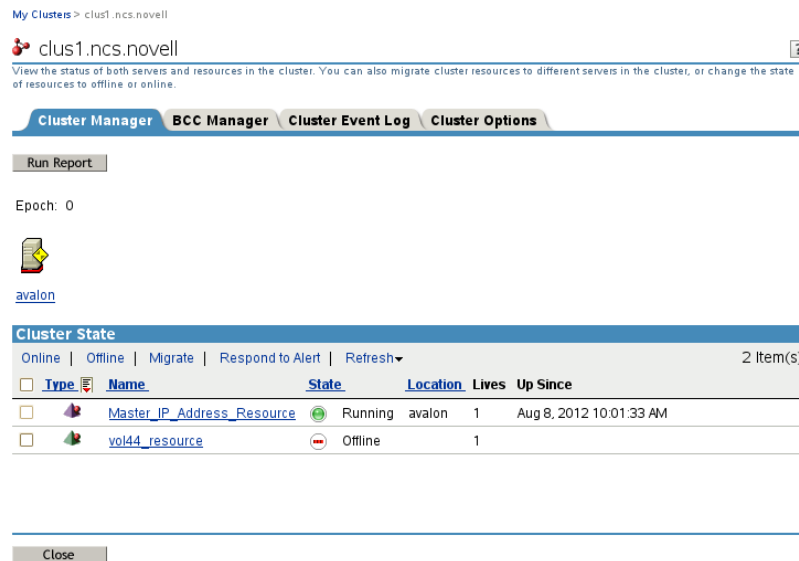
Before you delete a clustered LVM volume group, you must take the volume group cluster resource offline and delete the cluster resource object and related objects in Novell eDirectory.

All resource configuration must happen from the master node. On the Cluster Options page for Novell iManager, you are automatically connected to the Cluster object, which is associated with the master node.

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

3 Take the volume group cluster resource offline:

- 3a** On the Cluster Manager page, select the check box next to the volume group cluster resource, then click *Offline*. Wait until the resource is offline to continue.



My Clusters > clus1.ncs.novell

clus1.ncs.novell

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

Cluster Manager BCC Manager Cluster Event Log Cluster Options

Run Report

Epoch: 0

avalon

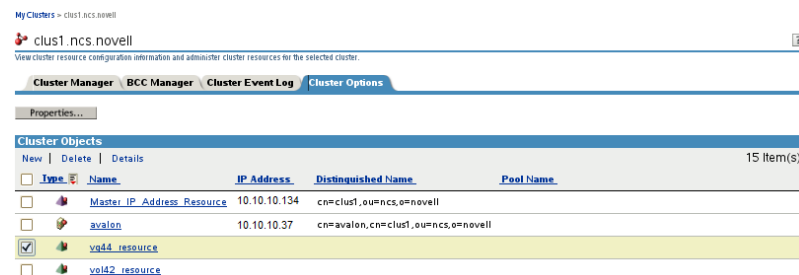
Cluster State						
Online Offline Migrate Respond to Alert Refresh 2 Item(s)						
<input type="checkbox"/>	Type	Name	State	Location	Lives	Up Since
<input type="checkbox"/>	Master IP Address Resource	Master_IP_Address_Resource	Running	avalon	1	Aug 8, 2012 10:01:33 AM
<input type="checkbox"/>	Volume Group Resource	vol44_resource	Offline		1	

Close

The unload script unbinds the NCP virtual server from the resource IP address, dismounts the NCP volume, dismounts the Linux LVM logical volume, and deactivates the Linux LVM volume group.

4 Delete the volume group cluster resource:

- 4a** Click the *Cluster Options* tab.
- 4b** Select the check box next to the volume group cluster resource, then click *Delete*.



My Clusters > clus1.ncs.novell

clus1.ncs.novell

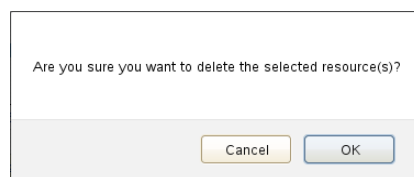
View cluster resource configuration information and administer cluster resources for the selected cluster.

Cluster Manager BCC Manager Cluster Event Log Cluster Options

Properties...

Cluster Objects				
New Delete Details 15 Item(s)				
<input type="checkbox"/>	Type	Name	IP Address	Distinguished Name
<input type="checkbox"/>	Master IP Address Resource	Master_IP_Address_Resource	10.10.10.134	cn=clus1,ou=ncs,o=novell
<input type="checkbox"/>	avalon	avalon	10.10.10.37	cn=avalon,cn=clus1,ou=ncs,o=novell
<input checked="" type="checkbox"/>	Volume Group Resource	vol44_resource		
<input type="checkbox"/>	Volume Group Resource	vol42_resource		

- 4c** When you are prompted to confirm the deletion, click *OK* to continue, or click *Cancel* to abort the deletion.



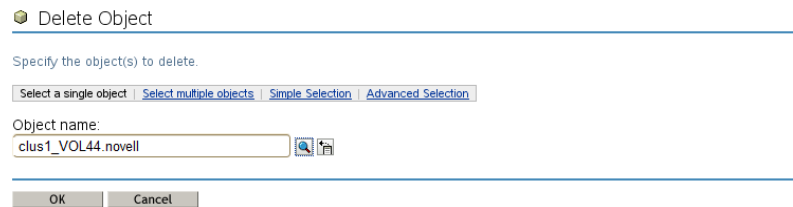
Are you sure you want to delete the selected resource(s)?

Cancel OK

5 If the Linux volume was enabled for NCP file access, delete the NCP volume object.

5a In iManager, select *Directory Administration > Delete Object*.

5b On the Delete Object page, browse to locate and select the resource's NCP Volume object, such as `clus1_VOL44.novell`, then click *OK*.

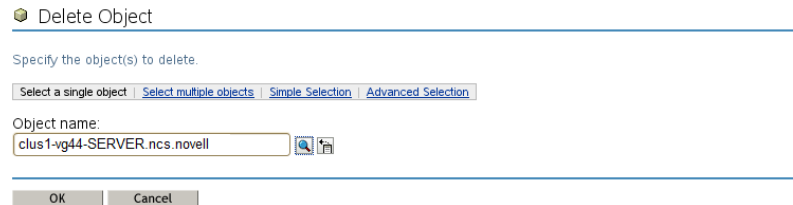


5c Click *OK* to delete the NCP volume object, then click *OK* again to close the success message.

6 If the Linux volume was NCP enabled, or if you manually created an NCP virtual server for the resource, delete the NCP virtual server object.

6a In iManager, select *Directory Administration > Delete Object*.

6b On the Delete Object page, browse to locate and select the resource's NCS:NCP Server object, such as `clus1-vg44-SERVER.ncs.novell`, then click *OK*.



6c Click *OK* to delete the NCP virtual server, then click *OK* again to close the success message.

7 (Optional) Delete the LVM volume group and logical volume:


7a Log in as the `root` user to the cluster node where the resource was online, then open a terminal console.

7b At a command prompt, launch NSSMU:

```
nssmu
```

7c In the NSSMU main menu, select *Linux Volumes*, then press Enter.

7d In the *Linux Volumes* list, select the clustered linux volume, then press Delete.



7e When you are prompted to confirm the delete action, press *y* (Yes) to continue, or press *n* (No) to cancel the delete action.


Deleting the volume and volume group puts the device in an uninitialized state.

7f Press *Esc* to return to the main menu

8 Re-initialize the device:

8a In the NSSMU main menu, select *Devices* and press *Enter*.

8b In the *Devices* list, select the device.



8c Press *F3* to initialize the device.

8d When you are prompted to confirm the initialization action, read the warning message, then press *y* (Yes) to continue, or press *n* (No) to cancel the action.

8e Select *DOS* or *GPT* as the partitioning type, then press *Enter*.

8f Visually verify that the device was initialized and the specified partition type is displayed.

8g Press *Esc* twice to exit NSSMU.

- 9 Verify that the LVM volume group is deleted by entering the following the native LVM command:

```
vgdisplay
```

8.11 Deleting a Clustered LVM Volume (Created in NSSMU or NLVM)

We strongly recommend that you delete a cluster-enabled LVM volume only from the master node in the cluster. This allows the cluster information to be automatically updated.

The procedures in this section assume that you created the clustered LVM volume in NSSMU or NLVM. There are default naming conventions applied by these tools that might not apply to LVM volume groups that you created and cluster-enabled by using native LVM tools and the Generic File System template.

WARNING: Deleting an LVM volume destroys all data on it.

NSSMU and the `nlvm delete linux volume <volume_name>` command delete the cluster-enabled LVM volume as well as the related objects in eDirectory:

- ♦ Linux LVM volume group and logical volume from the file system
- ♦ Cluster Resource object for the LVM resource
- ♦ If the LVM volume is NCP-enabled:
 - ♦ Volume object for the LVM volume
 - ♦ Virtual server for the LVM resource (NCS:NCP Server object)

When the LVM volume resides on the master node, the cluster information is automatically updated.

When the LVM volume resides on a non-master node, additional steps are required to update the cluster information. A cluster restart might be needed to force the information to be updated.

Use the following procedures to delete a cluster-enabled LVM volume:

- ♦ [Section 8.11.1, “Deleting a Cluster-Enabled LVM Volume on the Master Node,” on page 110](#)
- ♦ [Section 8.11.2, “Deleting a Cluster-Enabled LVM Volume on a Non-Master Node,” on page 111](#)

8.11.1 Deleting a Cluster-Enabled LVM Volume on the Master Node

- 1 If the LVM resource is on a non-master node in the cluster, migrate it to the master node. As the root user, open a terminal console, then enter

```
cluster migrate <resource_name> <master_node_name>
```

To migrate the resource, the master node must be in the resource's preferred nodes list.

- 2 Use the `cluster status` command to check the resource status. If the resource is online or comatose, take it offline.

As the root user, enter

```
cluster offline <resource_name>
```

Use the `cluster status` command to verify that the resource has a status of *Offline* before you continue.

- 3 Delete the LVM volume on the master node by using NSSMU.
You can alternatively use the `nlvm delete linux volume <lx_volume_name>` command.
 - 3a In NSSMU, select *Linux Volumes*, then press Enter.
 - 3b Select the unmounted LVM volume, then press Delete.
 - 3c Select OK to confirm, then press Enter.
- 4 In the *Tree View* in iManager, browse the objects to verify that the following objects were deleted:
 - ♦ LVM resource object (from the Cluster container)
 - ♦ If the LVM volume was NCP-enabled:
 - ♦ Volume object for the LVM volume
 - ♦ Virtual server for the LVM resource (NCS:NCP Server object)
- 5 Re-initialize the device that contained the LVM volume.
When NLVM or NSSMU removes the Linux LVM volume group, it leaves the device in an uninitialized state.
 - 5a In NSSMU, select *Devices*, then press Enter.
 - 5b Select the device, then press F3 (Initialize).
 - 5c Press y (Yes) to confirm.
 - 5d Select the partitioning scheme as DOS or GPT, then press Enter.
- 6 (Optional) Use a third-party SAN management tool to assign the device to only one desired server.

8.11.2 Deleting a Cluster-Enabled LVM Volume on a Non-Master Node

- 1 Log in as the root user to the non-master node where the cluster resource currently resides, then open a terminal console.
- 2 Use the `cluster status` command to check the resource status. If the resource is online or comatose, take it offline by using one of the following methods:


```
cluster offline <resource_name>
```

Use the `cluster status` command to verify that the resource has a status of *Offline* before you continue.
- 3 At the command prompt on the non-master node, enter


```
/opt/novell/ncs/bin/ncs-configd.py -init
```
- 4 Look at the file `/var/opt/novell/ncs/resource-priority.conf` to verify that it has the same information (REVISION and NUMRESOURCES) as the file on the master node.
- 5 Delete the LVM volume on the master node by using NSSMU.
You can alternatively use the `nlvm delete linux volume <lx_volume_name>` command.
 - 5a In NSSMU, select *Linux Volumes*, then press Enter.
 - 5b Select the unmounted LVM volume, then press Delete.
 - 5c Select OK to confirm, then press Enter.
- 6 In the *Tree View* in iManager, browse the objects to verify that the following objects were deleted:
 - ♦ LVM resource object (from the Cluster container)

- ♦ If the LVM volume was NCP-enabled:
 - ♦ Volume object for the LVM volume
 - ♦ Virtual server for the LVM resource (NCS:NCP Server object)
- 7** Re-initialize the device that contained the LVM volume.
- When NLVM or NSSMU removes the Linux LVM volume group, it leaves the device in an uninitialized state.
- 7a** In NSSMU, select *Devices*, then press Enter.
- 7b** Select the device, then press F3 (Initialize).
- 7c** Press y (Yes) to confirm.
- 7d** Select the partitioning scheme as DOS or GPT, then press Enter.
- 8** On the master node, log in as the root user, open a terminal console, then enter
- ```
/opt/novell/ncs/bin/ncs-configd.py -init
```
- 9** Look at the file `/var/opt/novell/ncs/resource-priority.conf` to verify that it has the same information (REVISION and NUMRESOURCES) as that of the non-master node where you deleted the cluster resource.
- 10** In iManager, select *Clusters > My Clusters*, select the cluster, then select the *Cluster Options* tab.
- 11** Click *Properties*, select the *Priorities* tab, then click *Apply* on the Priorities page.
- 12** At a command prompt, enter
- ```
cluster view
```
- The cluster view should be consistent.
- 13** Look at the file `/var/opt/novell/ncs/resource-priority.conf` on the master node to verify that the revision number increased.
- If the revision number increased, skip [Step 14](#).
- If the deleted resource is the only one in the cluster, the priority won't force the update. A phantom resource might appear in the interface. You need to restart Cluster Services to force the update, which also removes the phantom resource.
- 14** If the revision number did not automatically update in the previous steps, restart Novell Cluster Services by entering the following on one node in the cluster:
- ```
cluster restart [seconds]
```
- For *seconds*, specify a value of 60 seconds or more.
- For example:
- ```
cluster restart 120
```
- 15** (Optional) Use a third-party SAN management tool to assign the devices to only the desired server.

9 Backing Up Linux POSIX Volumes

Backup of Linux POSIX file systems requires that ACLS and POSIX permissions be set on the Linux path for the Linux User Management (LUM) enabled user performing the backup. The root user has all permissions needed to perform backup of any Linux path. You can use the Linux `chmod(8)` and `chown(8)` commands to give the backup user the Linux POSIX permissions to the directory being backed up.

The root user can back up any path in the Linux POSIX filesystem directory. For example, to use the `nbackup` command for Novell Storage Management Services to back up the data in path `/usr/novell/ncpvol2/Data/test*`, enter

```
nbackup -cvf ncptest.sidf -Uroot -Pnovell /usr/novell/ncpvol2/Data/test*
```

For information about the `nbackup` command options, see the `nbackup(1)` man page and the [OES 11 SP1: Storage Management Services Administration Guide for Linux](#).

For information about third-party products that support backup on Novell Open Enterprise Server, go to the [Open Enterprise Server Partners Web site \(http://www.novell.com/products/openenterpriseserver/partners.html\)](http://www.novell.com/products/openenterpriseserver/partners.html), then select *Backup* in the *Product Type* list.

A Documentation Updates

This section contains information about documentation content changes made to the *OES 11 SP1: Novell Linux POSIX Volumes Administration Guide* since the initial release of Novell Open Enterprise Server (OES) 11.

This document was updated on the following dates:

- [Section A.1, “December 3, 2013,” on page 115](#)
- [Section A.2, “September 19, 2013,” on page 115](#)
- [Section A.3, “July 31, 2013,” on page 115](#)
- [Section A.4, “April 2013 Scheduled Maintenance,” on page 116](#)
- [Section A.5, “January 2013 Scheduled Maintenance,” on page 116](#)
- [Section A.6, “January 11, 2013,” on page 116](#)
- [Section A.7, “November 2012 Scheduled Maintenance,” on page 117](#)
- [Section A.8, “September 2012 Scheduled Maintenance,” on page 119](#)
- [Section A.9, “August 28, 2012 \(OES 11 SP1\),” on page 119](#)
- [Section A.10, “January 18, 2012,” on page 121](#)

A.1 December 3, 2013

The new site for iManager 2.7.6 and earlier is at <https://www.netiq.com/documentation/imanager27>.

A.2 September 19, 2013

Modified Novell eDirectory 8.8 SP7 links to point to the [NetIQ eDirectory 8.8 SP7 Web site \(http://www.netiq.com/documentation/edir887\)](http://www.netiq.com/documentation/edir887).

A.3 July 31, 2013

Updates were made to the following section. The changes are described below.

A.3.1 Clustering LVM Volume Groups with Novell Cluster Services

Location	Change
“Deleting a Clustered LVM Volume (Created in NSSMU or NLVM)” on page 110	This section is new.

A.4 April 2013 Scheduled Maintenance

Updates were made to the following section. The changes are described below.

A.4.1 What's New for Linux POSIX Volumes

Location	Change
"What's New (OES 11 April 2013 Patches)" on page 17	This section is new.

A.5 January 2013 Scheduled Maintenance

Updates were made to the following section. The changes are described below.

A.5.1 What's New for Linux POSIX Volumes

Location	Change
Section 2.2, "What's New (OES 11 SP1 January 2013 Patches)," on page 17	This section is new.
"What's New (OES 11 January 2013 Patches)" on page 18	This section is new.

A.6 January 11, 2013

Updates were made to the following section. The changes are described below.

A.6.1 Planning for Linux POSIX Volumes

Location	Change
Section 5.2, “Using RAID Devices for Linux POSIX Volumes,” on page 29	<p>We recommend that you do not use Linux software RAIDs (such as MD RAIDs and Device Mapper RAIDs) for devices that you plan to use for storage objects that are managed by NSS management tools. The Novell Linux Volume Manager (NLVM) utility and the NSS Management Utility (NSSMU) list Linux software RAID devices that you have created by using Linux tools. Beginning with Linux Kernel 3.0 in OES 11 SP1, NLVM and NSSMU can see these devices, initialize them, and allow you to create storage objects on them. However, this capability has not yet been fully tested.</p> <p>IMPORTANT: In OES 11, a server hang or crash can occur if you attempt to use a Linux software RAID when you create storage objects that are managed by NSS management tools.</p>

A.7 November 2012 Scheduled Maintenance

Updates were made to the following sections. The changes are described below.

- [Section A.7.1, “Clustering LVM Volume Groups with Novell Cluster Services,” on page 117](#)
- [Section A.7.2, “Managing Linux Volumes with NLVM Commands,” on page 118](#)
- [Section A.7.3, “Managing Linux Volumes with NSSMU,” on page 118](#)
- [Section A.7.4, “Overview of Linux POSIX File Systems,” on page 118](#)
- [Section A.7.5, “What’s New for Linux POSIX Volumes,” on page 118](#)

A.7.1 Clustering LVM Volume Groups with Novell Cluster Services

Location	Change
Section 8.3.3, “Creating an LVM Volume Group Cluster Resource with NLVM Commands,” on page 88	<p>If you enable NCP for the volume, the LVM volume name must comply with the limitations for NCP volume names described in Section 5.4, “Naming Conventions for NCP Volume Names,” on page 31. If you use lowercase letters for the volume name, they are automatically changed to uppercase for the NCP volume name.</p>

A.7.2 Managing Linux Volumes with NLVM Commands

Location	Change
Section 7.7, “Creating a Linux POSIX Volume,” on page 66	If you enable NCP for the volume, the volume name must comply with the limitations for NCP volume names described in Section 5.4, “Naming Conventions for NCP Volume Names,” on page 31 . If you use lowercase letters for the volume name, they are automatically changed to uppercase for the NCP volume name.
Section 7.8, “Creating an LVM Logical Volume,” on page 67	

A.7.3 Managing Linux Volumes with NSSMU

Location	Change
Section 6.9, “Creating a Non-LVM Linux Volume,” on page 44	If you enable NCP for the volume, the volume name must comply with the limitations for NCP volume names described in Section 5.4, “Naming Conventions for NCP Volume Names,” on page 31 . If you use lowercase letters for the volume name, they are automatically changed to uppercase for the NCP volume name.
Section 7.8, “Creating an LVM Logical Volume,” on page 67	

A.7.4 Overview of Linux POSIX File Systems

Location	Change
“Default Ext3 Inode Size” on page 12	To allow space for extended attributes and ACLs for a file on Ext3 file systems, the default inode size for Ext3 was increased from 128 bytes on SLES 10 to 256 bytes on SLES 11.

A.7.5 What’s New for Linux POSIX Volumes

Location	Change
Section 2.4, “What’s New (OES 11 SP1 November 2012 Patches),” on page 19	This section is new.
“Linux Volume Support in SLES 11 SP2” on page 22	This section is new.
“Linux Volume Support in SLES 11 SP1” on page 22	This section is new.
“Linux Volume Support in SLES 11” on page 23	This section is new.

A.8 September 2012 Scheduled Maintenance

Updates were made to the following sections. The changes are described below.

- ♦ [Section A.8.1, “What’s New for Linux POSIX Volumes,” on page 119](#)

A.8.1 What’s New for Linux POSIX Volumes

Location	Change
Section 2.5, “What’s New (OES 11 SP1 September 2012 Patches),” on page 19	This section is new.
Section 2.6, “What’s New (OES 11 September 2012 Patches),” on page 20	This section is new.

A.9 August 28, 2012 (OES 11 SP1)

Updates were made to the following sections. The changes are described below.

- ♦ [Section A.9.1, “Backing Up Linux POSIX Volumes,” on page 119](#)
- ♦ [Section A.9.2, “Clustering LVM Volume Groups with Novell Cluster Services,” on page 119](#)
- ♦ [Section A.9.3, “Installing or Upgrading NSS Management Tools and Linux POSIX File Systems,” on page 120](#)
- ♦ [Section A.9.4, “Managing Linux POSIX Volumes with NLVM Commands,” on page 120](#)
- ♦ [Section A.9.5, “Managing Linux POSIX Volumes with NSSMU,” on page 120](#)
- ♦ [Section A.9.6, “Planning for Linux POSIX Volumes,” on page 120](#)
- ♦ [Section A.9.7, “Using NLVM and NSSMU in a Virtualized Environment,” on page 120](#)
- ♦ [Section A.9.8, “What’s New for Linux POSIX Volumes,” on page 121](#)

A.9.1 Backing Up Linux POSIX Volumes

The [Chapter 9, “Backing Up Linux POSIX Volumes,” on page 113](#) section is new.

A.9.2 Clustering LVM Volume Groups with Novell Cluster Services

Location	Change
Section 8.3, “Configuring an LVM Volume Group Cluster Resource with NSS Management Tools,” on page 81	Use the NCP option to enable the Linux POSIX volume to be accessed with the NetWare Core Protocol (NCP).
Section 8.8, “Renaming a Clustered LVM Logical Volume,” on page 102	This section is new.

A.9.3 Installing or Upgrading NSS Management Tools and Linux POSIX File Systems

The [Chapter 3, “Installing or Upgrading NSS Management Tools and Linux POSIX File Systems,”](#) on [page 25](#) section is new.

A.9.4 Managing Linux POSIX Volumes with NLVM Commands

Location	Change
Section 7.8, “Creating an LVM Logical Volume,” on page 67	Multiplier can also be T (terabytes). The default multiplier is G (gigabytes).
Section 7.7, “Creating a Linux POSIX Volume,” on page 66	Use the NCP option to enable the Linux POSIX volume to be accessed with the NetWare Core Protocol.
Section 7.8, “Creating an LVM Logical Volume,” on page 67	

A.9.5 Managing Linux POSIX Volumes with NSSMU

Location	Change
Section 6.9, “Creating a Non-LVM Linux Volume,” on page 44	Use the NCP option to enable the Linux POSIX volume to be accessed with the NetWare Core Protocol (NCP).
Section 6.10, “Creating an LVM Logical Volume,” on page 48	
Section 6.14, “Renaming the Mount Point Path for a Linux Volume,” on page 53	This section is new.
Section 6.15, “Renaming a Non-LVM Linux Volume,” on page 55	This section is new.
Section 6.16, “Renaming an LVM Volume,” on page 55	This section is new.

A.9.6 Planning for Linux POSIX Volumes

The [Chapter 5, “Planning for Linux POSIX Volumes,”](#) on [page 29](#) section is new.

A.9.7 Using NLVM and NSSMU in a Virtualized Environment

The [Chapter 4, “Using NLVM and NSSMU in a Virtualized Environment,”](#) on [page 27](#) section is new.

A.9.8 What's New for Linux POSIX Volumes

Location	Change
Section 2.7, "What's New (OES 11 SP1)," on page 20	This section is new.

A.10 January 18, 2012

The document format was updated to reflect newly revised corporate standards.

