

Linux POSIX Volume Administration Guide

Open Enterprise Server 11

July 25, 2012

Novell.

Legal Notices

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

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

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

Copyright © 2011–2012 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

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

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

Novell Trademarks

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

Third-Party Materials

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

Contents

About This Guide	5
1 Overview of Linux POSIX File Systems	7
1.1 Ext2	7
1.2 Ext3	8
1.3 ReiserFS	9
1.4 XFS	10
1.5 What's Next	11
2 What's New or Changed for Linux POSIX Volumes	13
2.1 What's New (May 2012 Patches)	13
2.2 What's New (April 2012 Patches)	13
2.3 What's New (OES 11)	14
3 Managing Linux POSIX Volumes with NSSMU	15
3.1 NSSMU Quick Reference for Linux Volumes	16
3.2 Launching NSSMU	17
3.3 Viewing a List of Devices	18
3.4 Viewing Device Details	18
3.5 Initializing a Device	20
3.6 Unsharing a Device	21
3.7 Viewing a List of Linux Volumes	22
3.8 Viewing Details for a Linux Volume	22
3.9 Creating a Linux POSIX Volume	25
3.10 Creating an LVM Logical Volume	27
3.11 Creating a Shared LVM Volume with Novell Cluster Services	30
3.12 Mounting a Linux Volume	30
3.13 Dismounting a Linux Volume	31
3.14 Renaming the Mount Point Path for a Linux Volume	31
3.15 Deleting a Linux POSIX Volume	33
4 Managing Linux POSIX Volumes with NLVM Commands	35
4.1 NLVM Commands Quick Reference for Linux Volumes	36
4.2 Viewing a List of Devices	38
4.3 Viewing Device Details	38
4.4 Initializing a Device	40
4.5 Unsharing a Device	41
4.6 Viewing a List of Linux Volumes	42
4.7 Creating a Linux POSIX Volume	42
4.8 Creating an LVM Logical Volume	43
4.9 Creating a Shared LVM Logical Volume with Novell Cluster Services	45
4.10 Mounting Linux Volumes	45
4.11 Dismounting Linux Volumes	46
4.12 Rescanning for Storage Objects on Known Devices	47

4.13	Deleting a Linux Volume	48
4.13.1	Deleting a Linux POSIX Volume	48
4.13.2	Deleting an LVM Volume Group and Logical Volume	49
4.14	Additional Information	50
5	Clustering LVM Volume Groups with Novell Cluster Services	51
5.1	Requirements for Creating LVM Cluster Resources	51
5.1.1	Novell Cluster Services	52
5.1.2	Linux Logical Volume Manager 2 (LVM2)	52
5.1.3	Clustered Logical Volume Manager Daemon (CLVMD)	52
5.1.4	Resource IP Address	52
5.1.5	Shared Storage Devices	52
5.1.6	All Nodes Must Be Present	52
5.1.7	Mixed Mode OES Clusters (Not Supported)	52
5.1.8	NCP File Access with Novell NCP Server	53
5.1.9	CIFS/Samba File Access with Novell Samba	53
5.2	Initializing a SAN Device	53
5.3	Creating an LVM Cluster Resource	54
5.3.1	Using NSSMU to Create the Resource	55
5.3.2	Using NLVM to Create the Resource	57
5.4	Configuring the LVM Cluster Resource Settings	58
5.5	Viewing or Modifying the LVM Resource Scripts	61
5.6	Sample LVM Resource Scripts	62
5.6.1	Sample LVM Resource Load Script	62
5.6.2	Sample LVM Resource Unload Script	63
5.6.3	Sample LVM Resource Monitor Script	63
5.7	Renaming the Mount Point Path for a Clustered LVM Volume	64
5.8	Disabling Clustering for an LVM Volume	65
5.9	Deleting a Clustered LVM Volume Group and Logical Volume	67
6	Backing Up Linux POSIX Volumes	69
A	Documentation Updates	71
A.1	July 25, 2012	71
A.1.1	Backing Up Linux POSIX Volumes	71
A.1.2	NLVM Commands	71
A.2	May 31, 2012	72
A.2.1	Clustering LVM Volume Groups with Novell Cluster Services	72
A.2.2	Managing Linux POSIX Volumes with NLVM Commands	72
A.2.3	Managing Linux POSIX Volumes with NSSMU	72
A.2.4	What's New or Changed for Linux POSIX Volumes	73
A.3	April 30, 2012	73
A.3.1	Clustering LVM Volume Groups with Novell Cluster Services	73
A.3.2	Managing Linux POSIX Volumes with NLVM Commands	73
A.3.3	Managing Linux POSIX Volumes with NSSMU	73
A.3.4	What's New or Changed for Linux POSIX Volumes	74
A.4	January 18, 2012	74

About This Guide

This guide describes how to manage Linux POSIX volumes on a Novell Open Enterprise Server (OES) 11 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 7](#)
- ♦ [Chapter 2, “What’s New or Changed for Linux POSIX Volumes,” on page 13](#)
- ♦ [Chapter 3, “Managing Linux POSIX Volumes with NSSMU,” on page 15](#)
- ♦ [Chapter 4, “Managing Linux POSIX Volumes with NLVM Commands,” on page 35](#)
- ♦ [Chapter 5, “Clustering LVM Volume Groups with Novell Cluster Services,” on page 51](#)
- ♦ [Chapter 6, “Backing Up Linux POSIX Volumes,” on page 69](#)
- ♦ [Appendix A, “Documentation Updates,” on page 71](#)

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: Linux POSIX Volume Administration Guide*, visit the [OES 11 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: NLVM Reference](#).

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

For information about clustering with Novell Cluster Services, see the [OES 11: Novell Cluster Services 2.0 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 1 (SP1), see the [SLES 11 SP1 Storage Administration Guide \(http://www.suse.com/documentation/sles11/stor_admin/data/bookinfo.html\)](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 provides several native Linux POSIX file systems. In Novell Open Enterprise Server (OES) 11, the Novell Storage Services (NSS) Management Utility (NSSMU) and the Novell Linux Volume Manager (NLVM) commands are designed to be used with the 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, “Ext2,” on page 7](#)
- ♦ [Section 1.2, “Ext3,” on page 8](#)
- ♦ [Section 1.3, “ReiserFS,” on page 9](#)
- ♦ [Section 1.4, “XFS,” on page 10](#)
- ♦ [Section 1.5, “What’s Next,” on page 11](#)

The NSSMU and NLVM tools also support the NSS file system on your OES 11 servers. 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: 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 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: Novell Cluster Services 2.0 for Linux Administration Guide](#).

1.1 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.

1.2 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 htree feature significantly improves its scalability over Ext2. However, it is still not as scalable as some of the other file systems in Linux. With htree, Ext3 scales similarly to NTFS. Without htree, 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.

1.3 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 htrees. 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.

1.4 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 multi-media 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.

1.5 What's Next

To use NSSMU to create and manage Linux POSIX volumes and LVM volumes, see [Chapter 3, "Managing Linux POSIX Volumes with NSSMU,"](#) on page 15.

To use NLVM commands to create and manage Linux POSIX volumes and LVM logical volumes, see [Chapter 4, "Managing Linux POSIX Volumes with NLVM Commands,"](#) on page 35.

To use NSSMU or NLVM to cluster-enable LVM volume groups and logical volumes, see [Chapter 5, "Clustering LVM Volume Groups with Novell Cluster Services,"](#) on page 51.

2 What's New or Changed for Linux POSIX Volumes

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

- ♦ [Section 2.1, “What’s New \(May 2012 Patches\),” on page 13](#)
- ♦ [Section 2.2, “What’s New \(April 2012 Patches\),” on page 13](#)
- ♦ [Section 2.3, “What’s New \(OES 11\),” on page 14](#)

2.1 What’s New (May 2012 Patches)

In addition to bug fixes, the following changes and enhancements were made for Novell Cluster Services in the OES 11 May 2012 Scheduled Maintenance patches:

NSSMU

- ♦ **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.

NLVM

- ♦ **Confirmation Prompt:** You are automatically prompted to confirm actions to initialize a device or to delete a Linux POSIX volume.
- ♦ **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.

2.2 What’s New (April 2012 Patches)

In addition to bug fixes, the following changes and enhancements were made for Novell Cluster Services in the OES 11 April 2012 Scheduled Maintenance patches:

NSSMU

The `ncp` option for creating a Linux POSIX volume can be used to enable the Linux POSIX file system on the volume to be accessed with the NetWare Core Protocol (NCP).

NLVM

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).

2.3 What's New (OES 11)

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
- ♦ 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

3 Managing Linux POSIX Volumes with NSSMU

The Novell Storage Services (NSS) Management Utility (NSSMU) for Novell Open Enterprise Server (OES) 11 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 11 server.

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

- ♦ [Section 3.1, “NSSMU Quick Reference for Linux Volumes,” on page 16](#)
- ♦ [Section 3.2, “Launching NSSMU,” on page 17](#)
- ♦ [Section 3.3, “Viewing a List of Devices,” on page 18](#)
- ♦ [Section 3.4, “Viewing Device Details,” on page 18](#)
- ♦ [Section 3.5, “Initializing a Device,” on page 20](#)
- ♦ [Section 3.6, “Unsharing a Device,” on page 21](#)
- ♦ [Section 3.7, “Viewing a List of Linux Volumes,” on page 22](#)
- ♦ [Section 3.8, “Viewing Details for a Linux Volume,” on page 22](#)
- ♦ [Section 3.9, “Creating a Linux POSIX Volume,” on page 25](#)
- ♦ [Section 3.10, “Creating an LVM Logical Volume,” on page 27](#)
- ♦ [Section 3.11, “Creating a Shared LVM Volume with Novell Cluster Services,” on page 30](#)
- ♦ [Section 3.12, “Mounting a Linux Volume,” on page 30](#)
- ♦ [Section 3.13, “Dismounting a Linux Volume,” on page 31](#)
- ♦ [Section 3.14, “Renaming the Mount Point Path for a Linux Volume,” on page 31](#)
- ♦ [Section 3.15, “Deleting a Linux POSIX Volume,” on page 33](#)

3.1 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 server. To launch NSSMU, log in to the server as the `root` user, open a terminal console, then enter `nssmu`.

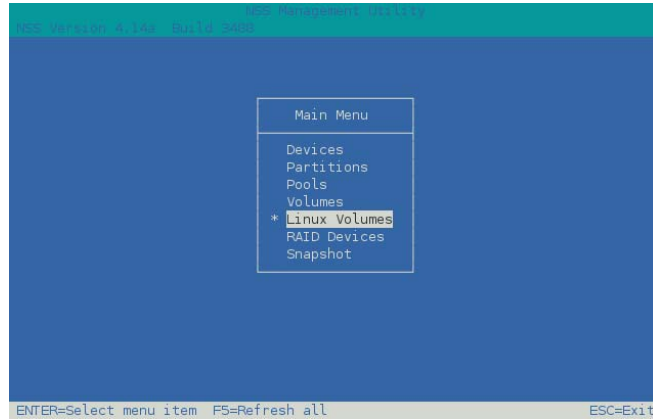
IMPORTANT: For information about the NSS file system management features of NSSMU, see the “NSS Management Utility (NSSMU) Quick Reference” in the *OES 11: NSS File System Administration Guide for Linux*.

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	
Enter = Show partitions (does not report the 4K partition created for Shareable for Clustering)	The DOS partitioning scheme supports devices up to 2 TB (terabytes) in size. It allows up to four partitions on a device.
Esc = Previous menu	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.
	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>Linux Volumes</i>	Use the Linux volumes page to create one of the following types of Linux volumes:
Ins = Create	
Del = Delete	<ul style="list-style-type: none">Native Linux POSIX volumes (non-LVM)LVM volume group and logical volumeClustered LVM volume group, logical volume, and Novell Cluster Services cluster resource
F3 = Rename	
F5 = Refresh details of the Linux volumes	NSSMU supports Ext2, Ext3, ReiserFS, and XFS file systems.
F7 = Mount or Dismount	You can mount or dismount the Linux volume.
Esc = Previous menu	You can delete a Linux volume. The volume must be dismounted (or offline for clustered volumes) before you can delete it.

3.2 Launching 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:

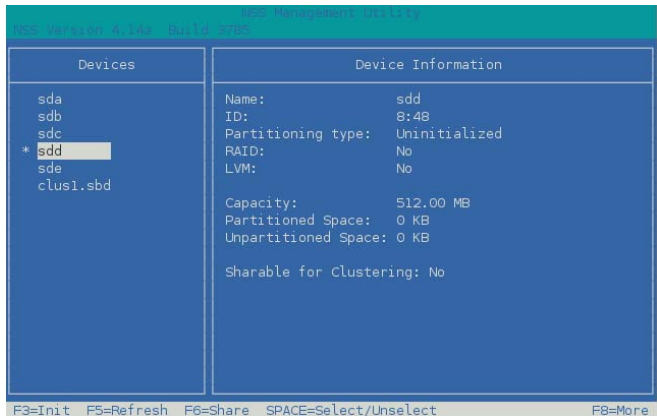
NSSMU Main Menu	Description
Devices	View information about devices. Initialize a device in MSDOS or GPT format. Enable or disable a device as <i>Shareable for Clustering</i> . This is used for devices that you plan to use as the Split Brain Detector (SBD) for a Novell Cluster Services cluster or for NSS pools.
Partitions	View a list of partitions on a device, including Linux and LVM partitions.
Pools	Create and manage NSS pools. Create and cluster-enable NSS pools on devices marked as Shareable for Clustering.
Volumes	Create and manage NSS volumes. Create and cluster-enable an LVM volume group and logical volume.
Linux Volumes	Create and manage Linux POSIX file systems, such as Ext2, Ext3, Reiser, and XFS.
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.

3.3 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 3.4, “Viewing Device Details,”](#) on page 18.
- 4 Press Esc twice to exit NSSMU.

3.4 Viewing Device Details

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

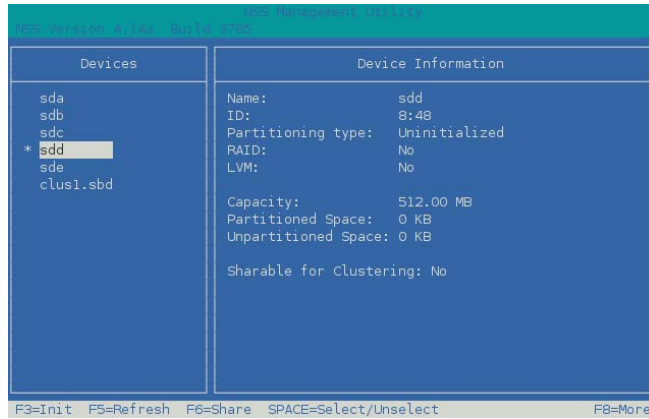
Parameter	Description
Name	Specifies the Linux device node name in the /dev directory. Examples: sda, sdb, sdc, sdd, and so on
ID	Specifies the major and minor number assigned to a block device on the server. Example: 8:0, 8:16, 8:32, 8:48

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 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.
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.

3.5 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.

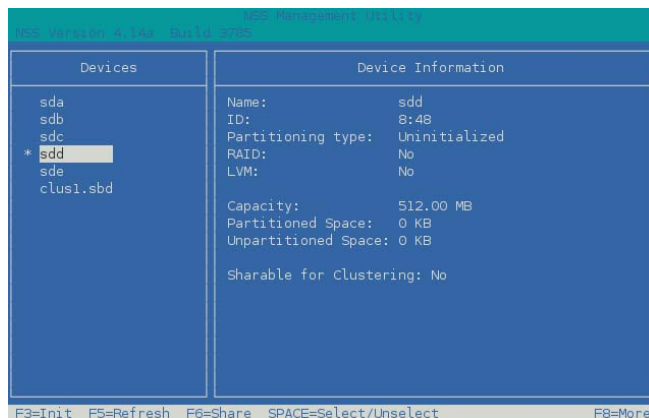
Initializing a device formats it with an MSDOS or a GPT partitioning scheme. MSDOS supports devices up to 2 TB in size. GPT supports devices of any size. The default partitioning scheme is MSDOS for devices less than 2 TB, and GPT for devices greater than or equal to 2 TB. NSSMU automatically selects 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*.

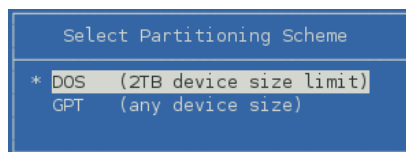
If a device has never been initialized, its *Partitioning type* is reported as *Uninitialized*.



- 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.

3.6 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*.

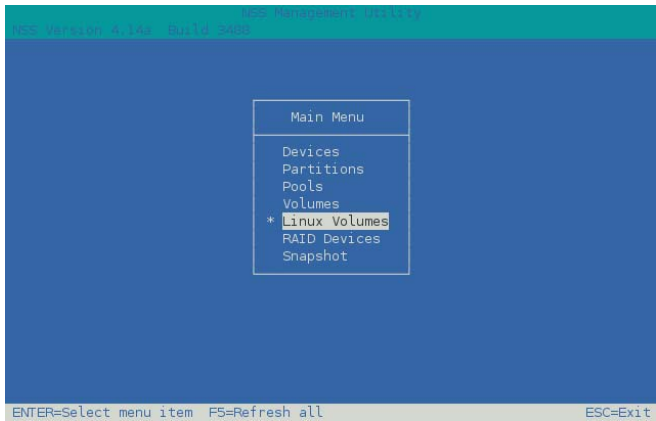
To unshare a device:

- 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.

3.7 Viewing a List of Linux Volumes

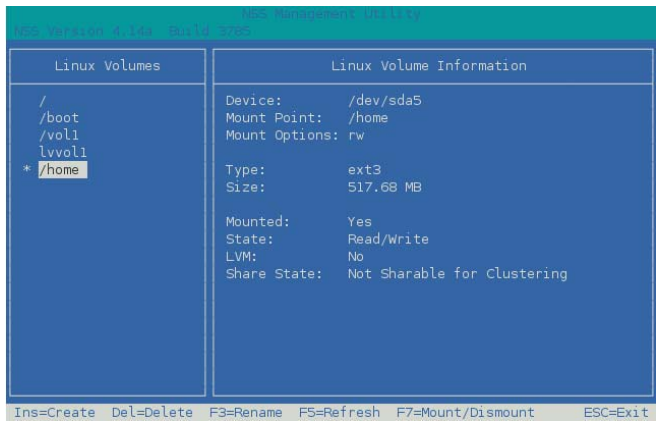
You can use the NSSMU Linux Volumes page to view information about Linux POSIX and LVM logical volumes that are mounted on the system.

- 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 as the final directory of the mount path, such as `/home`. LVM logical volumes show the volume name as the name you gave the volume when you created it, such as `lvvol1`.



- 3 Press Esc twice to exit NSSMU.

3.8 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: <code>/dev/sda2</code> Example for an LVM volume: <code>/dev/lvvol1/lvvol1</code>

Parameter	Description
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: ext2, ext3, reiserfs, xfs
Size	Specifies the amount of space reserved for this volume. Example: 30.58 GB
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	Specifies whether NCP is enabled for the volume. Values: Enabled or Not Enabled
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:

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

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

Linux Volumes	
/	Device: /dev/LXVOLD/LXVOLD
/boot	Mount Point: /usr/novell/LXVOLD
* LXVOLD	Mount Options: rw
	Type: ext3
	Size: 508.00 MB
	Mounted: Yes
	State: Read/Write
	LVM: Yes
	NCP: Not Enabled
	Share State: Sharable for Clustering

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*.

Linux Volumes	
/	Device: /dev/vol44/vol44
/boot	Mount Point: /mnt/vol44
* vol44	Mount Options: rw
	Type: ext3
	Size: 99.68 GB
	Mounted: Yes
	State: Read/Write
	LVM: Yes
	NCP: Not Enabled
	Share State: Shareable for Clustering

3 Press Esc twice to exit NSSMU.

3.9 Creating a Linux POSIX Volume




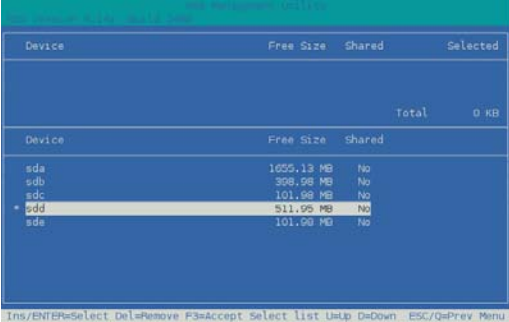
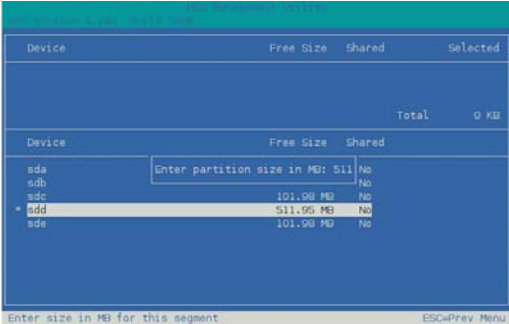
You can use NSSMU to create a native Linux POSIX volume on a device. You must have free unpartitioned space available on a device. The device should not be shareable for clustering.

- 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 3.5, “Initializing a Device,” on page 20](#).

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:

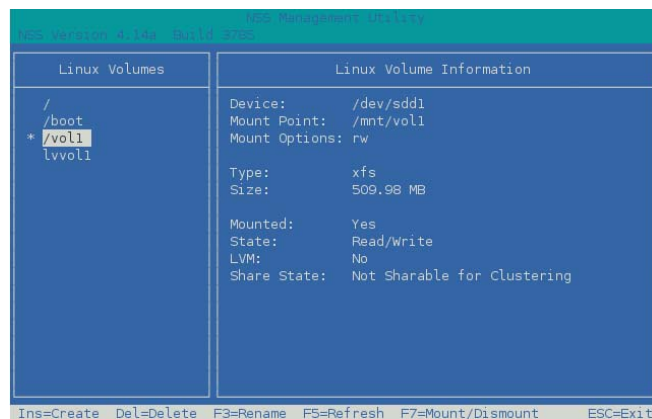
Parameter	Description
LVM Type	Select <i>Non-LVM2 Volume</i> from the following volume types to create a native Linux POSIX volume, then press Enter. <ul style="list-style-type: none">♦ Non-LVM2 Volume (local disk)♦ LVM2 Volume (local disk)♦ Cluster Enabled LVM2 Volume
NCP enable volume	Type Y (Yes) to enable NCP for the volume. An NCP share is added at the root of the Linux POSIX volume. The NCP share and volume name is the same as for the Linux POSIX volume. The default mount point is <code>/usr/novell/<volumename></code> . Type N (No) to not allow NCP access for the root of the volume at this time. After the Linux POSIX volume is created, you can enable NCP at any time by creating NCP shares on the root of the volume or on subdirectories. The share names can be different than the name of the Linux POSIX volume. In Novell Remote Manager, select <i>Manage NCP Services > Manage Shares</i> , then click <i>Create New Share</i> . For more information, see “Managing NCP Volumes” in the OES 11: NCP Server for Linux Administration Guide .
Volume type (file system to use for the volume)	Select one of the following file system options to make on the volume, then press Enter. <ul style="list-style-type: none">♦ ext2♦ ext3♦ reiserfs♦ xfs

Parameter	Description
<p>Mount Point</p> 	<p>Specify the full Linux path where you want to mount the device, then press Enter. For example, <code>/mnt/vol1</code>.</p> <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>Make Options</p> 	<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.ext2(8)</code>, <code>mkfs.ext3(8)</code>, <code>mkfs.reiser(8)</code>, and <code>mkfs.xfs(8)</code> man pages. If the make option is invalid for the specified file system, the file system creation fails.</p>
<p>Mount Options</p> 	<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>
<p>Devices</p> 	<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>Partition Size</p> 	<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>

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

- 7 View the details of the newly created volume.

For information, see [Section 3.4, “Viewing Device Details,”](#) on page 18.



- 8 Press Esc twice to exit NSSMU.

3.10 Creating an LVM Logical Volume

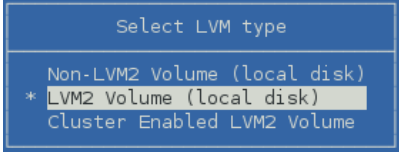
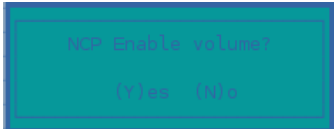

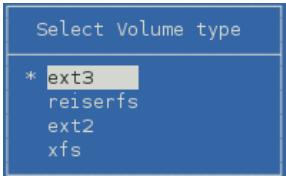

You can use NSSMU to create a Linux LVM volume group and logical volume on a device, make a file system on the volume, and mount the volume. You must have free unpartitioned space available on a device. The device should not be shareable for clustering.

- 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 3.5, “Initializing a Device,”](#) on page 20.

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

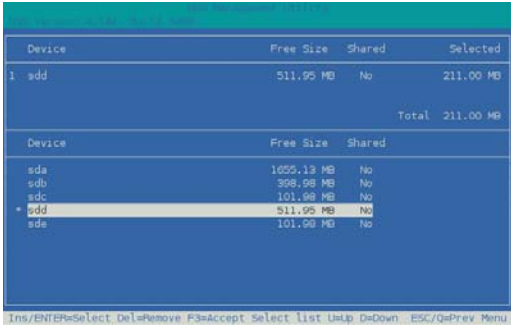
- 4 From the NSSMU main menu, select *Linux Volumes*, then press Enter.

- 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
<p>LVM Type</p> 	<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
<p>NCP enable volume</p> 	<p>Type Y (Yes) to enable NCP for the volume. An NCP share is added at the root of the Linux POSIX volume. The NCP share and volume name is the same as for the Linux POSIX volume. The default mount point is <code>/usr/novell/<volumename></code>.</p> <p>Type N (No) to not allow NCP access for the root of the volume at this time.</p> <p>After the Linux POSIX volume is created, you can enable NCP at any time by creating NCP shares on the root of the volume or on subdirectories. The share names can be different than the name of the Linux POSIX volume. In Novell Remote Manager, select <i>Manage NCP Services > Manage Shares</i>, then click <i>Create New Share</i>. For more information, see “Managing NCP Volumes” in the <i>OES 11: NCP Server for Linux Administration Guide</i>.</p>
<p>Volume name</p> 	<p>Type a name to use for the LVM volume group and logical volume (such as <code>lvvol1</code>), then press Enter.</p>
<p>Volume type (file system to use for the volume)</p> 	<p>Select one of the following file system options to make on the volume, then press Enter:</p> <ul style="list-style-type: none"> ♦ ext3 ♦ reiserfs ♦ ext2 ♦ xfs
<p>Mount Point</p> 	<p>Specify the full Linux path where you want to mount the device, then press Enter. For example, <code>/home</code>.</p> <p>The mount path's final directory can be the same or different than the name used for the logical volume.</p>

Parameter	Description
<div>Make Options</div> <div>Enter any make options: </div>	<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.ext2(8)</code>, <code>mkfs.ext3(8)</code>, <code>mkfs.reiser(8)</code>, and <code>mkfs.xfs(8)</code> man pages. If the make option is invalid for the specified file system, the file system creation fails.</p>
<div>Mount Options</div> <div>Enter any mount options: rw</div>	<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>
<div>Devices</div> <div></div>	<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>
<div>Partition Size</div> <div></div>	<p>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>

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



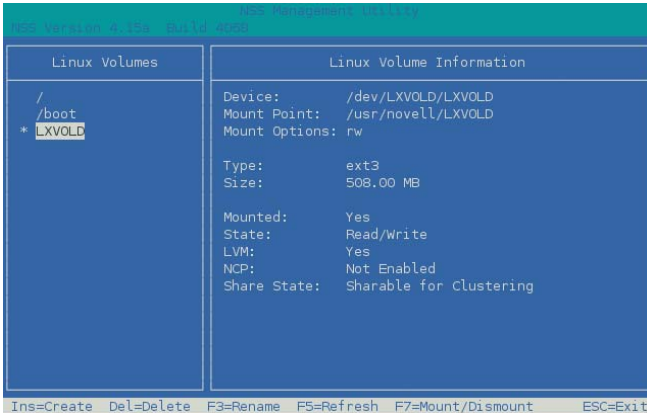
Device	Free Size	Shared	Selected
1 sdd	511.95 MB	No	211.00 MB
			Total 211.00 MB

Device	Free Size	Shared
sda	1055.13 MB	No
sdb	398.98 MB	No
sdc	101.88 MB	No
* sdd	511.95 MB	No
sde	101.98 MB	No

Ins/Enter=Select Del=Remove F3=Accept Select list Up/Down ESC/Q=Prev Menu

7 View the details of the newly created volume.

For information, see [Section 3.4, “Viewing Device Details,”](#) on page 18.



Linux Volumes	Linux Volume Information
/boot * LXVOLD	Device: /dev/LXVOLD/LXVOLD Mount Point: /usr/novell/LXVOLD Mount Options: rw Type: ext3 Size: 508.00 MB Mounted: Yes State: Read/Write LVM: Yes NCP: Not Enabled Share State: Sharable for Clustering

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

8 Press Esc twice to exit NSSMU.

3.11 Creating a Shared LVM Volume with Novell Cluster Services

For information about creating a clustered LVM volume group cluster resource with NSSMU, see [Section 5.3.1, “Using NSSMU to Create the Resource,”](#) on page 55 in [Chapter 5, “Clustering LVM Volume Groups with Novell Cluster Services,”](#) on page 51.

3.12 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.

3.13 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.

3.14 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:

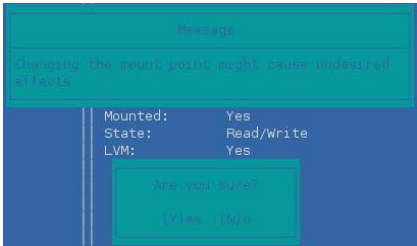
Volume Information	New Mount Point and Change	New Volume Name
Volume Type: Non-LVM Mount Point: <code>/projects/volp1</code> Volume Name: <code>/volp1</code>	Change the path to the final directory. Use the same name for the final directory. <code>/old/projects/volp1</code>	The volume name does not change.
Volume Type: Non-LVM Mount Point: <code>/home/janedoe</code> Volume Name: <code>/janedoe</code>	Change the final directory name. The preceding directories can be the same or different. <code>/home/janebuck</code> or <code>/users/janebuck</code>	The new name for the final directory becomes the new volume name: <code>/janebuck</code> 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 5.7, “Renaming the Mount Point Path for a Clustered LVM Volume,”](#) on page 64.

- 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*.

3.15 Deleting a Linux POSIX Volume

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

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 volumes. For information about deleting a clustered LVM volume group and logical volume, see [Section 5.9, “Deleting a Clustered LVM Volume Group and Logical Volume,”](#) on page 67.

- 1 From the NSSMU menu, select *Linux Volumes*, then press Enter.
- 2 In the *Linux Volumes* list, select the volume.
- 3 If the volume is mounted, press F7 to dismount it.
- 4 In the *Linux Volumes* list, select the volume, then verify that the Mounted state is *No*.
- 5 Press *Delete* to delete the volume.
- 6 Press y (yes) to confirm, or press n (no), or press Esc to cancel the delete process.
- 7 In the *Linux Volumes* list, verify that the volume is no longer displayed.
- 8 Press Esc to return to the main menu, then select *Devices*.
- 9 Select the device that contained the volume, then verify that the amount of free space increased.
- 10 Press Esc twice to exit NSSMU.

4 Managing Linux POSIX 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: NLVM Reference](#).

- ♦ [Section 4.1, “NLVM Commands Quick Reference for Linux Volumes,” on page 36](#)
- ♦ [Section 4.2, “Viewing a List of Devices,” on page 38](#)
- ♦ [Section 4.3, “Viewing Device Details,” on page 38](#)
- ♦ [Section 4.4, “Initializing a Device,” on page 40](#)
- ♦ [Section 4.5, “Unsharing a Device,” on page 41](#)
- ♦ [Section 4.6, “Viewing a List of Linux Volumes,” on page 42](#)
- ♦ [Section 4.7, “Creating a Linux POSIX Volume,” on page 42](#)
- ♦ [Section 4.8, “Creating an LVM Logical Volume,” on page 43](#)
- ♦ [Section 4.9, “Creating a Shared LVM Logical Volume with Novell Cluster Services,” on page 45](#)
- ♦ [Section 4.10, “Mounting Linux Volumes,” on page 45](#)
- ♦ [Section 4.11, “Dismounting Linux Volumes,” on page 46](#)
- ♦ [Section 4.12, “Rescanning for Storage Objects on Known Devices,” on page 47](#)
- ♦ [Section 4.13, “Deleting a Linux Volume,” on page 48](#)
- ♦ [Section 4.14, “Additional Information,” on page 50](#)

4.1 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: NLVM Reference*.

Task	Command
Create a clustered Linux LVM volume group and logical volume See also Chapter 5, “Clustering LVM Volume Groups with Novell Cluster Services,” on page 51.	<pre>nlvm create linux volume type=<fstype> device=<devicename> mp=<full_mount_point_path> [mntopt=<fs_mount_options>] lvm name=<lv_name> [group=<vg_name>] shared ip=<resource_ip_address></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>
Create a Linux LVM volume group and logical volume See also Section 4.8, “Creating an LVM Logical Volume,” on page 43.	<pre>nlvm create linux volume type=<fstype> device=<devicename> size=<value[K M G]> mp=<full_mount_point_path> [mntopt=<fs_mount_options>] lvm name=<lv_name> [group=<vg_name>]</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
Create a Linux POSIX volume See also Section 4.7, “Creating a Linux POSIX Volume,” on page 42.	<pre>nlvm create linux volume type=<fstype> device=<device_name anydisk> size=<value[K M G] max> mp=<full_mount_point_path> [mntopt=<fs_mount_options>]</pre> <p>Example:</p> <pre>nlvm create linux volume type=ext3 device=sda size=20G mp=/mnt/vol1 mntopt=rw</pre>
Delete a Linux volume See also Section 4.13.1, “Deleting a Linux POSIX Volume,” on page 48. See also Section 4.13.2, “Deleting an LVM Volume Group and Logical Volume,” on page 49. See also Section 5.9, “Deleting a Clustered LVM Volume Group and Logical Volume,” on page 67.	<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>
Initialize a device See also Section 4.4, “Initializing a Device,” on page 40.	<pre>nlvm [--force] init <device_name> [format=<gpt msdos>] unshared</pre> <p>Example:</p> <pre>nlvm init sde format=msdos unshared</pre>
Rescan for storage objects on known devices See also Section 4.12, “Rescanning for Storage Objects on Known Devices,” on page 47.	<pre>nlvm [-m] rescan</pre>
Unshare a device See also Section 4.5, “Unsharing a Device,” on page 41.	<pre>nlvm unshare <device></pre> <p>Example:</p> <pre>nlvm unshare sde</pre>
View a list of active devices See also Section 4.2, “Viewing a List of Devices,” on page 38.	<pre>nlvm list devices [more] [exclude=<raid nonraid shared nonshared lvm nonlvm>]</pre>
View a list of Linux volumes See also Section 4.6, “Viewing a List of Linux Volumes,” on page 42.	<pre>nlvm list linux volumes</pre>
View details for a device See also Section 4.3, “Viewing Device Details,” on page 38.	<pre>nlvm list <device></pre> <p>Example:</p> <pre>nlvm list device sde</pre>

4.2 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: 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
```

4.3 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: NLVM Reference*.

The command returns the following information:

Device name

Device major:minor
Size of device in MB or GB
Free space remaining on the device in KB, MB, or GB
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
```

4.4 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: NLVM Reference*.

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: Novell Cluster Services 2.0 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 initialization 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
```

4.5 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: 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*.

4.6 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: 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
```

4.7 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: 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=<fstype>
  device=<device_name|anydisk>
  size=<value[K|M|G]|max>
  mp=<full_mount_point_path>
  [mntopt=<fs_mount_options>]
  [-ncp]
```

Replace *fstype* with 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 `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). Multipliers are case insensitive and are multiples of 1024. If no multiplier is specified, it is assumed to be M. 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`.

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`.

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
```

Use the `ncp` option to give NCP users access to the files on the Linux POSIX file system. The NCP share is created at the root of the volume. The NCP volume name is the same as the name of the Linux POSIX volume.

- 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
```

4.8 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: 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)” (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=<fstype>
  device=<devicename>
  size=<value[K|M|G]>
  mp=<full_mount_point_path>
  [mntopt=<fs_mount_options>]
  lvm
  name=<lv_name>
  [group=<vg_name>]
  [-ncp]
```

Replace *fstype* with 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*.

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). Multipliers are case insensitive and are multiples of 1024. If no multiplier is specified, it is assumed to be M.

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 *lvm* option to create an LVM volume group and logical volume.

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

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.

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
```

Use the *ncp* option to give NCP users access to the files on the Linux POSIX file system. The NCP share is created at the root of the volume. The NCP volume name is the same as the name of the Linux POSIX volume.

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
```

- 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
```

4.9 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 5.3.2, “Using NLVM to Create the Resource,” on page 57](#) in [Chapter 5, “Clustering LVM Volume Groups with Novell Cluster Services,” on page 51](#).

4.10 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 3.12, “Mounting a Linux Volume,” on page 30](#).

You can also use the native Linux `mount(8)` command to mount Linux volumes on your OES 11 server 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
```

4.11 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 3.13, "Dismounting a Linux Volume," on page 31](#).

You can use the native Linux `umount (8)` command to dismount Linux volumes on your OES 11 server. 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
```

4.12 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 RAID(s)) 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: 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.

4.13 Deleting a Linux Volume

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

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

- ♦ [Section 4.13.1, “Deleting a Linux POSIX Volume,” on page 48](#)
- ♦ [Section 4.13.2, “Deleting an LVM Volume Group and Logical Volume,” on page 49](#)

4.13.1 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
```

4.13.2 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 5.9, "Deleting a Clustered LVM Volume Group and Logical Volume,"](#) on page 67.

- 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 [--no-prompt] 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
```

4.14 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:

- ♦ “LVM Management Tools” in the *OES 11: Novell Cluster Services 2.0 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)

5 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 Ext2/3, ReiserFS, or XFS.

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

- ♦ [Section 5.1, “Requirements for Creating LVM Cluster Resources,” on page 51](#)
- ♦ [Section 5.2, “Initializing a SAN Device,” on page 53](#)
- ♦ [Section 5.3, “Creating an LVM Cluster Resource,” on page 54](#)
- ♦ [Section 5.4, “Configuring the LVM Cluster Resource Settings,” on page 58](#)
- ♦ [Section 5.5, “Viewing or Modifying the LVM Resource Scripts,” on page 61](#)
- ♦ [Section 5.6, “Sample LVM Resource Scripts,” on page 62](#)
- ♦ [Section 5.7, “Renaming the Mount Point Path for a Clustered LVM Volume,” on page 64](#)
- ♦ [Section 5.8, “Disabling Clustering for an LVM Volume,” on page 65](#)
- ♦ [Section 5.9, “Deleting a Clustered LVM Volume Group and Logical Volume,” on page 67](#)

5.1 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: Novell Cluster Services 2.0 for Linux Administration Guide*.

- ♦ [Section 5.1.1, “Novell Cluster Services,” on page 52](#)
- ♦ [Section 5.1.2, “Linux Logical Volume Manager 2 \(LVM2\),” on page 52](#)
- ♦ [Section 5.1.3, “Clustered Logical Volume Manager Daemon \(CLVMD\),” on page 52](#)
- ♦ [Section 5.1.4, “Resource IP Address,” on page 52](#)
- ♦ [Section 5.1.5, “Shared Storage Devices,” on page 52](#)
- ♦ [Section 5.1.6, “All Nodes Must Be Present,” on page 52](#)
- ♦ [Section 5.1.7, “Mixed Mode OES Clusters \(Not Supported\),” on page 52](#)
- ♦ [Section 5.1.8, “NCP File Access with Novell NCP Server,” on page 53](#)
- ♦ [Section 5.1.9, “CIFS/Samba File Access with Novell Samba,” on page 53](#)

5.1.1 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.

5.1.2 Linux Logical Volume Manager 2 (LVM2)

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

5.1.3 Clustered Logical Volume Manager Daemon (CLVMD)

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

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

5.1.4 Resource IP Address

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

5.1.5 Shared Storage Devices

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

IMPORTANT: Do not mark the device as shareable for clustering because doing so adds a partition to the device, which makes it unavailable to LVM.

5.1.6 All Nodes Must Be Present

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

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

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

5.1.7 Mixed Mode OES Clusters (Not Supported)

LVM volume group cluster resources are not supported in mixed-mode OES clusters.

5.1.8 NCP File Access with Novell NCP Server

Novell NCP Server can be used to provide NCP file access to Linux POSIX file systems on OES 11 servers. Its NCP volumes feature can be used to provide NCP access to files on an LVM volume group cluster resource.

NCP Server must be installed, configured, and running on each node in the cluster. You must modify the Generic File System resource template scripts to add commands that mount and dismount the NCP volume, and create a NCP Virtual Server object for the volume group.

For information about setting up NCP volumes on a resource, see “[Configuring NCP Volumes with Novell Cluster Services](#)” in the *OES 11: NCP Server for Linux Administration Guide*.

5.1.9 CIFS/Samba File Access with Novell Samba

Novell Samba is the Linux Samba software that has been modified to work with Novell eDirectory. It can be used to provide CIFS/Samba access to files on an LVM volume group cluster resource.

Novell Samba must be installed and configured on each node in the cluster. The cluster load script starts the service when you online the Samba cluster resource, and the unload script stops the service when you offline it.

For information about using the Samba resource template to create a Samba cluster resource that is based on an LVM volume group, see “[Configuring Samba for LVM Volume Groups and Novell Cluster Services](#)” in the *OES 11: Novell Samba Administration Guide*.

5.2 Initializing a SAN Device

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 nodes in the cluster.
- 2 Log in to the master node of the cluster as the `root` user, then open a terminal console.
- 3 Launch NSSMU by entering

```
nssmu
```

- 4 In the NSSMU main menu, select *Devices*, then press Enter.
- 5 In the *Devices* list, select the SAN device (such as `sdd`), then view information about it.

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

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

- 6 Press F3 to initialize the selected device.
- 7 Read the advisory that this action destroys all data on the device, then press Y to confirm, or press Esc to cancel 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 Verify that the device is initialized and that it is unshared (that is, *Shareable for Clustering* is set to *No*.)
- 10 Press Esc twice to exit NSSMU.

5.3 Creating an LVM Cluster Resource

This section describes how to use NSSMU to create and cluster enable an LVM volume group.

After you create the resource, you can add lines to its load script, unload script, and monitor script to customize the resource for different services to use, such as Samba or MySQL. In Novell iManager, use the Clusters plug-in to compare the default scripts that are created to the resource template for your product to determine which lines need to be added or modified.

The examples in this section use following setup. Ensure that you replace the sample values with information for your configuration.

Parameter	Sample Value
Device name for the shared device	<div>/dev/sdd</div> <div>The device is initialized and contains no partitions. It is not enabled as shareable for clustering.</div>
Volume group name	<div>vol44</div> <div>By default, NSSMU uses the logical volume name as the LVM volume group name. If you use the <code>NLVM create linux volume</code> command to create the LVM volume group cluster resource, you can specify a different name for the volume group, such as <code>vg44</code>.</div>
Volume name	vol44
Linux POSIX file system type	<div>ext3</div> <div>Valid values are ext2, ext3, reiserfs, and xfs.</div>
Make options for the file system	<div>None (do not specify a value)</div> <div>For a list of the supported options for the file system type you are making, see the <code>mkfs(8)</code> man page and the man page for the specific file system: <code>mkfs.ext2(8)</code>, <code>mkfs.ext3(8)</code>, <code>mkfs.reiserfs(8)</code>, or <code>mkfs.xfs(8)</code>.</div> <div>IMPORTANT: The file system creation fails if you specify a make option that is not supported by the file system type.</div>
Mount options	<div>rw</div> <div>The read and write (rw) options are 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.</div>
Volume size	<div>100 GB</div> <div>A 100 GB LUN is prepared in the shared storage subsystem. It is attached to the nodes in the cluster. The LVM volume group and logical volume use the entire device.</div>

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 use the NLVM <code>create linux volume</code> command to specify a different name for the volume group, such as <code>vg44</code> , the mount device path is <code>/dev/vg44/vol44</code> .
Mount point	/mnt/vol44 You can use any valid Linux path as the mount point. NSSMU automatically creates the mount point path if it does not exist on this node. However, you must manually create the mount point on each of the other nodes in the cluster. If the path does not exist on a node when you fail over the resource to it, the resource goes comatose You can alternatively add the following line to the load script to create the path on a node if it does not exist: <code>ignore_error mkdir -p \$MOUNT_PATH</code>

- ♦ [Section 5.3.1, “Using NSSMU to Create the Resource,” on page 55](#)
- ♦ [Section 5.3.2, “Using NLVM to Create the Resource,” on page 57](#)

5.3.1 Using NSSMU to Create the Resource

Use the procedure in this section to create and cluster enable an LVM volume group cluster resource. The disk will contain one LVM volume group and one logical volume.

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

For information, see [Section 5.2, “Initializing a SAN Device,” on page 53](#).

- 3 Launch NSSMU by entering

```
nssmu
```

- 4 In the NSSMU main menu, select *Linux Volumes*.
- 5 Press *Insert* to begin creating an LVM volume group, then specify the following information:

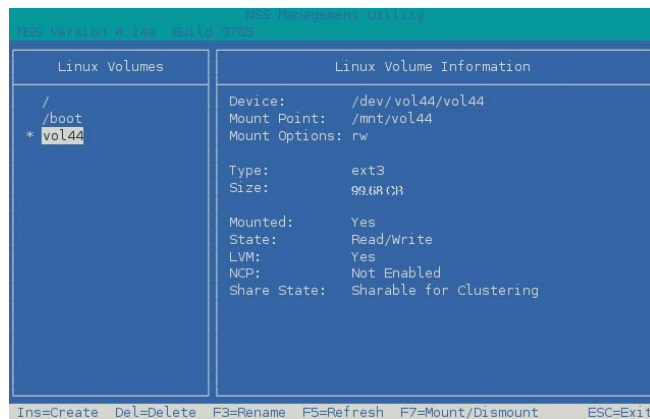
Parameter	Action
LVM type	Select <i>Cluster Enabled LVM2 Volume</i> .

Parameter	Action
NCP enable volume	<p>Type Y (Yes) to enable NCP for the volume. An NCP share is added at the root of the Linux POSIX volume. The NCP share and volume name is the same as for the Linux POSIX volume. The default mount point is <code>/usr/novell/<volume name></code>.</p> <p>Type N (No) to not allow NCP access for the root of the volume at this time.</p> <p>After the Linux POSIX volume is created, you can enable NCP at any time by creating NCP shares on the root of the volume or on subdirectories. The share names can be different than the name of the Linux POSIX volume. In Novell Remote Manager, select <i>Manage NCP Services > Manage Shares</i>, then click <i>Create New Share</i>. For more information, see “Managing NCP Volumes” in the <i>OES 11: NCP Server for Linux Administration Guide</i>.</p>
Volume name	Type the name for the LVM logical volume (such as <code>vol44</code>), then press Enter.
Volume IP address	Type the resource IP address for the LVM logical volume (such as <code>10.10.10.44</code>), then press Enter.
Volume file system type	<p>Select one of the following the Linux POSIX file systems, then press Enter:</p> <ul style="list-style-type: none"> ◆ <code>ext3</code> ◆ <code>resiserfs</code> ◆ <code>ext2</code> ◆ <code>xfs</code>
Mount point path	Type the full mount point path for the LVM logical volume (such as <code>/mnt/vol44</code>), then press Enter.
Make options	Press Enter. Specifying a make option is optional.
Mount options	Press Enter to accept the default read and write options (<code>rw</code>). You can add other mount options.
Device	Select the device that you initialized, such as <code>sdd</code> , then press Insert or Enter.

6 Press F3 to accept the setup you entered for the volume group cluster resource.

The resource is created and brought online on the node where you created it. The resource is named `<volume_group_name>_reference`. In the example, the name of the volume group is the same as the logical volume, so the reference name is `vol44_reference`.

- 7 In the *Linux Volumes* list, select the newly created volume and view information about it.
For information, see [Section 3.4, “Viewing Device Details,”](#) on page 18.



- 8 Press Esc twice to exit NSSMU.
- 9 Continue with [“Configuring the LVM Cluster Resource Settings”](#) on page 58.

5.3.2 Using NLVM to Create the Resource

You can alternatively use the Novell Linux Volume Manager `create linux volume` command with the `lvm` and `shared` options to automatically create and cluster enable an LVM volume group. The NLVM command allows you to specify a group name that is different than the volume name. For information about the options, see [“Create Linux Volume”](#) in the *OES 11: NLVM Reference*.

- 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 to ensure that it is available and contains no partitions.
For information, see [Section 5.2, “Initializing a SAN Device,”](#) on page 53.
- 3 List the devices by entering the following command, then verify that the device that you want to use has been formatted.

```
nlvm list devices
```

Note the device size. The LVM volume group will use the entire device for the clustered volume group. The device should be unshared and contain no partitions.

- 4 At the console prompt, enter the following NLVM command:

```
nlvm create linux volume
  type=<fstype>
  device=<devicename>
  mp=<full_mount_point_path>
  [mntopt=<fs_mount_options>]
  lvm
  name=<lv_name>
  [group=<vg_name>]
  shared
  ip=<resource_ip_address>
```

Replace *fstype* with 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`.

No size is specified. LVM automatically uses the entire device for the clustered volume group.

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 LVM to create an LVM volume group and logical volume.

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

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.

Specify the `shared` option to create a clustered general file system cluster resource using Clustered LVM (cLVM).

Replace *resource_ip_address* with a unique IP address to use for the LVM volume group cluster resource. Specify the address in IPv4 format. It must be in the same subnet as the cluster's IP address.

For example, to create a shared resource where the volume group assumes the name of the logical volume, enter

```
nlvm create linux volume type=ext3 device=sdd size=100G mp=/mnt/vol44 mntopt=rw
lvm name=vol44 shared ip=10.10.10.44
```

If the command is successful, the response is

```
Linux clustered volume vol44 created.
```

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

```
cluster status
```

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

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

- 6 Continue with [“Configuring the LVM Cluster Resource Settings”](#) on page 58.

5.4 Configuring the LVM Cluster Resource Settings

You can use Novell iManager to verify that the LVM volume group cluster resource was created and is online. Customize the resource policies, monitoring, and preferred nodes settings.

- 1 Open Novell iManager in a Web browser, then log in as an administrator user.
- 2 In *Roles and Tasks*, select *Clusters > Cluster Manager*.
- 3 Browse to select the Cluster object of the cluster where you created the volume group cluster resource.

- 4 In the list of resources, locate the new resource, such as vol144_resource, and notice the state of the resource. It should be online and running.

The screenshot shows the Cluster Manager interface. At the top, there's a 'Cluster Manager' header with a help icon. Below it, a 'Cluster' dropdown is set to 'clus1.ncs.novell', with a 'Run Report' button. The 'Epoch' is 0. A small icon of a server is shown. Below that, the 'Cluster State' section is active, displaying a table with 2 items. The table has columns: Type, Name, State, Location, Lives, and Up Since. The first row is 'Master_IP_Address_Resource' (Running, avalon, 1, Dec 8, 2011 1:33:11 PM). The second row is 'vol144_resource' (Running, avalon, 1, Dec 8, 2011 2:01:41 PM), which is highlighted with a red border. A 'Close' button is at the bottom.

Type	Name	State	Location	Lives	Up Since
	Master_IP_Address_Resource	Running	avalon	1	Dec 8, 2011 1:33:11 PM
	vol144_resource	Running	avalon	1	Dec 8, 2011 2:01:41 PM

- 5 Click the resource's name link to open its Properties page.
- 6 On the Resource Policies page, view and modify the resource's Policy settings if needed. Click *Apply* if you make changes.

The screenshot shows the 'Cluster Resource Properties: vol144_resource.clus1.ncs.novell' page. It has tabs for 'Policies', 'Monitoring', 'Preferred Nodes', 'Scripts', and 'Business Continuity'. The 'Policies' tab is active, showing settings for 'Resource Behavior', 'Failover Mode', 'Start Mode', and 'Failback Mode'. 'Resource Behavior' has 'Resource Follows Master' and 'Ignore Quorum' checkboxes. 'Failover Mode' has 'Auto' (selected) and 'Manual' radio buttons. 'Start Mode' has 'Auto' (selected) and 'Manual' radio buttons. 'Failback Mode' has 'Auto', 'Disable' (selected), and 'Manual' radio buttons. 'OK', 'Cancel', and 'Apply' buttons are at the bottom.

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

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 in the event of 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 Click the *Monitoring* tab, enable and configure monitoring for the resource, then click *Apply*.

For information, see “[Enabling Monitoring and Configuring the Monitor Script](#)” in the *OES 11: Novell Cluster Services 2.0 for Linux Administration Guide*.

Clusters > Cluster Manager

Cluster Resource Properties: vol44_resource.clus1.ncs.novell

Policies **Monitoring** Preferred Nodes Scripts Business Continuity

To monitor the health of this resource, enable Resource monitoring. You can set the interval to poll the resource's health, and set an action if it fails to successfully load on the maximum number of local restarts.

☒ Enable Resource Monitoring

Polling Interval: 0 Minutes

Failure Rate

Maximum Local Failures: 0

Time Interval: 0 Minutes
(Example: 3 Failures in 5 minutes)

Failure Action

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

☒ Set Resource as Comatose

☐ Migrate the Resource based on the Preferred Nodes List

☐ Reboot the Hosting Node without Syncing or Unmounting Disks

OK Cancel Apply

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

Clusters > Cluster Manager

Cluster Resource Properties: vol44_resource.clus1.ncs.novell

Policies **Monitoring** **Preferred Nodes** Scripts Business Continuity

View or change the preferred nodes for this cluster resource.

Assigned:

- avalon
- london
- newcastle
- cardiff

Unassigned:

OK Cancel Apply

- 9 At the bottom of the page, click *OK* to close the Properties page and save your changes.
The changes do not take effect until the resource is taken offline and brought online again.
- 10 If you modified the settings on any of the pages, you must take the resource offline and bring it online in order for the changes to take effect:
 - 10a In *Roles and Tasks*, select *Clusters > Cluster Manager*.
 - 10b Select the check box next to the resource, then click *Offline*. Wait for the status to report that it is offline, then continue.
 - 10c Select the check box next to the resource, then click *Online*.
- 11 Continue with [“Viewing or Modifying the LVM Resource Scripts” on page 61](#).

5.5 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 > Cluster Manager*.
- 2 Browse to select the Cluster object of the cluster where you created the volume group cluster resource.
- 3 From the list of Cluster objects, click the name link to open the Properties page, then click the *Scripts* tab.
The *Scripts* tab opens to the load script.
- 4 On the *Load Script* page, view or modify the load script. Click *Apply* if you make changes.
See the [“Sample LVM Resource Load Script” on page 62](#).
- 5 Click the *Unload Script* link to view or modify the unload script. Click *Apply* if you make changes.
See the [“Sample LVM Resource Unload Script” on page 63](#).
- 6 Click the *Monitor Script* link to view or modify the monitor script. Click *Apply* if you make changes.
See the [“Sample LVM Resource Monitor Script” on page 63](#).
- 7 At the bottom of the page, click *OK* to close the Properties page and save your changes.
The changes do not take effect until the resource is taken offline and brought online again.
- 8 If you modified the settings on any of the pages, you must take the resource offline and bring it online in order for the changes to take effect:
 - 8a In *Roles and Tasks*, select *Clusters > Cluster Manager*.
 - 8b Select the check box next to the resource, then click *Offline*. Wait for the status to report that it is offline, then continue.
 - 8c Select the check box next to the resource, then click *Online*.
 - 8d Verify that the resource comes online and reports a *Running* state.
If the resource goes into a *Comatose* state, it is probably because you made a mistake in the lines you added or modified in the scripts. Take the resource offline, go back and correct the scripts, then try to bring it online again.

5.6 Sample LVM Resource Scripts

This section contains sample scripts for the LVM resource. The sample values are the ones used to create the generic LVM cluster resource in [Section 5.3, “Creating an LVM Cluster Resource,” on page 54](#). Sample values and added lines appear in bold font.

- [Section 5.6.1, “Sample LVM Resource Load Script,” on page 62](#)
- [Section 5.6.2, “Sample LVM Resource Unload Script,” on page 63](#)
- [Section 5.6.3, “Sample LVM Resource Monitor Script,” on page 63](#)

5.6.1 Sample LVM Resource Load Script

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

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

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

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

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

# add the IP address
exit_on_error add_secondary_ipaddress $RESOURCE_IP

exit 0
```

5.6.2 Sample LVM Resource Unload Script

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

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

# del the IP address
ignore_error del_secondary_ipaddress $RESOURCE_IP

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

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

exit 0
```

5.6.3 Sample LVM Resource Monitor Script

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

# define the IP address
RESOURCE_IP=10.10.10.44
# define the file system type
MOUNT_FS=ext3
# define the volume group name
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
```

5.7 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 an administrator user.
- 2 In *Roles and Tasks*, select *Clusters > Cluster Manager*.
- 3 Browse to select the Cluster object of the cluster where you created the volume group cluster resource.
- 4 Select the check box next to the resource, then click *Offline*. Wait for the status to report that it is offline, then continue.
- 5 Modify the mount point path value in the load, unload, and monitor scripts for the LVM volume group cluster resource:

- 5a On the Cluster Manager page, select the resource's name link to open its Cluster Properties page, then click the *Scripts* tab.

The *Scripts* tab automatically displays the load script.

- 5b Modify the load script:

- 5b1 In the load script, type the new value for the mount point in the MOUNT_POINT variable:

```
MOUNT_POINT=/media/ext3/vol144
```

- 5b2 Ensure that the following command is added above the mount command line in the load script in order to create the path on nodes if it does not exist.

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

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

- 5b3 Click *Apply*.

- 5c Click the *Unload Script* link, type the new value for the mount point in the MOUNT_POINT variable, then click *Apply*.

```
MOUNT_POINT=/media/ext3/vol144
```

- 5d Click the *Monitor Script* link, type the new value for the mount point in the MOUNT_POINT variable, then click *Apply*.

```
MOUNT_POINT=/media/ext3/vol144
```

- 5e At the bottom of the page, click *OK* to close the Properties page and save your changes.

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

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

- 6a In *Roles and Tasks*, select *Clusters > Cluster Manager*.

- 6b Select the check box next to the resource, then click *Online*.

6c Verify that the resource comes online and reports a *Running* state.

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

7 In NSSMU, verify that the new mount point is used when the clustered LVM volume resource is brought online:

7a Log in as the `root` user to the node that is hosting the resource, then start NSSMU by entering:

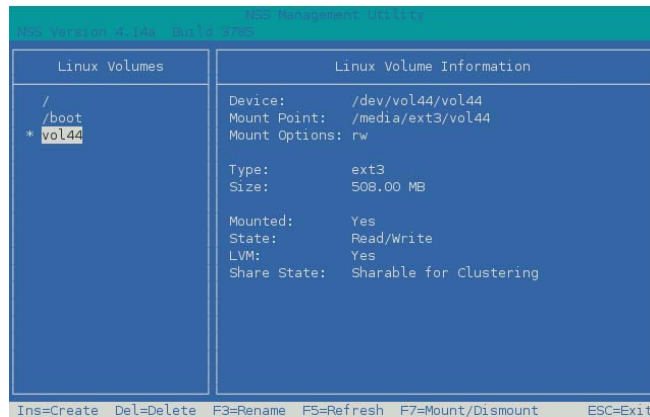
```
nssmu
```

7b From the NSSMU main menu, select *Linux Volumes*, then press Enter.

7c In the *Linux Volumes* list, select the clustered LVM volume.

7d View the volume details to verify that the mount point has changed.

For example, the mount point is `/media/ext3/vol44`.



7e Press Esc twice to exit NSSMU.

5.8 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 Log in to Novell iManager as a cluster administrator user.

1b Select *Clusters > Cluster Manager*, then browse to select the Cluster object.

1c On the Cluster Manager page, select the check box next to the LVM volume group cluster resource (such as `vol44_resource`), then click *Offline*. Wait until the resource is offline to continue.

1d Select *Clusters > Cluster Options*.

1e On the Cluster Options page, select the check box next to the LVM volume group cluster resource, click *Delete*, then click *OK* to confirm.

1f Exit iManager.

2 Log in as the `root` user on the server where the resource was online, then open a terminal console.

- 3 Remove the shareable state from the clustered LVM volume group 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

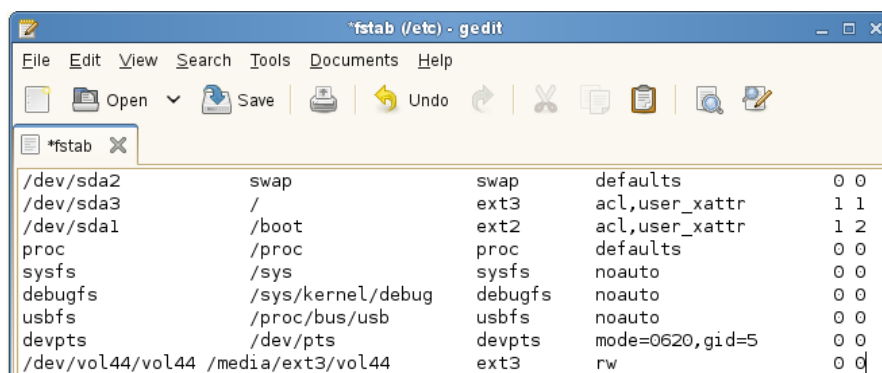
```
mount -t ext3 /dev/vol44/vol44 /media/ext3/vol44
```

- 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.

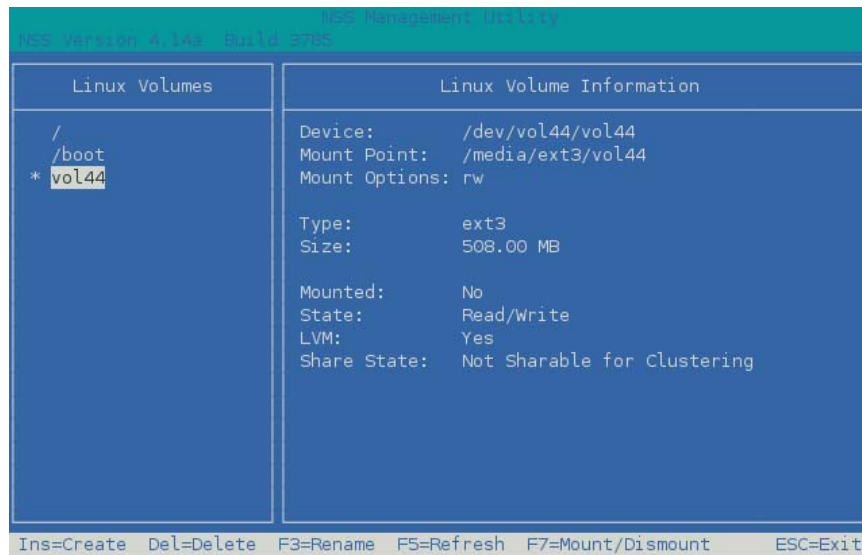
For example, complete the line for the volume's device path, such as:

```
/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.



5.9 Deleting a Clustered LVM Volume Group and Logical Volume

Before you can delete a clustered LVM volume group, you must take the volume group cluster resource offline.

- 1 Log in to the server as the root user, then open a terminal console.
- 2 Take the resource offline by entering

```
cluster offline <resource_name>
```

For example, enter

```
cluster offline vol44_resource
```

- 3 Delete the LVM logical volume, the clustered LVM volume group, and the Cluster Resource objects, then re-initialize the device:
 - 3a In the terminal console, launch NSSMU by entering


```
nssmu
```
 - 3b In the NSSMU main menu, select *Linux Volumes*, then press Enter.
 - 3c In the *Linux Volumes* list, select the clustered linux volume, then press Delete.

Deleting the volume and volume group sets the device in an uninitialized state.
 - 3d When you are prompted to confirm the delete action, read the warning message, then press Y (yes) to continue, or press N (no) to cancel the delete action.
 - 3e Press Esc to return to the main menu, then select *Devices* and press Enter.
 - 3f In the *Devices* list, select the device, then press F3 to initialize it.
 - 3g When you are prompted to confirm the initialization action, read the warning message, then press Y (yes) to continue, or press N (no) to cancel the action.
 - 3h Press Esc twice to exit NSSMU.

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

```
vgdisplay
```

6 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 “[nbackup\(1\)](#)” in the *OES 11: 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: 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, “July 25, 2012,” on page 71](#)
- ♦ [Section A.2, “May 31, 2012,” on page 72](#)
- ♦ [Section A.3, “April 30, 2012,” on page 73](#)
- ♦ [Section A.4, “January 18, 2012,” on page 74](#)

A.1 July 25, 2012

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

- ♦ [Section A.1.1, “Backing Up Linux POSIX Volumes,” on page 71](#)
- ♦ [Section A.1.2, “NLVM Commands,” on page 71](#)

A.1.1 Backing Up Linux POSIX Volumes

Location	Change
Chapter 6, “Backing Up Linux POSIX Volumes,” on page 69	This section is new.

A.1.2 NLVM Commands

Location	Change
Chapter 4, “Managing Linux POSIX Volumes with NLVM Commands,” on page 35	Corrections for the <code>mntopt</code> option.

A.2 May 31, 2012

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

- [Section A.2.1, “Clustering LVM Volume Groups with Novell Cluster Services,” on page 72](#)
- [Section A.2.2, “Managing Linux POSIX Volumes with NLVM Commands,” on page 72](#)
- [Section A.2.3, “Managing Linux POSIX Volumes with NSSMU,” on page 72](#)
- [Section A.2.4, “What’s New or Changed for Linux POSIX Volumes,” on page 73](#)

A.2.1 Clustering LVM Volume Groups with Novell Cluster Services

Location	Change
Section 5.2, “Initializing a SAN Device,” on page 53	You are automatically prompted to confirm the initialization action.
Section 5.9, “Deleting a Clustered LVM Volume Group and Logical Volume,” on page 67	You are automatically prompted to confirm the delete action.

A.2.2 Managing Linux POSIX Volumes with NLVM Commands

Location	Change
Section 4.4, “Initializing a Device,” on page 40	You are automatically prompted to confirm the initialization action. Use the --no-prompt NLVM option to suppress the confirmation.
Section 4.13, “Deleting a Linux Volume,” on page 48	You are automatically prompted to confirm the delete action. Use the --no-prompt NLVM option to suppress the confirmation.

A.2.3 Managing Linux POSIX Volumes with NSSMU

Location	Change
Section 3.5, “Initializing a Device,” on page 20	You are automatically prompted to confirm the initialization action.
Section 3.15, “Deleting a Linux POSIX Volume,” on page 33	You are automatically prompted to confirm the delete action.

A.2.4 What's New or Changed for Linux POSIX Volumes

Location	Change
Section 2.2, "What's New (April 2012 Patches)," on page 13	This section is new.

A.3 April 30, 2012

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

- ♦ [Section A.3.1, "Clustering LVM Volume Groups with Novell Cluster Services," on page 73](#)
- ♦ [Section A.3.2, "Managing Linux POSIX Volumes with NLVM Commands," on page 73](#)
- ♦ [Section A.3.3, "Managing Linux POSIX Volumes with NSSMU," on page 73](#)
- ♦ [Section A.3.4, "What's New or Changed for Linux POSIX Volumes," on page 74](#)

A.3.1 Clustering LVM Volume Groups with Novell Cluster Services

Location	Change
Section 5.3, "Creating an LVM Cluster Resource," on page 54	Modified to add information about the NCP option.

A.3.2 Managing Linux POSIX Volumes with NLVM Commands

Location	Change
Section 4.7, "Creating a Linux POSIX Volume," on page 42	Modified to add information about the NCP option.

A.3.3 Managing Linux POSIX Volumes with NSSMU

Location	Change
Section 3.9, "Creating a Linux POSIX Volume," on page 25	Modified to add information about the NCP option.
Section 3.8, "Viewing Details for a Linux Volume," on page 22	

A.3.4 What's New or Changed for Linux POSIX Volumes

Location	Change
Section 2.2, "What's New (April 2012 Patches)," on page 13	This section is new.

A.4 January 18, 2012

The document format was updated to reflect newly revised corporate standards.