

# Open Enterprise Server 2023

## Linux Volume Manager Reference

October 2022

## **Legal Notices**

**© Copyright 2022 Micro Focus or one of its affiliates.**

The only warranties for products and services of Micro Focus and its affiliates and licensors (“Micro Focus”) are as may be set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Micro Focus shall not be liable for technical or editorial errors or omissions contained herein. The information contained herein is subject to change without notice.

---

# Contents

|  |           |
|--|-----------|
| <b>About This Guide</b>  | <b>7</b>  |
| <b>1 Overview of NLVM</b>  | <b>9</b>  |
| <b>2 What's New or Changed in Novell Linux Volume Manager</b>              | <b>11</b> |
| 2.1 What's New or Changed (OES 2023) .....                                 | 11        |
| <b>3 Installing or Upgrading NLVM</b>                                      | <b>13</b> |
| <b>4 Using NLVM in a Virtualized Environment</b>                           | <b>15</b> |
| <b>5 Planning for NLVM</b>   | <b>17</b> |
| 5.1 Root User .....  | 17        |
| 5.2 Naming Conventions for Storage Objects .....                           | 17        |
| 5.2.1 NSS Pool and Volume Names .....                                      | 17        |
| 5.2.2 NSS Pool Snapshot Names .....  | 18        |
| 5.2.3 NSS Software RAID Names .....  | 18        |
| 5.2.4 NCP Volume Names .....   | 19        |
| 5.2.5 Linux LVM Volume Group and Logical Volume Names .....                | 19        |
| 5.3 NSS Pools on the System Device .....                                   | 20        |
| 5.4 NSS Pools Created on NetWare Servers .....                             | 20        |
| 5.5 NSS Pools Created on Older Versions of OES Servers .....               | 20        |
| 5.6 Linux LVM Volume Group .....   | 20        |
| 5.7 Linux LVM Volume Group Cluster Resources .....                         | 21        |
| 5.8 Using NLVM with NSS Software RAID5 .....                               | 21        |
| 5.9 Using NLVM with Linux Software RAID5 .....                             | 21        |
| 5.9.1 Linux Software RAID5 .....   | 21        |
| 5.9.2 Linux Software RAID5 Are Not Cluster Aware .....                     | 22        |
| 5.9.3 Linux Software RAID5 Are Not Recommended for the System Device ..... | 22        |
| 5.10 Using iSCSI Devices with NSS Software RAID5 .....                     | 22        |
| 5.11 Using Antivirus Software with NCP Volumes .....                       | 22        |
| <b>6 NLVM Commands</b>   | <b>23</b> |
| 6.1 Syntax Overview .....  | 27        |
| 6.1.1 Syntax .....   | 27        |
| 6.1.2 Syntax Conventions .....   | 27        |
| 6.1.3 Documentation Conventions .....                                      | 28        |
| 6.1.4 Files .....  | 28        |
| 6.2 NLVM Options .....   | 29        |
| 6.3 Common Options .....   | 30        |
| 6.4 Complete Move .....  | 31        |
| 6.5 Create Linux Volume .....  | 31        |

|      |                     |     |
|------|---------------------|-----|
| 6.6  | Create Partition    | 36  |
| 6.7  | Create Pool         | 38  |
| 6.8  | Create RAID         | 42  |
| 6.9  | Create Snap         | 45  |
| 6.10 | Create Volume       | 46  |
| 6.11 | Delete Linux Volume | 48  |
| 6.12 | Delete Move         | 49  |
| 6.13 | Delete Partition    | 50  |
| 6.14 | Delete Pool         | 51  |
| 6.15 | Delete RAID         | 51  |
| 6.16 | Delete RAID Segment | 52  |
| 6.17 | Delete Snap         | 53  |
| 6.18 | Delete Volume       | 54  |
| 6.19 | Expand Partition    | 54  |
| 6.20 | Expand Pool         | 55  |
| 6.21 | Expand RAID         | 56  |
| 6.22 | Fstrim Pool         | 57  |
| 6.23 | Fstrim Status       | 57  |
| 6.24 | Init Device         | 58  |
| 6.25 | Label               | 59  |
| 6.26 | Linux Mount         | 60  |
| 6.27 | Linux Unmount       | 61  |
| 6.28 | List Device         | 62  |
| 6.29 | List Devices        | 64  |
| 6.30 | List Linux Volume   | 68  |
| 6.31 | List Linux Volumes  | 70  |
| 6.32 | List Move           | 72  |
| 6.33 | List Moves          | 74  |
| 6.34 | List Partition      | 76  |
| 6.35 | List Partitions     | 78  |
| 6.36 | List Pool           | 84  |
| 6.37 | List Pools          | 87  |
| 6.38 | List Snap           | 91  |
| 6.39 | List Snaps          | 92  |
| 6.40 | List Volume         | 94  |
| 6.41 | List Volumes        | 96  |
| 6.42 | Mount               | 99  |
| 6.43 | Move                | 99  |
| 6.44 | Pause Move          | 101 |
| 6.45 | Pool Activate       | 101 |
| 6.46 | Pool Deactivate     | 102 |
| 6.47 | RAID                | 102 |
| 6.48 | Rename Pool         | 105 |
| 6.49 | Rename RAID         | 105 |
| 6.50 | Rename Volume       | 106 |
| 6.51 | Rescan              | 107 |
| 6.52 | Resume Move         | 107 |
| 6.53 | Share               | 108 |
| 6.54 | Unmount             | 108 |

|           |  |            |
|-----------|--|------------|
| 6.55      | Unshare .....  | 109        |
| 6.56      | Volume Mount .....   | 109        |
| 6.57      | Volume Unmount .....   | 110        |
| <b>7</b>  | <b>NLVM Examples for the NSS File System</b>                               | <b>111</b> |
| 7.1       | Creating an NSS Pool and Volume .....                                      | 111        |
| 7.2       | Mirroring a Pool Partition .....   | 111        |
| 7.3       | Recovering a Mirror where All Elements Report 'Not in Sync' .....          | 112        |
| 7.4       | Logging Out of an iSCSI Device that Contains an NSS Pool .....             | 113        |
| 7.5       | Creating a Linux Volume on a Device that Contains a Novell Partition ..... | 114        |
| <b>8</b>  | <b>NLVM Examples for Clustering with OES Cluster Services</b>              | <b>117</b> |
| 8.1       | Creating or Mirroring an SBD Partition .....                               | 117        |
| 8.1.1     | Requirements and Guidelines for Creating an SBD Partition .....            | 118        |
| 8.1.2     | Creating a Non-Mirrored SBD Partition with NLVM .....                      | 120        |
| 8.1.3     | Mirroring an Existing SBD Partition with NLVM .....                        | 124        |
| 8.1.4     | Creating a Mirrored SBD Partition with NLVM .....                          | 126        |
| 8.2       | Unmirroring a Mirrored SBD Partition with NLVM .....                       | 129        |
| 8.3       | Deleting an SBD Partition with NLVM .....                                  | 130        |
| <b>9</b>  | <b>Troubleshooting NLVM</b>  | <b>133</b> |
| 9.1       | Viewing Error Code Messages .....  | 133        |
| 9.2       | Failure to Create an LVM Volume Group .....                                | 133        |
| 9.3       | Failure to Create a Clustered LVM Volume Group .....                       | 134        |
| 9.4       | Device Is Not Available for Use in an LVM Volume Group .....               | 134        |
| 9.5       | NLVM Pool Move Fails and Deactivates the Pool .....                        | 134        |
| 9.6       | Error 20897 - This node is not a cluster member .....                      | 134        |
| 9.7       | NLVM Error Codes .....   | 135        |
| 9.7.1     | NLVM Error List .....  | 135        |
| 9.7.2     | NLVM Error Descriptions .....  | 138        |
| 9.8       | NSS Error Codes .....  | 144        |
| <b>10</b> | <b>Security Considerations</b>   | <b>145</b> |
| 10.1      | Root User Privileges .....   | 145        |
| 10.2      | Files .....  | 145        |
| <b>A</b>  | <b>Configuring Settings for the NLVM Library</b>                           | <b>147</b> |



# About This Guide

The Linux Volume Manager (NLVM) allows you to use NetWare partitions on an Open Enterprise Server (OES) 2023 server. This guide describes NLVM and how to use it with OES Storage Services (NSS) file systems, Linux POSIX file systems, and OES Cluster Services.

- ♦ Chapter 1, “Overview of NLVM,” on page 9
- ♦ Chapter 2, “What’s New or Changed in Novell Linux Volume Manager,” on page 11
- ♦ Chapter 3, “Installing or Upgrading NLVM,” on page 13
- ♦ Chapter 4, “Using NLVM in a Virtualized Environment,” on page 15
- ♦ Chapter 5, “Planning for NLVM,” on page 17
- ♦ Chapter 6, “NLVM Commands,” on page 23
- ♦ Chapter 7, “NLVM Examples for the NSS File System,” on page 111
- ♦ Chapter 8, “NLVM Examples for Clustering with OES Cluster Services,” on page 117
- ♦ Chapter 9, “Troubleshooting NLVM,” on page 133
- ♦ Chapter 10, “Security Considerations,” on page 145
- ♦ Appendix A, “Configuring Settings for the NLVM Library,” on page 147

## Audience

This guide is intended for storage and cluster 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: NLVM Reference*, visit the [OES 2023 website \(https://www.microfocus.com/documentation/open-enterprise-server/2023/\)](https://www.microfocus.com/documentation/open-enterprise-server/2023/).

## Additional Documentation

For documentation on OES, see the [OES 2023 Documentation website \(https://www.microfocus.com/documentation/open-enterprise-server/2023/\)](https://www.microfocus.com/documentation/open-enterprise-server/2023/).





# 1 Overview of NLVM

The Novell Linux Volume Manager (NLVM) provides management of OES Storage Services (NSS) storage objects in Open Enterprise Server (OES) server. The command line interface (CLI) commands can be used in a Linux console or in a script. The NSS management tools use the NLVM library of APIs to create and manage NSS storage objects. NLVM also provides options to create Linux POSIX file systems, such as Btrfs, Ext2, Ext3, ReiserFS, and XFS.

This command reference describes how to use command line commands to manage the following storage objects:

- ◆ [Devices and Partitions](#)
- ◆ [Linux POSIX Volumes](#)
- ◆ [NSS Pools](#)
- ◆ [NSS Pool Snapshots](#)
- ◆ [NSS Software RAIDs](#)
- ◆ [NSS Volumes](#)



# 2 What's New or Changed in Novell Linux Volume Manager

This section describes the changes made to Novell Linux Volume Manager (NLVM) since the Open Enterprise Server (OES) 2023 release.

- ♦ [Section 2.1, “What’s New or Changed \(OES 2023\),” on page 11](#)

## 2.1 What's New or Changed (OES 2023)

NLVM in OES 2023 has been modified for bug fixes. There are no new features or enhancements in OES 2023.



# 3 Installing or Upgrading NLVM

The Novell Linux Volume Manager command line tool and libraries are installed and upgraded by default whenever you install or upgrade OES Storage Services (NSS) on your Open Enterprise Server (OES) 11 or later server. No action is required.

For information about installing NSS on your OES 11 or later server, see “[Installing and Configuring OES Storage Services](#)” in the *OES 2023: NSS File System Administration Guide for Linux*.

For general information about installing, upgrading, and patching OES Services on your OES 11 or later server, see the *OES 2023: Installation Guide*.



# 4 Using NLVM in a Virtualized Environment

The Novell Linux Volume Manager (NLVM) utility runs in a virtualized environment just as it does on a physical server running Open Enterprise Server 11 and later, and requires no special configuration or other changes.

- ♦ For information on setting up virtualized OES, see “[Installing OES on a VM](#)” in the *OES 2023: Installation Guide*.
- ♦ To get started with Xen virtualization, see the [Virtualization Guide \(https://documentation.suse.com/sles/15-SP4/html/SLES-all/book-virtualization.html\)](https://documentation.suse.com/sles/15-SP4/html/SLES-all/book-virtualization.html).
- ♦ To get started with KVM virtualization, see the [Virtualization Guide \(https://documentation.suse.com/sles/15-SP4/html/SLES-all/book-virtualization.html\)](https://documentation.suse.com/sles/15-SP4/html/SLES-all/book-virtualization.html).
- ♦ To get started with third-party virtualization platforms, such as Hyper-V from Microsoft and the different VMware product offerings, refer to the documentation for the product you are using.

For information about using the Novell Linux Volume Manager for OES Storage Services (NSS) volumes in a virtualized environment with Open Enterprise Server 11 and later, refer to the guidelines and requirements in “[Using NSS in a Virtualization Environment](#)” in the *OES 2023: NSS File System Administration Guide for Linux*.





# 5 Planning for NLVM

Consider the requirements and caveats in this section when planning to use Novell Linux Volume Manager (NLVM) command line commands on Open Enterprise Server (OES) servers.

- ♦ [Section 5.1, “Root User,” on page 17](#)
- ♦ [Section 5.2, “Naming Conventions for Storage Objects,” on page 17](#)
- ♦ [Section 5.3, “NSS Pools on the System Device,” on page 20](#)
- ♦ [Section 5.4, “NSS Pools Created on NetWare Servers,” on page 20](#)
- ♦ [Section 5.5, “NSS Pools Created on Older Versions of OES Servers,” on page 20](#)
- ♦ [Section 5.6, “Linux LVM Volume Group,” on page 20](#)
- ♦ [Section 5.7, “Linux LVM Volume Group Cluster Resources,” on page 21](#)
- ♦ [Section 5.8, “Using NLVM with NSS Software RAID5,” on page 21](#)
- ♦ [Section 5.9, “Using NLVM with Linux Software RAID5,” on page 21](#)
- ♦ [Section 5.10, “Using iSCSI Devices with NSS Software RAID5,” on page 22](#)
- ♦ [Section 5.11, “Using Antivirus Software with NCP Volumes,” on page 22](#)

## 5.1 Root User

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

## 5.2 Naming Conventions for Storage Objects

Consider the naming conventions in this section when you create or rename storage objects with NLVM.

- ♦ [Section 5.2.1, “NSS Pool and Volume Names,” on page 17](#)
- ♦ [Section 5.2.2, “NSS Pool Snapshot Names,” on page 18](#)
- ♦ [Section 5.2.3, “NSS Software RAID Names,” on page 18](#)
- ♦ [Section 5.2.4, “NCP Volume Names,” on page 19](#)
- ♦ [Section 5.2.5, “Linux LVM Volume Group and Logical Volume Names,” on page 19](#)

### 5.2.1 NSS Pool and Volume Names

OES Storage Services (NSS) pool names and volume names must be unique from other pools and volumes on the server. In a cluster, the names of shared pools and volumes must be unique across all nodes in the cluster.

Pool and volume names can be 2 to 15 characters.

Uppercase letters A to Z, number characters 0 to 9, and underscore ( `_` ) are valid characters for all pools and volumes. Names cannot start or end in an underscore, and cannot contain double underscores. When you create an NSS pool or volume, the name you specify is automatically converted to uppercase.

If the pool is not shared, the pool name or volume name can also contain special characters:

`!@#$%&()`

Names that contain special characters must be enclosed in quotation marks in all commands and scripts.

The names cannot be reserved names such as `con`, `com`, `lpt`, `pipe`, `all`, and so on.

## 5.2.2 NSS Pool Snapshot Names

An NSS pool snapshot name must be a unique snap name on the server.

Pool snapshot names are 2 to 15 characters.

The naming conventions for a pool snapshot are the same as for NSS pools and volumes. When you create an NSS pool snapshot, the name you specify is automatically converted to uppercase.

## 5.2.3 NSS Software RAID Names

An NSS software RAID name must be unique from other devices on the server. In a cluster, the names of shared software RAID names must be unique across all nodes in the cluster.

RAID names are 2 to 58 characters.

Names are preferred to use characters A to Z, a to z, 0 to 9, and underscore ( `_` ). Names cannot start or end in underscore, and cannot contain double underscores. Printable ASCII characters (see decimal codes 33 to 122 in a code chart) are valid. The name is case sensitive; it can contain uppercase and lowercase characters.

RAID names can contain special characters such as:

`!@#$%&()`

Names that contain special characters must be enclosed in quotation marks in all commands and scripts. On the BASH command line, each special character must be escaped by preceding it with a backslash character ( `\` ).

The RAID names cannot be reserved names such as `con`, `com`, `lpt`, `pipe`, `all`, and so on.

## 5.2.4 NCP Volume Names

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

If you NCP enable a Linux volume as you create it with NSSMU or the `nlvm create linux volume` command, the NCP volume name is based on the specified Linux volume name, but all letters are capitalized. Ensure that the specified Linux volume name does not exceed 14 characters and does not use special characters. Letters A-Z, letters a-z, numbers 0-9, and underscores are supported.

## 5.2.5 Linux LVM Volume Group and Logical Volume Names

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

- ♦ [“NLVM Requirements for LVM Names” on page 19](#)
- ♦ [“LVM2 Requirements for LVM Names” on page 19](#)
- ♦ [“Clustered LVM Requirements for LVM Names” on page 20](#)

### NLVM Requirements for LVM Names

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

When you create a Linux LVM logical volume without specifying an LVM volume group name, NLVM assigns the volume name to the volume group.

### LVM2 Requirements for LVM Names

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

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

A logical volume cannot be named the following reserved words:

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

The logical volume name also cannot contain the following strings:

```
_mlog
_mimage
```

## Clustered LVM Requirements for LVM Names

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

### 5.3 NSS Pools on the System Device

You can create an NSS pool on the system device where you installed the operating system if there is free space available on the device. This capability is not supported at install time. When you create the pool, select the system device (such as `sda`) and specify the amount of free space to use for the pool.

### 5.4 NSS Pools Created on NetWare Servers

NLVM is compatible with NSS pools that were created on NetWare servers.

For information about relocating a pool from a standalone NetWare server to an OES 2018 or later server, see “[Migrating NSS Devices to OES 2023](#)” in the *OES 2023: NSS File System Administration Guide for Linux*.

For information about cluster migrating a shared pool cluster resource to an OES 2018 or later node during a rolling cluster conversion, see the *OES 2015 SP1: Novell Cluster Services NetWare to Linux Conversion Guide*.

### 5.5 NSS Pools Created on Older Versions of OES Servers

NLVM is compatible with NSS pools that were created on older versions of OES servers.

For information about relocating a pool from NetWare to an OES 2018 or later server, see “[Migrating NSS Devices to OES 2023](#)” in the *OES 2023: NSS File System Administration Guide for Linux*.

For information about cluster migrating a shared pool cluster resource to an OES 2018 or later node during a rolling cluster upgrade, see “[Upgrading OES Clusters](#)” in the *OES 2023: OES Cluster Services for Linux Administration Guide*.

### 5.6 Linux LVM Volume Group

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

## 5.7 Linux LVM Volume Group Cluster Resources

OES Cluster Services, NLVM, and NSSMU use the Clustered Logical Volume Manager (CLVM) to manage LVM volume group cluster resources. CLVM requires the Linux kernel 2.6.32.45-0.3 or later. You can get the latest kernel version by using the SLES 15 SP4 update channel.

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

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

## 5.8 Using NLVM with NSS Software RAIDs

NSS software RAIDs are supported for use with NSS pools. You can use the `nlvm create raid` command with `type=sbd` to mirror an SBD partition on two shared LUN devices for the OES Cluster Services SBD (split-brain-detector). The `sbd` type for a software RAID1 is also used by the OES Cluster Services SBD Utility (`sbdutil`) to mirror the SBD partition.

---

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

---

## 5.9 Using NLVM with Linux Software RAIDs

Linux Software RAIDs are intended to be used with Linux tools and file systems. Consider the caveats in this section before implementing Linux Software RAIDs on your OES server.

- ♦ [Section 5.9.1, “Linux Software RAIDs,” on page 21](#)
- ♦ [Section 5.9.2, “Linux Software RAIDs Are Not Cluster Aware,” on page 22](#)
- ♦ [Section 5.9.3, “Linux Software RAIDs Are Not Recommended for the System Device,” on page 22](#)

### 5.9.1 Linux Software RAIDs

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

---

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

---

For NSS pools, you can use hardware RAID devices or NSS Software RAID devices to achieve disk fault tolerance.

For Linux POSIX volumes, LVM volume groups, and cLVM volume groups, you can use hardware RAID devices on your storage subsystem to achieve disk fault tolerance.

## 5.9.2 Linux Software RAID5 Are Not Cluster Aware

Do not use Linux Software RAID5 for devices that you plan to use for shared storage objects. Linux Software RAID5 devices do not support concurrent activation on multiple nodes; that is, they are not cluster aware. They cannot be used for shared-disk storage objects, such as the OCFS2 file system, cLVM volume groups, and OES Cluster Services SBD (split-brain-detector) partitions.

For shared disks, you can use hardware RAID devices on your storage subsystem to achieve fault tolerance.

## 5.9.3 Linux Software RAID5 Are Not Recommended for the System Device

We recommend that you do not use Linux software RAID5 (such as MD RAID5 and Device Mapper RAID5) on the system device if you plan to use free space on the device later for storage objects managed by NSS tools. During the SLES and OES installation, if you create a Linux software RAID5 device to use as the system device for the root (/) file system, the free space on the system device cannot be used later for NSS pools because the configuration of NSS storage objects on Linux software RAID5 has not yet been fully tested.

---

**IMPORTANT:** In OES 11, a server hang or crash can occur if you attempt to use a Linux software RAID5 when you create storage objects that are managed by NSS management tools.

---

For the Linux system device, you can use a hardware RAID device to achieve fault tolerance. This allows NSS tools to see and use any available free space on the system device for unshared NSS pools.

## 5.10 Using iSCSI Devices with NSS Software RAID5

Using iSCSI devices on the iSCSI initiator server to create NSS software RAID5 devices can cause poor performance. If you would like RAID5 protection, create the RAID5 on the target server and present that RAID device to the initiator as a single iSCSI device.

## 5.11 Using Antivirus Software with NCP Volumes

For information about using antivirus software with NCP volumes, see “[List of Antivirus Software](#)” in the *OES 2023: Planning and Implementation Guide*.

# 6 NLVM Commands

The Novell Linux Volume Manager (NLVM) command line interface (CLI) for Open Enterprise Server (OES) provides commands that can be used in a Linux console or in a script. The OES Storage Services (NSS) management tools use NLVM to create and manage NSS storage objects. NLVM provides options to create Linux POSIX file systems, such as Btrfs, Ext2, Ext3, ReiserFS, and XFS.

This section describes the syntax and usage for NLVM commands.

- ◆ [General Options](#)
- ◆ [Devices and Partitions](#)
- ◆ [Linux POSIX Volumes](#)
- ◆ [NSS Pools](#)
- ◆ [NSS Pool Snapshots](#)
- ◆ [NSS Software RAIDs](#)
- ◆ [NSS Volumes](#)
- ◆ [All NLVM Commands \(A to Z\)](#)

## General Options

- ◆ [Section 6.1, “Syntax Overview,” on page 27](#)
- ◆ [Section 6.2, “NLVM Options,” on page 29](#)
- ◆ [Section 6.3, “Common Options,” on page 30](#)

## Devices and Partitions

- ◆ [Section 6.6, “Create Partition,” on page 36](#)
- ◆ [Section 6.13, “Delete Partition,” on page 50](#)
- ◆ [Section 6.19, “Expand Partition,” on page 54](#)
- ◆ [Section 6.24, “Init Device,” on page 58](#)
- ◆ [Section 6.25, “Label,” on page 59](#)
- ◆ [Section 6.28, “List Device,” on page 62](#)
- ◆ [Section 6.29, “List Devices,” on page 64](#)
- ◆ [Section 6.34, “List Partition,” on page 76](#)
- ◆ [Section 6.35, “List Partitions,” on page 78](#)
- ◆ [Section 6.47, “RAID,” on page 102](#)
- ◆ [Section 6.51, “Rescan,” on page 107](#)
- ◆ [Section 6.53, “Share,” on page 108](#)
- ◆ [Section 6.55, “Unshare,” on page 109](#)

## Linux POSIX Volumes

- ◆ Section 6.5, “Create Linux Volume,” on page 31
- ◆ Section 6.11, “Delete Linux Volume,” on page 48
- ◆ Section 6.26, “Linux Mount,” on page 60
- ◆ Section 6.27, “Linux Unmount,” on page 61
- ◆ Section 6.30, “List Linux Volume,” on page 68
- ◆ Section 6.31, “List Linux Volumes,” on page 70

## NSS Pools

- ◆ Section 6.4, “Complete Move,” on page 31
- ◆ Section 6.7, “Create Pool,” on page 38
- ◆ Section 6.12, “Delete Move,” on page 49
- ◆ Section 6.14, “Delete Pool,” on page 51
- ◆ Section 6.20, “Expand Pool,” on page 55
- ◆ Section 6.22, “Fstrim Pool,” on page 57
- ◆ Section 6.32, “List Move,” on page 72
- ◆ Section 6.33, “List Moves,” on page 74
- ◆ Section 6.36, “List Pool,” on page 84
- ◆ Section 6.37, “List Pools,” on page 87
- ◆ Section 6.42, “Mount,” on page 99
- ◆ Section 6.43, “Move,” on page 99
- ◆ Section 6.44, “Pause Move,” on page 101
- ◆ Section 6.45, “Pool Activate,” on page 101
- ◆ Section 6.46, “Pool Deactivate,” on page 102
- ◆ Section 6.48, “Rename Pool,” on page 105
- ◆ Section 6.51, “Rescan,” on page 107
- ◆ Section 6.52, “Resume Move,” on page 107
- ◆ Section 6.54, “Unmount,” on page 108

## NSS Pool Snapshots

- ◆ Section 6.9, “Create Snap,” on page 45
- ◆ Section 6.17, “Delete Snap,” on page 53
- ◆ Section 6.38, “List Snap,” on page 91
- ◆ Section 6.39, “List Snaps,” on page 92

## NSS Software RAIDs

- ◆ Section 6.8, “Create RAID,” on page 42
- ◆ Section 6.15, “Delete RAID,” on page 51



- ◆ Section 6.16, “Delete RAID Segment,” on page 52
- ◆ Section 6.21, “Expand RAID,” on page 56
- ◆ Section 6.47, “RAID,” on page 102
- ◆ Section 6.49, “Rename RAID,” on page 105

## **NSS Volumes**

- ◆ Section 6.10, “Create Volume,” on page 46
- ◆ Section 6.18, “Delete Volume,” on page 54
- ◆ Section 6.40, “List Volume,” on page 94
- ◆ Section 6.41, “List Volumes,” on page 96
- ◆ Section 6.50, “Rename Volume,” on page 106
- ◆ Section 6.56, “Volume Mount,” on page 109
- ◆ Section 6.57, “Volume Unmount,” on page 110

## **All NLVM Commands (A to Z)**

- ◆ Section 6.1, “Syntax Overview,” on page 27
- ◆ Section 6.2, “NLVM Options,” on page 29
- ◆ Section 6.3, “Common Options,” on page 30
- ◆ Section 6.4, “Complete Move,” on page 31
- ◆ Section 6.5, “Create Linux Volume,” on page 31
- ◆ Section 6.6, “Create Partition,” on page 36
- ◆ Section 6.7, “Create Pool,” on page 38
- ◆ Section 6.8, “Create RAID,” on page 42
- ◆ Section 6.9, “Create Snap,” on page 45
- ◆ Section 6.10, “Create Volume,” on page 46
- ◆ Section 6.11, “Delete Linux Volume,” on page 48
- ◆ Section 6.12, “Delete Move,” on page 49
- ◆ Section 6.13, “Delete Partition,” on page 50
- ◆ Section 6.14, “Delete Pool,” on page 51
- ◆ Section 6.15, “Delete RAID,” on page 51
- ◆ Section 6.16, “Delete RAID Segment,” on page 52
- ◆ Section 6.17, “Delete Snap,” on page 53
- ◆ Section 6.18, “Delete Volume,” on page 54
- ◆ Section 6.19, “Expand Partition,” on page 54
- ◆ Section 6.20, “Expand Pool,” on page 55
- ◆ Section 6.21, “Expand RAID,” on page 56
- ◆ Section 6.22, “Fstrim Pool,” on page 57
- ◆ Section 6.23, “Fstrim Status,” on page 57

- ◆ Section 6.24, “Init Device,” on page 58
- ◆ Section 6.25, “Label,” on page 59
- ◆ Section 6.26, “Linux Mount,” on page 60
- ◆ Section 6.27, “Linux Unmount,” on page 61
- ◆ Section 6.28, “List Device,” on page 62
- ◆ Section 6.29, “List Devices,” on page 64
- ◆ Section 6.30, “List Linux Volume,” on page 68
- ◆ Section 6.31, “List Linux Volumes,” on page 70
- ◆ Section 6.32, “List Move,” on page 72
- ◆ Section 6.33, “List Moves,” on page 74
- ◆ Section 6.34, “List Partition,” on page 76
- ◆ Section 6.35, “List Partitions,” on page 78
- ◆ Section 6.36, “List Pool,” on page 84
- ◆ Section 6.37, “List Pools,” on page 87
- ◆ Section 6.38, “List Snap,” on page 91
- ◆ Section 6.39, “List Snaps,” on page 92
- ◆ Section 6.40, “List Volume,” on page 94
- ◆ Section 6.41, “List Volumes,” on page 96
- ◆ Section 6.42, “Mount,” on page 99
- ◆ Section 6.43, “Move,” on page 99
- ◆ Section 6.44, “Pause Move,” on page 101
- ◆ Section 6.45, “Pool Activate,” on page 101
- ◆ Section 6.46, “Pool Deactivate,” on page 102
- ◆ Section 6.47, “RAID,” on page 102
- ◆ Section 6.48, “Rename Pool,” on page 105
- ◆ Section 6.49, “Rename RAID,” on page 105
- ◆ Section 6.50, “Rename Volume,” on page 106
- ◆ Section 6.51, “Rescan,” on page 107
- ◆ Section 6.52, “Resume Move,” on page 107
- ◆ Section 6.53, “Share,” on page 108
- ◆ Section 6.54, “Unmount,” on page 108
- ◆ Section 6.55, “Unshare,” on page 109
- ◆ Section 6.56, “Volume Mount,” on page 109
- ◆ Section 6.57, “Volume Unmount,” on page 110

## 6.1 Syntax Overview

Novell Linux Volume Manager can be used to manage NSS file systems or Linux POSIX file systems on your OES server. This section describes the general syntax and conventions for NLVM.

- ♦ [Section 6.1.1, “Syntax,” on page 27](#)
- ♦ [Section 6.1.2, “Syntax Conventions,” on page 27](#)
- ♦ [Section 6.1.3, “Documentation Conventions,” on page 28](#)
- ♦ [Section 6.1.4, “Files,” on page 28](#)

### 6.1.1 Syntax

Using commands for the NLVM program requires `root` user privileges. NLVM options must follow immediately after `nlvm`.

```
nlvm [nlvm_options] <command> <command_options>
```

### 6.1.2 Syntax Conventions

When issuing NLVM commands, consider the following general syntax conventions:

- ♦ [“NSS Pool and Volume Names” on page 27](#)
- ♦ [“NSS Software RAID Names” on page 27](#)
- ♦ [“NCP Volume Names” on page 27](#)
- ♦ [“Order of Command Options” on page 27](#)
- ♦ [“Sizes” on page 28](#)
- ♦ [“Name Format” on page 28](#)

#### NSS Pool and Volume Names

All NSS pool names and NSS volume names are automatically converted to uppercase.

#### NSS Software RAID Names

NSS software RAID names are case sensitive.

#### NCP Volume Names

When you create an NCP volume, the name is automatically converted to uppercase.

#### Order of Command Options

Command options can be specified in any order except where it is otherwise noted. Options with an equal sign (=) can be in any order.

## Sizes

All sizes are in bytes and can be specified with one of the following multipliers: K, M, G, and T. Multipliers are case insensitive and are multiples of 1024. If no multiplier is specified, it is assumed to be G by default. If 'max' is entered, all of the free unpartitioned space on the device is used. All sizes can be entered as whole numbers or with fractional parts such as 200.45G and 3.98T.

Examples for common command options:

`size=20` (If no multiplier is used, it is assumed to be G (gigabytes).)

`size=20G` (You can also specify `max` instead of a value and multiplier.)

`size=3.98T` (You can specify a value with decimal places.)

## Name Format

Examples for common name formats used in command options:

`device=sdb` (You can specify the leaf node name of the device, including multipath names.)

`device=/dev/mapper/mpatha` (You can specify the full Linux path of the device.)

`device=anydisk` (You can specify `anydisk` or `anyshared` keywords if the command allows it.)

`part=sdcl.1` (You can specify only the partition node name, not the full Linux path.)

`part=cluster1.sbd`

`name=MYPOOL1` (All NSS pool names and NSS volume names are converted to uppercase.)

### 6.1.3 Documentation Conventions

In the command syntax for NLVM, the mandatory command options are surrounded by angle brackets (<>). The optional command options are surrounded by square brackets ([]). The brackets are not used when you issue the command. For example, the command syntax conventions are:

```
nlvm [nlvm_options] command <mandatory_value> [options]
```

### 6.1.4 Files

The following are key files used by NLVM:

**`/etc/opt/novell/nss/nlvm.conf`**

Location of the NLVM configuration file.

**`/opt/novell/nss/sbin/nlvm`**

Location of the NLVM utility. It also has a link in the `sbin` directory so that it is in the search path.

**`/var/opt/novell/nss/debug`**

Location of the debug log files.

## 6.2 NLVM Options

The NLVM options can be used as needed with any command, except where it is otherwise noted. NLVM options can appear in any order in the command after `nlvm`.

```
nlvm [nlvm_option] <command> [command_options]
```

- ♦ `-d, --debug`
- ♦ `-f, --force`
- ♦ `-l, --getlock`
- ♦ `-m`
- ♦ `--no-prompt`
- ♦ `-r, --rescan`
- ♦ `-s, --share`
- ♦ `-t, --terse`
- ♦

### **-d, --debug**

This option causes a `/var/opt/novell/log/nss/debug/nlvm_debug.log` file to be created so that the operations can be reviewed. This is helpful in diagnosing problems in running the NLVM utility. Up to 10 debug files can be created; they are numbered automatically.

---

**NOTE:** The debug can be turned on always by using the `/etc/opt/novell/nss/nlvm.conf` file.

---

### **-f, --force**

This option can be used with certain commands to force the command to complete. Support for this NLVM option is indicated in the individual commands.

### **-l, --getlock**

This option forces the command to get the `nlvm` lock. The lock protects multiple users from modifying things at the same time.

Use with caution! This option is to be used only if the lock does not get released properly due to a segment fault or other operation aborts.

### **-m**

This option prevents pools that have been unmounted from being mounted.

Pools are by design auto mounted. Therefore, running the `nssmu` utility, or running most `nlvm` commands without the `-m` option can cause an unmounted pool to be remounted if underlying devices and partitions still exist. To execute an `nlvm` command without mounting the unmounted pools, you must include the `-m` option. The `nlvm mount` command internally sets the `-m` flag, so only the specified pool is mounted.

#### **--no-prompt**

This option can be used with certain commands to prevent a confirmation message from being displayed, such as when you initialize a device or delete Linux POSIX volumes, pool moves, partitions, pools, RAIDs, RAID segments, snapshots, and NSS volumes. Support for this NLVM option is indicated in the individual commands.

#### **-r, --rescan**

This option forces a fresh rescan of the system before executing a command to update the device and partition objects. Use this if something changed the information outside the NSSMU, iManager, or `nlvm` utility.

#### **-s, --share**

This option sets the shared override bit for the command being executed.

In a OES Cluster Services cluster, NLVM uses the cluster's SBD to detect if a node is a cluster member and to lock against concurrent changes to physically shared storage. Without an SBD, NLVM cannot detect whether a node is a member of the cluster and cannot acquire the locks it needs to execute tasks. In this state, you can use the `-s` option with NLVM commands to prepare a device and create an SBD partition. To minimize the risk of corruption, you must ensure that nobody else is changing any storage on any nodes at the same time.

#### **-t, --terse**

This option can be used with `nlvm list` commands to display the output in a format for parsing. Values are labeled in the format `ParameterName=value`. Information about a storage object is output in a single line. The line wraps automatically if the output exceeds the console width.

A request might return multiple lines if the target object contains storage objects, such as partitions on a device or segments in a software RAID. The target object's information appears on the first line, and subsequent lines contain information about each of its member objects. A single blank line separates output for some target objects.

#### **-p**

This option can be used to display all size outputs in the specified human-readable unit of size, as follows: `-pk=kilobytes (KB)`, `-pm=megabytes (MB)`, `-pg=gigabytes (GB)`, `-pt=terabytes (TB)`, `-pp=petabytes (PB)`, `-ps=sectors`, or `-pb=bytes`.

## 6.3 Common Options

Common options can be used as noted with specific commands. Common options are specified at the end of the command.

- ♦ `all`
- ♦ `more`

#### **all**

This option can be used with `nlvm list` commands to display detailed information for all objects of that type on the server. It displays the same information as a specific `nlvm list` request against an object. It can be used with the `-t` or `--terse` NLVM option to format the detailed output for parsing.

**more**

This option can be used with `nlvm list` commands to display more information than appears in the standard output. It can be used with the `-t` or `--terse` NLVM option to format the enhanced output for parsing.

## 6.4 Complete Move

**complete move <move\_name>**

Check to see if an NSS pool move is complete. If the move is complete, the old location is deleted. If the move is not completed, it will return an error 11 (EAGAIN).

If a pool is cluster-enabled, issue the command on the node where its pool cluster resource is currently online.

```
nlvm [nlvm_options] complete move <move_name>
```

### Command Option

**move\_name**

Mandatory. Specify the name of the move object to check. The move name typically looks like `POOLNAME_move`.

### Command Example

```
nlvm complete move MYPOOL1_move
```

Verify that the move `MPOOL1_move` is complete. If it is, delete the old location of the pool.

## 6.5 Create Linux Volume

**create linux volume <type> <<device> [size] | <part>> [mp] [mkopt] [mntopt] [lvm] [name] [group] [shared] [ip] [ncp] [valid]**

Create a Linux POSIX volume on a device.

```
nlvm [nlvm_options] create linux volume <type> <<device> [size] |  
<part>> [mp] [mkopt] [mntopt] [lvm] [name] [group] [shared] [ip] [ncp]  
[valid]
```

For a cluster-enabled LVM volume, issue the command from the master node in the cluster.

### Command Options

**type=fstype**

Mandatory. Specify the type of Linux POSIX file system to use for `mkfs`.

Supported file system types are `btrfs` (in OES 11 SP1 and later; requires the `btrfsprogs` package), `ext2`, `ext3`, `ext4`, `reiserfs`, and `xf`s.

### Examples

```
type=ext3  
type=reiserfs
```

**device=<devicename|anydisk>**

Mandatory unless the `part` option is used. Specify the device to use for the Linux POSIX volume, or specify the keyword `anydisk`.

---

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

---

If the device is seen by a single server, or a single node in a cluster, do not use the `shared` option.

If the device is seen by multiple nodes in a OES Cluster Services cluster, you must specify the `devicename` and use the `shared`, `ip`, `name`, `lvm`, and `group` (optional) options to create the Linux volume group cluster resource. Specify an unshared initialized device. For OES 11 SP2 and later, you can alternatively specify a shared device with no data partitions or an uninitialized device. The cluster-enabled LVM volume group uses the entire device. OES Cluster Services mounts the cluster resource exclusively on one node at a time.

**Examples**

```
device=sdb
device=/dev/sdb
device=anydisk
device=mpatha
device=/dev/mapper/mpatha
```

**size=<value[K|M|G|T]|max>**

Mandatory unless the `shared` option is used, or unless the `part` option is used instead of the `device` option. Specify a size of the partition to create for the Linux volume, or specify `max` to use all of the free unpartitioned space for the volume. The minimum allowed size is 8 MB.

If the `shared` option is used, the entire device is dedicated to the LVM volume group. If the `size` option is specified, it is ignored.

If the `part` option is used, the entire partition is dedicated to the volume. If the `size` option is specified, it is ignored.

**Examples**

```
size=20G
size=100m
size=max
```

**part=partition\_name**

Specify the node name (such as `sdC2`) for the partition you want to use for a non-clustered volume. The partition must exist; it is not created with this command. The partition type must be compatible with the type of Linux volume you want to create on it, such as type 83 for a Linux native volume or type 8E for a Linux LVM volume. The entire partition is used for the volume you create.

Do not specify the `part` option in combination with the `device` option. The `size` option is ignored.

Do not specify the `part` option in combination with the `shared` option. You can use a partition only for non-clustered volumes.



### Example

```
part=sdc2
```

### **mp=</mount\_path>**

Specify the path of the mount point where the volume is to be mounted. If the path does not currently exist, it will be created.

For LVM volumes, the name option must be used with the `lvm` option to specify a volume name. The full mount point path can specify a directory path that is the same or different than the specified volume name. If a mount path is not specified for an LVM volume or a clustered LVM volume, the utility assigns a default mount path of `/usr/novell/<volume_name>`.

For Linux POSIX volumes, the final directory of the full mount point path is used as the volume name. For example, if the mount point is `/home/users/bob`, the volume name is `bob`. The final directory name must be unique as a volume name on the server. If you use the `ncp` option, the NCP volume name is based on the final directory name, but all letters are capitalized. Ensure that the final directory name does not exceed 14 characters and does not use special characters. Letters A-Z, letters a-z, numbers 0-9, and underscores are supported.

If a mount path is not specified for a Linux POSIX volume, the utility assigns a default mount path of `/usr/novell/<file_system_type>_<next_available_number>`. For example, if the file system type is `ext3`, the default mount path is `/usr/novell/ext3_0`. If that path is not available, the path is `/usr/novell/ext3_1`, and so forth until a unique volume name is achieved.

### Example

```
mp=/home
```

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

Specify the options to use when running `mkfs`. For a list of available options, see the `mkfs(8)` man page. No default option is specified.

### Example

```
mkopt=-v
```

### **mntopt=<option1[,option2[,...]>**

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

### Example

```
mntopt=rw
```

### **lvm**

Used to specify that an LVM volume and volume group is to be created.

If the `lvm` option is used, the `name` option must be provided to specify a name for the LVM volume. Specifying a different name for the LVM volume group is optional.

### Example

```
lvm
```

**name=<lvm\_volume\_name>**

Used with the `lvm` option to specify a name for the LVM volume.

If you do not specify the `group` option, this name is also used as the LVM volume group name.

For LVM logical volume naming conventions, see [Section 5.2.5, “Linux LVM Volume Group and Logical Volume Names,” on page 19](#). If you use the `nep` option, the NCP volume name is based on the LVM volume name, but all letters are capitalized. Ensure that the name does not exceed 14 characters and does not use special characters. Letters A-Z, letters a-z, numbers 0-9, and underscores are supported.

If the `lvm` option is not specified, this option is ignored.

**Example**

```
name=mylvmvol1
```

**group=<lvm\_volume\_group\_name>**

Optional. Used with the `lvm` option to specify a name for the LVM volume group. If the `group` option is not specified, the volume group name is the same as the LVM volume name.

For LVM volume group naming conventions, see [Section 5.2.5, “Linux LVM Volume Group and Logical Volume Names,” on page 19](#).

If the `lvm` option is not specified, this option is ignored.

**Example**

```
group=clustervg01
```

**shared**

Used to cluster-enable an LVM volume group. This creates an LVM volume group cluster resource, including its load, unload, and monitoring scripts, for use in an existing OES Cluster Services cluster. The cluster resource name is the LVM volume group name plus `_resource`; that is, `<lvm_vg_name>_resource`. For example, `mylvmvg01_resource`. The resource is created and set to an Offline state. You can use the Clusters plug-in in iManager to modify the scripts and resource settings as needed, and then use iManager or cluster commands to online the resource.

If the `shared` option is used, the `ip`, `name`, and `lvm` options must also be provided. You can use the `group` option to specify a different name for the LVM volume group.

The device must be uninitialized and not marked as shareable. If the device is already initialized, then the device must be uninitialized using the following command:

```
/sbin/wipefs -fa <device>
```

---

**CAUTION:** This command completely removes all the existing data or file system.

---

The LVM volume group uses the entire device. Use OES Cluster Services tools or commands to online the cluster resource exclusively on one node at a time

**Examples**

```
shared lvm ip=10.10.10.101 name=mylvmvol1
shared lvm ip=10.10.10.101 name=mylvmvol1 group=mylvmvg1
```

### **ip=<IP\_address\_for\_LVM\_volgroup\_cluster\_resource>**

Used with the `shared` option to specify the IP address to use for the Linux volume group cluster resource. This is required for cluster-enabled Linux volume groups on OES Cluster Services clusters. Specify the IP address in IPv4 format.

If the `shared` option is not specified, this option is ignored.

#### **Example**

```
ip=10.10.10.101
```

### **ncp**

Used to enable the Linux POSIX file system on the volume to be accessed with the NetWare Control Protocol (NCP). An NCP volume ID is assigned automatically to the volume. You can use the `valid` option in combination with the `shared` and `ncp` options to assign an NCP volume ID for a clustered LVM volume.

If you use the `ncp` option, the volume name used for the `name` option must comply with the name limitations described in [Section 5.2.4, “NCP Volume Names,” on page 19](#).

### **valid=value**

(Optional) Used in combination with the `shared` and `ncp` options to assign an NCP volume ID for a clustered LVM volume. If the `valid` option is not used, a volume ID is automatically assigned. For clustered volumes, the valid range is 254 to 0, in descending order.

In a OES Cluster Services cluster, the volume ID must be unique across all member nodes. In a Business Continuity Cluster, the volume ID must be unique across all nodes in every peer cluster.

#### **Example**

```
lvm shared ip=10.10.10.134 name=lvmvol40 ncp valid=240
```

The `valid` option requires the `shared` and `ncp` options. The `shared` option requires the `lvm`, `ip`, and `name` options.

### **Command Examples**

```
nlvm create linux volume type=ext3 device=sdf size=10G mp=/home/bob  
mntopt=rw
```

Create a 10 GB Linux POSIX volume using the Ext3 file system on the `/dev/sdf` device. Mount the volume on path `/home/bob` with the Read/Write mount option.

```
nlvm create linux volume type=ext3 device=/dev/sdf mp=/home/bob  
mntopt=rw lvm shared ip=10.10.10.101 group=clustervgbob  
name=clustervolbob
```

Create and cluster-enable an LVM volume group on the `/dev/sdf` device with a resource IP address of 10.10.10.101, an LVM volume name of `clustervolbob`, and an LVM volume group name of `clustervgbob`. Create a Linux POSIX volume on the LVM volume using the Ext3 file system. The entire device is dedicated to the LVM volume.

This command automatically creates an LVM volume group cluster resource called `clustervgbob_resource` in a OES Cluster Services cluster where the node is a member. It creates its resource load, unload, and monitoring scripts; sets the resource to offline; and waits to be brought online by using the cluster commands. You manage the resource by using OES Cluster Services tools and commands.

```
nlvm create linux volume type=ext3 device=sdf mp=/home/bob mntopt=rw
lvm shared ip=10.10.10.101 group=clustervgbob name=clustervolbob ncp
valid=240
```

Create and cluster-enable an LVM volume group on the `/dev/sdf` device with a resource IP address of 10.10.10.101, an LVM volume group name of `clustervgbob`, and an LVM volume name of `clustervolbob`. Create a Linux POSIX volume on the LVM volume using the Ext3 file system. The entire device is dedicated to the LVM volume.

NCP-enable the volume and automatically assign it the NCP name of `CLUSTERVOLBOB`, which is the assigned LVM name in all capital letters. Assign it the NCP volume ID of 240, which the administrator knows to be unique across all member nodes in the OES Cluster Services cluster and across all peer clusters in a Business Continuity Cluster.

This command automatically creates an LVM volume group cluster resource called `clustervgbob_resource` in a OES Cluster Services cluster where the node is a member. It creates its resource load, unload, and monitoring scripts; sets the resource to offline; and waits to be brought online by using the cluster commands. You manage the resource by using OES Cluster Services tools and commands.

## 6.6 Create Partition

**create partition <type> <device> <size> [label] [dm]**

Create a partition on a disk.

```
nlvm [nlvm_options] create partition <type> <device> <size> [label]
[dm]
```

The number of partitions per device can be limited by the device partitioning scheme, the partition type, or the device driver, whichever is the most restrictive.

- ◆ **Partitioning scheme:** The MS-DOS format allows up to 4 primary partitions, where 1 can be an extended partition with logical partitions. The GPT format allows up to 128 partitions.
- ◆ **Partition type:** If a device contains only Novell type partitions, the number of partitions is limited only by the space on the disk. If there are any non-Novell partitions on the device, each partition created, including Novell type partitions, will be a physical partition and limited by Linux to 255 partitions.
- ◆ **Device driver:** Check your device vendor's documentation to determine driver restrictions. For example, the Hewlett-Packard CCISS device driver supports up to 15 partitions per device, regardless of the partition type.

### Best Practices for Creating Partitions

- ◆ Disks using Novell partitions should have only Novell partitions on the device.
- ◆ Do not create more than 15 partitions on a device.

### Command Options

**type=partition\_type**

Mandatory. You must specify the partition type in hexadecimal, without the leading `0x`.

Before you create a OES Cluster Services SBD (split brain detector) partition with `type=1ad`, you must take the cluster down, and stop OES Cluster Services from running on all nodes.

## Examples

```
type=83 (partition type for Linux)
type=8e (partition type for Linux LVM)
type=169 (partition type for NSS)
type=1ad (partition type for OES Cluster Services SBD partition)
type=1ac (partition type for snapshots)
```

### **device=<devicename|anydisk|anyshared>**

Mandatory. Specify the device to use for the partition, or specify the keyword `anydisk` or `anyshared`.

If you use NLVM to create an SBD, the `nlvm create partition` command can accept an initialized or uninitialized device when you use the `type=1ad` option. NLVM checks the specified device to see if it is initialized, and takes the following actions:

- ♦ **Uninitialized device:** NLVM initializes the device, marks it as Shareable for Clustering, and creates the requested SBD partition.
- ♦ **Initialized and shared device:** NLVM creates the requested SBD partition.
- ♦ **Initialized and unshared device:** NLVM creates the requested SBD partition, but does not alter the shared state. It returns an error warning that the SBD partition is not shared.

You must manually mark the device as Shareable for Clustering after the partition is created. You can use the `nlvm share` command to share the device.

## Examples

```
device=sdb
device=/dev/sdb
device=anydisk
device=anyshared
```

### **size=<value[K|M|G|T]|max>**

Mandatory. Specify the size of the partition to create, or specify `max` to use all free unpartitioned space. The minimum allowed size is 1 MB.

Because a physical partition must end on a cylinder boundary, its size might be slightly different than the size you specify. If the size does not fall naturally on a cylinder boundary, the partition size is rounded up or down, depending on the partition type, the size specified, and the amount of free space. For a Novell type partition (NSS or SBD), the size is rounded down. For a Linux type partition, the size is rounded up if enough free space is available; otherwise, the size is rounded down.

## Examples

```
size=20G
size=100.45M
size=max
```

### **label="Label for the partition"**

Specify the label to be added to a Novell partition type. This option is ignored for other partition types. If the label contains spaces, you must put quotation marks around it. If the label contains a special character, you must escape the character by adding a backslash character (`\`) in front of it.

If you create a OES Cluster Services SBD partition, the label should be the cluster name. For example, if the cluster name is `cluster1`, NLVM creates a partition named `cluster1.sbd`. If an SBD partition already exists for the cluster, the new partition is named `cluster1.sbd1`, and the cluster does not recognize it. To use the new partition for the cluster, you must delete the old partition. Then the new partition is automatically renamed as `cluster1.sbd`, and is used by the cluster.

### Examples

```
label="This label has spaces"
label=engineering
label=special\character
label=cluster1
```

### dm

Create a device mapper object for this partition in the `/dev/nss` directory. This is useful when creating Novell partition types that need to be accessed directly.

### Example

```
dm
```

### Command Examples

```
nlvm create partition type=169 device=sdb size=20G dm
```

Create an NSS partition on the `/dev/sdb` device of size 20 GB. Also create a device mapper object for the partition, `/dev/nss/sdb1.1`.

```
nlvm create partition type=83 device=sdc size=200G
```

Create a Linux partition on the `/dev/sdc` device of size 200 GB.

```
nlvm create partition type=8e device=sdf size=200G
```

Create a Linux LVM partition on the `/dev/sdf` device of size 200 GB.

```
nlvm -s create partition type=1ad device=sdg size=max label=cluster1
```

Take the cluster down and stop OES Cluster Services. Create a OES Cluster Services SBD partition on the `/dev/sdg` device, and use all available free space on the device. Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

## 6.7 Create Pool

```
create pool <name> <size [device]|part> [ip] [vsn] [csn] [cifs]
```

Create an NSS pool.

```
nlvm [nlvm_options] create pool <name> <size [device]|part> [ip] [vsn]
[csn] [cifs] [type]
```

For a cluster-enabled pool, issue the command from the master node in the cluster.

---

**NOTE:** NSS64-bit pools are by default Trustee Index media upgraded. Creating NSS64-bit pools in a mixed-node cluster environment is not recommended, because the pools will not be accessible from nodes older than OES 2015 SP1. You can still go ahead and force the creation of the pool using the `-f` or `--force` options. As a workaround, configure preferred nodes for each media-upgraded cluster resource so that these resources load on OES 2015 SP1 or later nodes.

For more information on creating preferred nodes, see “[Configuring Preferred Nodes and Node Failover Order for a Resource](#)” in the *OES 2023: OES Cluster Services for Linux Administration Guide*.

---

## Command Options

### **name=*pool\_name***

Mandatory. Specify the name of the pool to create. This name must be unique from other pools. The pool name is automatically converted to uppercase.

Pool names are 2 to 15 characters. Uppercase letters A to Z, number characters 0 to 9, and underscore ( `_` ) are valid characters for all pools. Names cannot start or end in underscore, and cannot contain double underscores.

If the pool is not shared, the pool can also contain special characters:

```
!@#%& ( )
```

Names that contain special characters must be enclosed in quotation marks in all commands and scripts.

The names cannot be reserved names such as `con`, `com`, `lpt`, `pipe`, `all`, and so on.

### **Example**

```
name=MYPOOL1
```

### **size=<value[K|M|G|T]|max>**

Mandatory. Specify the amount of space to be used on the associated device. The size is not used if you specify the `part=` option instead of `device=`. The total pool size must be greater than 10 megabytes.

If multiple devices are specified, each `device` option instance must have a matching `size` option instance. The first `size` instance is matched to the first `device` instance, and so on.

### **Examples**

```
size=200G
size=3.98T
```

### **device=<devicename|anydisk|anyshared>**

Specify the device to use for the pool, or specify the keyword `anydisk` or `anyshared`.

Do not specify the `device` option in combination with the `part` option.

You can specify multiple device instances to create a pool composed of multiple segments. Each `device` option instance must have a matching `size` option instance. The first `device` instance is matched to the first `size` instance, and so on. When specifying multiple devices, device names must be provided for each instance.

### **Examples**

```
device=sdb
device=sde device=sdf device=sdg      (Specify a size for each
instance.)
device=anydisk
device=anyshared
```

**part=partition\_name**

Specify the node name (such as `sdcl.1`) for the partition where you want to create the pool. The partition must exist; it is not created with this command. The entire partition is used for the pool.

Do not specify the `part` option in combination with the `device` option.

**Example**

```
part=sdcl.1
```

**ip=ip\_address**

Specify this option to create a cluster enabled pool. If using this option, the device or partition must be shared. This option is mandatory if you are creating a cluster enabled pool.

**Example**

```
ip=10.10.10.41
```

**vsn=virtual\_server\_name**

Specify the virtual server name for a cluster enabled pool. It is optional and used only for cluster enabled pools. If a name is not supplied, the default name will be used in the format of `<clusterName>-<poolName>-SERVER`. Underscores in the cluster name or pool name are changed to hyphens. If you customize the virtual server name, you can use letters, numbers, hyphens, and underscores.

**Examples**

```
vsn=CLUSTER2-POOL-2-SERVER
vsn=C1-P1-SERVER
vsn=MY-CUSTOM-NAME
```

**csn=cifs\_virtual\_server\_name**

Specify the CIFS virtual server name for a cluster enabled pool. It is optional and used for cluster enabled pools where CIFS is enabled as an advertising protocol. The name can be up to 15 characters, which is a restriction of the CIFS protocol.

For users to collaborate effectively, all paths for user access should be identical, independent of the access protocol used. This is possible only if the same name is used for the NCP virtual server name and the CIFS virtual server name, and the name can be only up to 15 characters.

If the `cifs` option is used without the `csn` option, the NCP virtual server name is used as the CIFS virtual server name. In this case, if the name is more than 15 characters, the CIFS virtual server name uses the rightmost 13 characters and adds `-W`. For example, an NCP virtual server name of `CLUSTER1-P_USERS` is modified to `STER1-P_USERS-W` for the CIFS virtual server name. If a default NCP virtual server name was used in the form of `<clusterName>-<poolName>-SERVER` and the name exceeds 15 characters, the CIFS virtual server name uses the rightmost 13 characters of the `<clusterName>-<poolName>` part of the name and adds `-W`. For example, an NCP virtual server name of `CLUS1-P123-SERVER` is modified to `CLUS1-P123-W` for the CIFS virtual server name.

To use the NCP virtual server name for the CIFS server name, use the `nlvm` command as follows without the `csn` option:



```
nlvm create pool name=a4 size=15M device=sdb ip=10.10.10.39 vsn=pqr  
cifs
```

In this example, `pqr` is used as the NCP virtual server name and CIFS virtual server name.

If an administrator user later changes the NCP virtual server name in NSSMU or iManager, NSSMU automatically applies the name change to the CIFS virtual server name, so that the administrator does not need to make the change twice. To use a different server name for the CIFS virtual server name, you can change the CIFS virtual server name by using the CIFS management tools. This change will not affect the NCP virtual server name.

### Examples

```
csn=CLUS1-P1  
csn=c1-p123
```

### cifs

Specify this option to enable CIFS as an advertising protocol when you create a cluster enabled pool. By default, CIFS is disabled as an advertising protocol. OES CIFS must be installed on the machine in order for this option to work.

You can use the `csn` option to specify a CIFS virtual server name. Without the `csn` option, the NCP virtual server name is used as the CIFS virtual server name. See the `csn` option for details.

### Example

```
cifs
```

### type

NSS supports two types of pools: NSS64 and NSS32. NSS32-bit pools use 32-bit block addressing and supports up to 8 TB, whereas, NSS64 pools use 64-bit block addressing and supports up to 8 EB (exabyte). When creating a pool, specify the pool type. Beginning with OES 2018, if you do not specify the type, NLVM defaults to NSS64. All pools prior to OES 2015 use 32-bit block addressing and they are of type NSS32. You cannot change the pool type later.

---

**NOTE:** NSS64-bit pools are by default Trustee Index media upgraded. Creating NSS64-bit pools in a mixed-node cluster environment is not recommended, because the pools will not be accessible from nodes older than OES 2015 SP1. You can still go ahead and force the creation of the pool using the `-f` or `--force` options. As a workaround, configure preferred nodes for each media-upgraded cluster resource so that these resources load on OES 2015 SP1 or later nodes.

---

### Example

```
type=nss64
```

### `-f, --force`

This option, when used with the `create pool` command, forcefully creates an NSS64-bit pool in a mixed-node cluster environment.

### Command Examples

```
nlvm create pool name=MYPOOL1 size=20G device=sdb
```

Create a pool named MYPOOL1 on device `/dev/sdb` that is 20 GB in size.

```
nlvm create pool name=MYPOOL2 size=20G device=sdb size=100G device=sdg
```

Create a pool named MYPOOL2 that is a total of 120 GB in size. Use 20 GB of free space from device /dev/sdb. Use 100 GB of free space from device /dev/sdg.

```
nlvm create pool name=MYPOOL2 size=200G device=anydisk
```

Create a pool named MYPOOL2 on any device that has 200 GB of free unpartitioned space available.

```
nlvm create pool name=MYPOOL3 size=100G device=anyshared
```

Create a pool named MYPOOL3 on any shared device that has 100 GB of free unpartitioned space available.

```
nlvm create pool name=MYPOOL4 part=sdcl.1
```

Create a pool named MYPOOL4 on partition /dev/sdcl.1 and use all of the partition.

```
nlvm -f create pool name=MYPOOL6 size=10G device=sd type=NSS64  
ip=192.168.1.1
```

Forcefully creates the NSS64-bit pool named MYPOOL6 in a mixed-node cluster environment.

## 6.8 Create RAID

```
create raid <name> <raid> [type] <size> <device> [stripe] [part]
```

Create an NSS software RAID device or an SBD software RAID device.

```
nlvm [nlvm_options] create raid <name> <raid> [type] <size> <device>  
[stripe] [part]
```

### Command Options

**name=raid\_name**

Mandatory except when you mirror an existing SBD partition. This name must be unique from other RAID devices. The RAID name is case sensitive.

When you create an NSS software RAID device, you must specify the name of the device to create.

When you create a new OES Cluster Services SBD RAID 1 device, you must specify the name of the device to create. The name must match the name of an existing cluster (such as `cluster1`) that has a Cluster object in NetIQ eDirectory. This allows the SBD to be used by the cluster. The name is case sensitive.

When you mirror an existing OES Cluster Services SBD partition, the name is optional. If you specify a name (which should be the cluster name), the RAID 1 is given that name. If the name is not specified, the RAID 1 name defaults to the SBD partition's name.

RAID names are 2 to 58 characters. Names are preferred to use characters A to Z, a to z, 0 to 9, and underscore (\_). Names cannot start or end in underscore, and cannot contain double underscores. Printable ASCII characters (see decimal codes 33 to 122 in a code chart) are valid.

RAID names can also contain the following special characters:

```
!@#$%&()
```

Names that contain special characters must be enclosed in quotation marks in all commands and scripts. On the BASH command line, each special character must be escaped by preceding it with a backslash character (\).

The names cannot be reserved names such as `con`, `com`, `lpt`, `pipe`, `all`, and so on.

#### Example

```
name=MYRAID1
```

#### **raid=<0|1|5>**

Mandatory. Specify the RAID type. Valid options are 0 for striping, 1 for mirrored, or 5 for striping with parity.

#### Example

```
raid=1
```

#### **type=<nss|sbd>**

Mandatory except when you mirror an existing partition. Specify the type of partition to mirror. This option is used only for RAID 1. Valid options are `nss` and `sbd` (OES Cluster Services split-brain detector). The default mirror type is `nss`.

Before you create a new SBD RAID 1, you must take the cluster down, and stop OES Cluster Services from running on all nodes. This is not necessary when you mirror an existing SBD partition.

#### Examples

```
type=nss
```

```
type=sbd
```

#### **size=<value[K|M|G|T]|max>**

Mandatory except when you mirror an existing partition. Specify the size of each segment of the RAID. The minimum size is 12 megabytes.

Because a physical partition must end on a cylinder boundary, its size might be slightly smaller than the size you specify. If the size does not fall naturally on a cylinder boundary, the partition size is rounded down for Novell type partitions.

#### Examples

```
size=20G
```

```
size=1.45T
```

#### **device=devicename**

Mandatory. Specify the device to create a RAID segment on. This option is used multiple times, once for each segment to create. RAID 0 or RAID 1 requires a minimum of two devices. RAID 5 requires a minimum of three devices. Devices must be unique for each instance.

#### Example

```
device=sdb device=sdc device=sdd
```

**stripe=stripe\_size**

Specify the RAID stripe size in bytes. This option is applicable only for RAID 0 and RAID 5. The stripe size must be a power of 2, with a minimum size of 4 KB and a maximum size of 256 KB. The default stripe size is 64 KB.

**Example**

```
stripe=64K
```

**part=partition\_name**

Specify the node name for the partition to be mirrored. Use this option to mirror an existing NSS partition (such as `sdcl.1`) or OES Cluster Services SBD partition (such as `cluster1.sbd`). The existing partition is the first segment of a RAID 1 mirror.

If the `part` option is used, the RAID size option is ignored. Each segment's size is the size of the existing partition. The data on the original partition is mirrored on up to three specified devices.

After you mirror the partition, you manage the RAID 1 device by using the normal NSS software RAID management tools and commands.

**Examples**

```
part=sdcl.1
part=cluster1.sbd
```

**Command Examples**

```
nlvm create raid name=MYRAID5 size=20G raid=5 device=sdb device=sdcl.1 device=sdd
```

Create a RAID 5 (striping with parity) device that has segments of 20 GB each on devices `/dev/sdb`, `/dev/sdc`, and `/dev/sdd`. The default stripe size of 64 KB is automatically applied. The default partition type is `nss`.

```
nlvm create raid name=MYRAID1 raid=1 device=sdf part=sdcl.1
```

Create a RAID 1 (mirror) for the existing NSS pool partition `/dev/sdc1.1` on the `/dev/sdf` device. The partition type is the same as the existing partition's type. The pool's existing partition becomes the first segment of the RAID, and its existing data is mirrored to device `/dev/sdf`.

```
nlvm -s create raid name=cluster1 raid=1 type=sbd device=sdcl.1 size=max device=sde
```

Before you issue the command, take the cluster down, and then stop OES Cluster Services on all nodes. Create a new OES Cluster Services SBD RAID 1 device for a cluster named `cluster1`. Use devices `sdcl.1` and `sde`. Use the maximum space available as the partition size, based on the smaller of the two devices. Specify the size only once. Use the `-s` NLVM option to override the shared locking requirement and force the command to execute. Afterwards, join the nodes to the cluster.

```
nlvm create raid name=cluster2 raid=1 part=cluster2.sbd device=sdf
```

Mirror an existing OES Cluster Services SBD partition named `cluster2.sbd`. The RAID type is RAID 1. The name `cluster2` is the same name as the cluster that uses the SBD partition. This name is also the same as the label on the existing SBD partition. The partition is mirrored on the previously initialized and shared device `/dev/sdf`. Device `sdf`

is at least the size of the existing partition, and can be formatted as MSDOS or GPT. The new SBD RAID 1 device is named `cluster2.sbd`. The mirrored SBD partitions are named `cluster2.msbd0` and `cluster2.msbd1`.

## 6.9 Create Snap

**create snap <name> <pool> <<device> <size>|<part>> [chunk]**

Create a snapshot of an NSS pool.

```
nlvm [nlvm_options] create snap <name> <pool> <<device> <size>|<part>>
[chunk]
```

For the stored-on location, you can specify the device and size, or specify an existing snap partition (type 1AC).

### Command Options

**name=***snapshot\_name*

Mandatory. Specify the name of the NSS snapshot. This name must be a unique snap name on the server. The snap name is automatically converted to uppercase.

Pool snapshot names are 2 to 15 characters. The naming conventions are the same as for pools.

#### Example

```
name=POOL1SNAP
```

**pool=***pool\_name*

Mandatory. Specify the name of an existing pool that you want to snap.

#### Example

```
pool=MYPOOL1
```

**device=***devicename*

Specify the device where you want to store the copy-on-write data for this snapshot. Use the `size` option to specify the amount of space to use on the device.

The `device` and `size` options are used instead of the `part` option.

#### Example

```
device=sdb
```

**size=***<value[K|M|G|T]|max>*

Specify the amount of space to use on the specified device. The minimum size is 50 MB; there is no maximum. A snap partition (type 1AC) of the specified size is created on the specified device.

NSSMU restricts the maximum snapshot size to 8 TB.

#### Examples

```
size=20G
size=100.50M
```

**part=<snap\_partition>**

Specify an existing, but currently unused, snap partition (type 1AC) where you want to store the copy-on-write data for this snapshot. Because the partition will be re-initialized and associated with this snapshot, it must not belong to any current snapshot. A snap partition can be used by only one snapshot.

Only a partition of type 1AC (snapshot) is allowed; all other partition types result in an error.

The `part` option is used instead of the `device` and `size` options.

**Example**

```
part=sdd3
```

**chunk=chunk\_size**

Specify the chunk size of the snapshot in bytes. The default size is 64 KB. The chunk size must be a power of 2, with the minimum size of 512 bytes, and a maximum size of 256 KB.

**Example**

```
chunk=128K
```

**Command Example**

```
nlvm create snap name=POOL1SNAP pool=MYPOOL1 device=sdb size=20G  
chunk=128K
```

Create a snapshot named `POOL1SNAP` of pool `MYPOOL1`. The copy-on-write partition is on device `/dev/sdb` and of size 20 GB, and the snapshot chunk size is 128 KB.

```
nlvm create snap name=POOL2SNAP pool=MYPOOL2 part=sdd3 chunk=128K
```

Create a snapshot named `POOL2SNAP` of pool `MYPOOL2`. The copy-on-write partition uses an existing but unused partition `/dev/sdd3` of type 1AC (snapshot), and the snapshot chunk size is 128 KB. The specified partition is re-initialized and assigned to snap `POOL2SNAP`.

## 6.10 Create Volume

```
create volume <name> <pool> [passw] [keytype] [quota] [valid] [adenable]
```

Create an NSS volume on an existing pool. NSS volumes are always mounted at `/media/nss/<VolumeName>` unless otherwise specified.

```
nlvm [nlvm_options] create volume <name> <pool> [passw] [keytype]  
[quota] [valid] [adenable]
```

**Command Options**

```
name=volume_name
```

Mandatory. Specify the name of the NSS volume to create. This name must be unique from other volumes. The volume name is automatically converted to uppercase.

Volume names are 2 to 15 characters. The naming conventions are the same as for pools.

**Example**

```
name=MYVOL1
```

**pool=pool\_name**

Mandatory. Specify the name of an existing NSS pool where you want to create the volume.

**Example**

```
pool=MYPOOL1
```

**passwd=password**

Specify a password if the volume is an encrypted volume.

**Example**

```
passwd=novell
```

**keytype=<AES128/AES256>**

Optional. Specify a keytype value along with the password if the volume is an encrypted volume. If the `keytype` option is not used, AES128 encryption algorithm is used by default on the encrypted volume. Valid options are AES128 and AES256. This option should be used only with the `passwd` option.

The AES256 value should be used only on the NSS64 pool type with the pool media upgraded to AES.

**Example**

```
keytype=AES256
```

**quota=size**

Optional. Specify a quota for the volume. A quota is the maximum amount of space in the pool that can be used by the volume. If no quota is specified or if the quota value exceeds the size of the pool, the volume can grow to the size of the pool.

If the maximum pool size is smaller than the specified volume quota, the volume can grow only to the size of the pool. If you later expand the size of the pool, then the volume quota is again the limiting factor.

**Example**

```
quota=500G
```

**valid=value**

(Optional) Used in combination with a clustered NSS pool to assign an NCP volume ID for a clustered NSS volume. If the `valid` option is not used, a volume ID is automatically assigned. For clustered volumes, the valid range is 254 to 0, in descending order.

In a Cluster Services cluster, the volume ID must be unique across all member nodes. In a Business Continuity Cluster, the volume ID must be unique across all nodes in every peer cluster.

**Example**

```
pool=MYPOOL50 valid=250
```

MYPOOL50 is a clustered NSS pool.

**adenable=yes**

(Optional) To create a volume with AD enabled, enter the value as 'yes'.

## Example

```
adenable=yes
```

### Command Examples

```
nlvm create volume name=MYVOL1 pool=MYPOOL1
```

Create a non-encrypted NSS volume on an existing pool named MYPOOL1.

```
nlvm create volume name=MYVOL1 pool=MYPOOL1 passw=novell keytype=AES256
```

Create an NSS volume on an existing pool named MYPOOL1, and encrypt the volume using the AES256 encryption algorithm and password as novell.

```
nlvm create volume name=MYVOL1 pool=MYPOOL1 quota=500G
```

Create a non-encrypted NSS volume on an existing pool named MYPOOL1. The volume has a quota of 500 GB.

```
nlvm create volume name=MYVOL50 pool=MYPOOL50 volid=250
```

Create a non-encrypted, clustered NSS volume on an existing clustered NSS pool named MYPOOL50. Assign it the NCP volume ID of 250, which the administrator knows to be unique across all member nodes in the OES Cluster Services cluster and across all peer clusters in a Business Continuity Cluster.

```
nlvm create volume name=MYVOL1 pool=MYPOOL1 adenable=yes
```

Create a non-encrypted NSS volume on an existing NSS pool named MYPOOL1 with AD media enabled.

## 6.11 Delete Linux Volume

### delete linux volume <volume\_name>

Delete an existing Linux POSIX volume. You cannot delete the root (/) volume. You must unmount the volume before you can delete it.

If the volume is a clustered LVM volume group and logical volume, you must take the cluster resource offline, and then delete the resource before you can delete the volume.

```
nlvm [nlvm_options] delete linux volume <volume_name>
```

You are automatically prompted to confirm the delete action. Respond by typing `yes` or `no`, then press Enter. Use the `--no-prompt` NLVM option to suppress the confirmation prompt.

You can use the `nlvm list linux volumes` command to find the *volume\_name*. A Linux POSIX volume is preceded by a forward slash, such as `/vol1`. This is the last directory of the mount point path that you provided when you created the Linux POSIX volume with NLVM or NSSMU. An LVM volume name is the volume name you used when you created the volume, such as `lvvol1`.

### Command Options

#### *volume\_name*

Mandatory. Specify the name of the volume to delete.

### Examples

For a Linux POSIX volume mounted at `/home/bob`, the volume name is `/bob`.



For an LVM logical volume that you named `lvvol1` that is mounted at `/mnt/lvvol1`, the volume name is `lvvol1` (with no forward slash).

For an LVM logical volume that you named `lvvol2` that is mounted at `/home/users`, the volume name is `lvvol2` (not `/users`).

#### **--no-prompt**

Optional. Specify this NLVM option to prevent a confirmation message from being displayed.

#### **Example**

```
--no-prompt
```

#### **Command Examples**

```
nlvm delete linux volume /bob
```

Delete the Linux POSIX volume that is mounted at `/home/bob`.

```
nlvm delete linux volume lvvol1
```

Delete the Linux LVM logical volume `lvvol1` that is mounted at `/mnt/lvvol1`.

```
nlvm --no-prompt delete linux volume lvvol2
```

Delete the Linux LVM logical volume `lvvol2` that is mounted at `/home/users`. The confirmation message is not displayed.

## 6.12 Delete Move

**delete move <<move\_name>|<pool\_name>>**

Delete an NSS pool move. This command deletes the move request, returns the pool back to its original location, and removes the new location. You can delete the move at any time while the move is in progress, even if it is pending only the `complete move` command to be finalized.

Use the `complete move` command if you want to keep the new location and remove the original location.

If a pool is cluster-enabled, issue the command on the node where its pool cluster resource is currently online.

```
nlvm [nlvm_options] delete move <<move_name>|<pool_name>>
```

You are automatically prompted to confirm the delete action. Respond by typing `yes` or `no`, then press Enter. Use the `--no-prompt` NLVM option to suppress the confirmation prompt.

#### **Command Options**

##### ***move\_name* or *pool\_name***

Mandatory. Specify the name of the NSS pool move to delete, such as `POOLNAME_move`. You can alternatively specify the pool name.

#### **--no-prompt**

Optional. Specify this NLVM option to prevent a confirmation message from being displayed.

#### **Example**

```
--no-prompt
```

### Command Example

```
nlvm delete move MYPOOL_move
```

Delete the pool move named `MYPOOL_move`. This removes the new location, and sets the pool to the original location.

## 6.13 Delete Partition

### **delete partition** <partition\_name>

Delete an existing partition by name.

```
nlvm [nlvm_options] delete partition <partition_name>
```

You are automatically prompted to confirm the delete action. Respond by typing `yes` or `no`, then press Enter. Use the `--no-prompt` NLVM option to suppress the confirmation prompt.

Before you create a OES Cluster Services SBD partition, you must take the cluster down, and stop OES Cluster Services from running on all nodes.

### Command Options

#### **partition\_name**

Mandatory. Specify the node name (such as `sdcl.1`) of the partition to be deleted.

#### Example

```
sdcl.1
```

#### **-f, --force**

Optional. The `force` NLVM option can be used with the `delete partition` command if the partition is part of a pool or move. If the partition is part of a pool, deleting the partition automatically deletes the pool. If the partition is part of a move destination, deleting the partition automatically deletes the pool move.

#### Examples

```
-f  
--force
```

#### **--no-prompt**

Optional. Specify this NLVM option to prevent a confirmation message from being displayed.

#### Example

```
--no-prompt
```

### Command Examples

```
nlvm delete partition sdcl.1
```

Delete the partition `/dev/sdcl.1`.

```
nlvm --force delete partition sdd1.2
```

Delete the partition `/dev/sdd1.2` that is part of an NSS pool move destination. The pool move is deleted as well.

## 6.14 Delete Pool

### **delete pool <pool\_name>**

Delete an existing NSS pool by name.

```
nlvm [nlvm_options] delete pool <pool_name>
```

You are automatically prompted to confirm the delete action. Respond by typing `yes` or `no`, then press Enter. Use the `--no-prompt` NLVM option to suppress the confirmation prompt.

#### **Command Options**

##### **pool\_name**

Mandatory. Specify the name of the NSS pool to be deleted.

##### **Example**

```
MYPOOL1
```

##### **--no-prompt**

Optional. Specify this NLVM option to prevent a confirmation message from being displayed.

##### **Example**

```
--no-prompt
```

#### **Command Example**

```
nlvm delete pool MYPOOL1
```

Delete the NSS pool named `MYPOOL1`.

## 6.15 Delete RAID

### **delete raid <raid\_name>**

Delete an existing NSS software RAID device by name.

If the RAID device is a single element RAID 1, this command removes the RAID 1 mirror object from the pool partition and leaves the pool on the corresponding partition. The pool is not deleted and no data is destroyed.

When you delete a single element RAID1 mirror for an SBD (split-brain detector) partition, it removes the mirror object and leaves the SBD in the corresponding partition. The SBD is not deleted and no data is destroyed.

For a RAID1 that contains multiple elements, deleting the RAID1 deletes all mirrors and the pool partitions or SBD partitions on them. All data is destroyed. If you want to keep the pool or SBD on one of the member devices, use the `nlvm delete partition` command to delete the partitions for mirror elements you do not want to keep. For the remaining single-element mirror, go to the RAIDs page and delete the RAID1 mirror element. This removes the RAID1 object and leaves the pool partition or SBD partition.

```
nlvm [nlvm_options] delete raid <raid_name>
```

You are automatically prompted to confirm the delete action. Respond by typing `yes` or `no`, then press Enter. Use the `--no-prompt` NLVM option to suppress the confirmation prompt.

For single element RAID 1 devices, this command duplicates the `nlvm raid delete` command, which deletes a single element mirror from a pool, leaving the pool on the corresponding partition.

Before you delete a OES Cluster Services SBD RAID 1, you must take the cluster down, and stop OES Cluster Services from running on all nodes.

### Command Options

#### *raid\_name*

Mandatory. Specify the name of the NSS software RAID device to be deleted.

#### Example

```
MYRAID1
```

#### `--no-prompt`

Optional. Specify this NLVM option to prevent a confirmation message from being displayed.

#### Example

```
--no-prompt
```

### Command Example

```
nlvm delete raid MYRAID1
```

Delete the NSS software RAID device named `MYRAID1`.

## 6.16 Delete RAID Segment

**`delete raid <raid_name> segment <number>`**

Delete a specified segment of an existing NSS software RAID device. This is valid only for RAID 1 and RAID 5 devices. RAID 5 can remove only 1 segment, but it must be replaced by another segment in order to have redundancy.

```
nlvm [nlvm_options] delete raid <raid_name> segment <number>
```

You are automatically prompted to confirm the delete action. Respond by typing `yes` or `no`, then press Enter. Use the `--no-prompt` NLVM option to suppress the confirmation prompt.

Use the `--force` NLVM option to remove out-of-sync segments.

### Command Options

#### *raid\_name*

Mandatory. Specify the name of the NSS software RAID device that contains the segment to be deleted.

#### Example

```
MYRAID1
```

#### *number*

Mandatory. Specify the segment index (zero relative) to be removed. For RAID 1, the value must be 0 to 3. For RAID 5, the value must be 0 to 13.

### Example

```
0
```

### **--no-prompt**

Optional. Specify this NLVM option to prevent a confirmation message from being displayed.

### Example

```
--no-prompt
```

### **-f, --force**

Optional. Specify this NLVM option to force the command to delete out-of-sync segments.

### Command Example

```
nlvm delete raid MYPOOL1 segment 0
```

Delete the first segment of the NSS software RAID device named MYRAID1.

```
nlvm --force delete raid MYPOOL1 segment 1
```

Delete the second segment of the NSS software RAID device named MYRAID1. Use the `--force` option to force the deletion of an out-of-sync segment.

## 6.17 Delete Snap

### **delete snap <snap\_name>**

Delete an existing NSS pool snapshot by name.

```
nlvm [nlvm_options] delete snap <snap_name>
```

You are automatically prompted to confirm the delete action. Respond by typing `yes` or `no`, then press Enter. Use the `--no-prompt` NLVM option to suppress the confirmation prompt.

### Command Options

#### **snap\_name**

Mandatory. Specify the name of the NSS pool snapshot to be deleted.

### Example

```
POOL1SNAP
```

### **--no-prompt**

Optional. Specify this NLVM option to prevent a confirmation message from being displayed.

### Example

```
--no-prompt
```

### Command Example

```
nlvm delete snap POOL1SNAP
```

Delete the NSS pool snapshot named POOL1SNAP.

## 6.18 Delete Volume

**delete volume** <volume\_name>

Delete an existing NSS volume by name.

```
nlvm [nlvm_options] delete volume <volume_name>
```

You are automatically prompted to confirm the delete action. Respond by typing `yes` or `no`, then press Enter. Use the `--no-prompt` NLVM option to suppress the confirmation prompt.

### Command Options

#### **volume\_name**

Mandatory. Specify the name of the NSS volume to be deleted.

#### **Example**

```
MYVOL1
```

#### **--no-prompt**

Optional. Specify this NLVM option to prevent a confirmation message from being displayed.

#### **Example**

```
--no-prompt
```

### Command Example

```
nlvm delete volume MYVOL1
```

Delete the NSS volume named `MYVOL1`.

## 6.19 Expand Partition

**expand partition** <partition\_name> <size>

Expand an existing partition.

This command does not add a partition, but expands the existing partition. There must be free space contiguously following this partition in order to expand it.

```
nlvm [nlvm_options] expand partition <partition_name> <size>
```

### Command Options

#### **partition\_name**

Mandatory. Specify the node name (such as `sdcl.1`) of the partition to be expanded. This must be the first command option.

#### **Example**

```
sdcl.1
```

**size=<value[K|M|G|T]|max>**

Mandatory. Specify the amount of space to add to the existing partition.

### Examples

```
size=20G  
size=200.45G
```

### Command Example

```
nlvm expand partition sdc1.1 size=20G
```

Expand the `/dev/sdc1.1` partition by adding the next 20 GB of contiguous free unpartitioned space. For example, if the original partition is 20 GB, the expanded size is 40 GB.

## 6.20 Expand Pool

```
expand pool <pool_name> <size device | part >
```

Expand an existing NSS pool by adding a new partition. Either a partition must be specified, or the device and size must be specified.

If the specified device is the same device as the last segment of the existing pool, and free space exists following the last segment, the utility tries to expand the partition first before trying to add a new partition.

```
nlvm [nlvm_options] expand pool <pool_name> <size device | part >
```

### Command Options

#### **pool\_name**

Mandatory. Specify the name of the NSS pool to be expanded. This must be the first command option.

#### **Example**

```
MYPOOL1
```

#### **device=device\_name**

Specify the device to use for the expanded space.

You can specify multiple `device` option instances to create a pool comprised of multiple segments. Each `device` option instance must have a matching `size` option instance. The first `device` instance is matched to the first `size` instance, and so on.

#### **Example**

```
device=sdb
```

#### **size=<value[K|M|G|T]|max>**

Specify the amount of space to add to the existing pool.

If multiple devices are specified, each `device` option instance must have a matching `size` option instance. The first `size` instance is matched to the first `device` instance, and so on.

### Examples

```
size=20G  
size=100.50M
```

**part=partition\_name**

Specify the name of a partition to add to the pool. The entire partition size is added to the pool's capacity. The partition must be of type 0X169 (nss).

**Example**

```
part=sdc1.1
```

**Command Examples**

```
nlvm expand pool MYPOOL1 device=sdf size=20G
```

Expand the NSS pool named MYPOOL1 by adding the 20 GB of free space from device /dev/sdf.

```
nlvm expand pool MYPOOL2 device=sdf size=20G device=sdg size=100G
```

Expand the NSS pool named MYPOOL1 by adding the 20 GB of free space from device /dev/sdf and 100 GB of free space from device /dev/sdg.

```
nlvm expand pool MYPOOL1 part=sdc1.1
```

Expand the NSS pool named MYPOOL1 by adding the /dev/sdc1.1 partition to it. The partition is type 0X169. The entire partition size is added to the pool's capacity.

## 6.21 Expand RAID

**expand raid <raid\_name> <device>**

Expand an existing NSS software RAID device by adding a new segment. Specify the RAID name and the device to use. The device option can be specified multiple times to specify additional segments. Each device must have a free space area at least as big as the segment size of the RAID.

```
nlvm [nlvm_options] expand raid <raid_name> <device>
```

**Command Options****raid\_name**

Mandatory. Specify the name of the NSS software RAID device to be expanded. This must be the first command option.

**Example**

```
MYRAID1
```

**device=device\_name**

Specify the device to use for the expanded space.

**Example**

```
device=sdb
```

**Command Examples**

```
nlvm expand raid MYRAID1 device=sdf
```

Expand the NSS software RAID device named MYRAID1 by adding the device /dev/sdf.



```
nlvm expand raid MYRAID5 device=sdg device=sdh
```

Expand the NSS software RAID device named MYRAID5 by adding the `/dev/sdg` and `/dev/sdh` devices as two new segments.

## 6.22 Fstrim Pool

```
fstrim pool <pool_name>
```

It is used on a mounted pool to discard (or trim) blocks which are not in use by the pool. You must use this option only on thin-provisioned storage devices.

```
nlvm [nlvm_options] fstrim pool <pool_name>
```

### Command Option

#### *pool\_name*

Mandatory. Specify the name of the NSS pool.

#### Example

```
MYPOOL1
```

### Command Examples

```
nlvm fstrim pool MYPOOL1
```

Trims the unused blocks in the NSS pool named MYPOOL1 and then displays the trimmed size in bytes.

```
nlvm fstrim pool MYPOOL1 async
```

Trims the unused blocks in the NSS pool named MYPOOL1 asynchronously. For more information on the trimmed size in bytes, see logs at `/var/log/messages`.

## 6.23 Fstrim Status

```
fstrim status <pool_name>
```

Displays the fstrim state, last executed time, and last trimmed size on the specified pool.

```
nlvm [nlvm_options] fstrim status <pool_name>
```

### Command Option

#### *pool\_name*

Mandatory. Specify the name of the NSS pool.

#### Example

```
MYPOOL1
```

### Command Example

```
nlvm fstrim status MYPOOL1
```

Displays the fstrim state, last executed time, and last trimmed size on the NSS pool named MYPOOL1.

## 6.24 Init Device

**init** <device\_name> [format] [shared|unshared]

Initialize a device by deleting all partitions on the device and setting the partitioning scheme.

```
nlvm [nlvm_options] init <device_name> [format] [shared|unshared]
```

You are automatically prompted to confirm the initialize action. Respond by typing *yes* or *no*, then press Enter. Use the `--no-prompt` NLVM option to suppress the confirmation prompt.

You can optionally specify whether to set the device as shared or unshared. If neither the `shared` nor `unshared` option is added, the device is initialized, the partitioning scheme is set, and the shared state remains what it was before the initialize command.

### Command Options

#### **device\_name**

Mandatory. Specify the name of the device to be initialized. This must be the first command option.

You can enter multiple devices by separating the device names with a comma and no spaces.

#### **Examples**

```
sdb  
sde, sdf, sdg
```

#### **format=<gpt|msdos>**

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

#### **Example**

```
format=msdos
```

#### **shared**

After initializing the device, the device is set as shared. A small partition is created on the device to store the shared setting. The remainder of the device is free space.

For example, use this option to mark a device as Shareable for Clustering if you plan to use it for a shared NSS pool. NSS looks for this setting to cluster enable the pool.

#### **unshared**

After initializing the device, the device is not marked as shared. The device is unpartitioned free space.

Use this option to remove all partitions from a device. For example, LVM requires that a device contains no partitions before it creates a volume group on it.

**-f, --force**

Optional. Specify this NLVM option to force the initialization. This option is required if the device contains a root (/), swap, or boot partition, or if the `init` command cannot delete any pools on the disk.

**Examples**

```
-f
--force
```

**--no-prompt**

Optional. Specify this NLVM option to prevent a confirmation message from being displayed.

**Example**

```
--no-prompt
```

**Command Examples**

```
nlvm --force init sdb
```

Force the initialization of a previously formatted device `/dev/sdb`, and set its partitioning scheme to use the default setting of `msdos`. If the device size is greater than 2TB, the device has only 2 TB of usable space. If the device was previously set as shared, the shared setting remains after the initialization. Otherwise, the device is unshared.

```
nlvm init sdd format=gpt unshared
```

Initialize the device `/dev/sdd`, and set its partitioning scheme to GPT. If the device was previously set as shared, this removes the shared setting from the device. The device is unpartitioned free space.

```
nlvm init sde format=gpt shared
```

Initialize the device `/dev/sde`, set its partitioning scheme to GPT, and mark the device as shared. The device contains a small partition to hold the shared setting, and the rest is free space.

```
nlvm --no-prompt init sde,sdf,sdg format=gpt unshared
```

Initialize multiple devices at a time. Set each device's partitioning scheme to GPT. If a device was previously set as shared, the `unshared` option removes its shared setting. The devices are each unpartitioned free space. The confirmation message is not displayed.

## 6.25 Label

```
label <partition_name> <"label text">
```

Modify or add a label to a Novell type partition (NSS, SBD, or RAID).

```
nlvm [nlvm_options] label <partition_name> <"label text">
```

**Command Options**

***partition\_name***

Mandatory. Specify the node name (such as `sd1.1`) of the partition. This must be the first command option.

### Example

```
sdc1.1
```

### "label text"

Mandatory. Specify the text word or phrase to use for the label. If the text has spaces, use quotation marks.

### Example

```
"This is the label"  
engineering
```

### Command Example

```
nlvm label sdc1.1 "This is the label"
```

Add the label "This is the label" to the /dev/sdc1.1 partition.

## 6.26 Linux Mount

### linux mount <lx\_volume\_name>

Mount a specified Linux POSIX volume on Linux. If the volume is NCP-enabled, this command also mounts the volume for NCP, and NCP assigns it a volume ID.

```
nlvm [nlvm_options] linux mount <lx_volume_name>
```

### Command Options

#### lx\_volume\_name

Mandatory. Specify the name of the Linux POSIX volume to mount.

Use the name format as it is displayed in NSSMU or with the `nlvm list volumes` command. For a non-LVM volume that is not NCP-enabled, specify the name as a forward slash with the name of the final directory of the mount point (`/<dir_name>`). For an LVM volume that is not NCP-enabled, specify the volume name of the LVM logical volume. For an NCP-enabled volume, specify the NCP name in all capital letters. Volume names are case sensitive.

### Examples

|          |  |
|----------|--|
| LV_VOL1  | [ex: an LVM volume that is NCP-enabled]        |
| lv_voll  | [ex: an LVM volume that is not NCP-enabled]    |
| HOME     | [ex: a non-LVM volume that is NCP-enabled]     |
| /home    | [ex: a non-LVM volume that is not NCP-enabled] |
| enabled] |  |

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

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

### Example

```
mntopt=rw
```

## Command Examples

```
nlvm linux mount LV_VOL1
```

Mounts the NCP-enabled LVM volume LV\_VOL1 in Linux using the parameters from the `/etc/fstab` file, and then mounts it in NCP. NCP automatically assigns a volume ID.

```
nlvm linux mount /home
```

Mounts the non-LVM volume using the parameters from the `/etc/fstab` file.

```
nlvm linux mount HOME
```

Mounts the NCP-enabled non-LVM volume in Linux using the parameters from the `/etc/fstab` file, and then mounts it in NCP. NCP automatically assigns a volume ID.

```
nlvm linux mount HOME mntopt=rw,user_xattr
```

Mounts the NCP-enabled non-LVM volume in Linux using the specified mount parameters for an Ext3 file system type, and then mounts the volume in NCP. NCP automatically assigns a volume ID.

## 6.27 Linux Unmount

```
linux unmount <lx_volume_name>
```

Dismount a specified Linux volume. If the volume is NCP-enabled, it also dismounts it from NCP before it dismounts it from Linux.

```
nlvm [nlvm_options] linux unmount <lx_volume_name>
```

### Command Option

***lx\_volume\_name***

Mandatory. Specify the name of the Linux POSIX volume to dismount.

Use the name format as it is displayed in NSSMU or with the `nlvm list volumes` command. For information, see `nlvm linux mount`.

### Examples

|         |  |
|---------|--|
| LV_VOL1 | [ex: an LVM volume that is NCP-enabled]        |
| lv_vol1 | [ex: an LVM volume that is not NCP-enabled]    |
| HOME    | [ex: a non-LVM volume that is NCP-enabled]     |
| /home   | [ex: a non-LVM volume that is not NCP-enabled] |

### Command Example

```
nlvm volume unmount HOME
```

Dismounts the NCP-enabled non-LVM volume HOME from NCP, and then dismounts it from Linux.

```
nlvm volume unmount /home
```

Dismounts the non-LVM volume /home from Linux.

```
nlvm volume unmount lv_vol1
```

Dismounts the LVM volume lv\_vol1 from Linux.

## 6.28 List Device

**list device** <device\_name>

Print the details of a specified device.

```
nlvm [nlvm_options] list device <device_name>
```

### Command Option

**device\_name**

Mandatory. Specify the desired device.

### Example

```
sdb
```

### Command Example

```
nlvm list device sdb
```

Print the details for the /dev/sdb device.

### Response Parameters

The device details include the following values. Most labels are self-explanatory.

| Label                         | Description  |
|-------------------------------|--|
| Name                          | Device name such as sdb or raid1   |
| Size                          | Total amount of space on the device in KB, MB, GB, or TB, and the number of whole sectors in that space<br><br>Size=623.91MB (1277773) |
| Used                          | Used space on the device in KB, MB, GB, or TB, and the number of whole sectors in that space.  |
| Free                          | Available space on the device in KB, MB, GB, or TB, and the number of whole sectors in that space                                      |
| Format                        | MSDOS, GPT, CSM (legacy EVMS Cluster Segment Manager), LVM (clustered Linux LVM volume), None (not initialized)                        |
| Shared                        | Yes or No; whether this device is marked as Shareable for Clustering   |
| RAID                          | Yes or No; whether this is an NSS software RAID device   |
| M:M                           | Major:Minor numbers, such as 8:112   |
| H:S                           | Heads:Sectors geometry per track, such as 255:32   |
| SectorSize (logical/physical) | Logical or physical sector size of a device  |
| Provisioning                  | Thick or Thin provisioning type  |

If the device contains partitions, it provides the following information:

| Label      | Description   |
|------------|---|
| Part       | Partition name such as <code>sdb1.1</code> , <code>sdc2</code> , or <code>cluster.sbd</code>  |
| Partition= |   |
| Type       | Partition type, including NSS, NSS RAID, SBD, Linux, Linux_swap, LVM                          |
| Size       | Amount of space allocated to the partition in KB, MB, GB, or TB                               |
| Sectors    | Number of whole sectors allocated to the partition  |
| Pool       | If the partition is the NSS type, the name of the pool that resides on the partition (if any) |

For NSS software RAID devices, it provides the following information:

| Label   | Description   |
|---------|---|
| RAID    | No or RAID type (0, 1, or 5)                                      |
| Sync    | Yes or %; whether the RAID is in sync or if a sync is in progress |
| Segs    | Number of segments defined for the RAID                           |
| Enbl    | Yes or No; whether the RAID is enabled on this node               |
| Missing | Segment number (if any) that is missing in the RAID               |
| Stripe  | RAID stripe size in bytes (typically KB) for RAID types 0 and 5   |

For RAID segments, it provides the following information:

| Label   | Description  |
|---------|--|
| Segment | Segment index number   |
| Name    | Segment name, such as <code>sdb1.4</code>                              |
| Device  | Name of the device that contains the segment, such as <code>sdb</code> |
| Size    | Segment size in KB, MB, GB, or TB                                      |
| Sectors | Number of whole sectors allocated to the partition                     |

## Sample Command Responses

### Sample 1: Standard Device

```
nlvm list device sdf
Name=sdf
Size=1.63TB(3515088896) Used=801.00GB(1679818816) Free=875.12GB(1835269816)
Format=GPT Shared=Yes RAID=No M:M=8:80 H:S=255:32 SectorSize(logical/
physical)=4096b/4096b Provisioning=Thin
Partitions on the device:
Part  Type          Size      Sectors  Pool
sdf1.1 NSS            500.00GB  104857600  DATA
sdf1.2 NSS            300.00GB   629145600  WEB
sdf1.3 NSS             1.00GB    2097152    DATA3
```

## Sample 2: NSS RAID 1 (Mirror) Device (not initialized)

```
nlvm list device RD1
Name=RD1
Size=1023.98MB(2097120) Used=1023.98MB(2097120) Free=0KB(0)
Format=None Shared=Yes RAID=1 Sync=NoEbl M:M=254:3 H:S=255:32
SectorSize(logical/physical)=4096b/4096b Provisioning=Thin
Segs=2 Enbl=No
Segments of the RAID:
Segment Name Device Size Sectors Sync
0 sde1.1 sde 1.00GB 2097152 NoEbl
1 sdf1.3 sdf 1.00GB 2097152 NoEbl
```

## Sample 3: NSS RAID 0 Device

```
nlvm list device MYRAID0
Name=MYRAID0
Size=15.00GB(31457280) Used=10.00GB(20971552) Free=4.99GB(10485728)
Format=MSDOS Shared=Yes RAID=0 Sync=NoEbl M:M=254:4 H:S=255:32
SectorSize(logical/physical)=512b/512b Provisioning=Thick
Segs=3 Enbl=No Missing=None Stripe=64k
Segments of the RAID:
Segment Name Device Size Sectors
0 sdj1.1 sdj 5.00GB 10485792
1 sdi1.1 sdi 5.00GB 10485792
2 sdh1.1 sdh 5.00GB 10485792
Partitions on the device:
Part Type Size Sectors Pool
MYRAID0p1.1 NSS 10.00GB 20971520 DATA4
```

## 6.29 List Devices

### list devices [exclude] [more|all]

Print a list of the devices. For each device, display the device name, size, free available space, partitioning type, if it is marked as Shareable for Clustering, and if it is an NSS software RAID device. If no other options are specified, this prints a list of all devices and software RAID devices.

```
nlvm [-t] list devices [exclude] [more|all]
```

### Command Options

**exclude=<raid|nonraid|shared|nonshared|lvm|nonlvm>**

Exclude the specified type of devices. This option can be used multiple times to add exclusions for different types. Valid device types are `raid`, `nonraid`, `shared`, `nonshared`, `lvm`, or `nonlvm`.

### Example

```
exclude=raid exclude=nonshared
```

### -t, --terse

Use this NLVM option to format the output for parsing.

### more

Prints more information than appears in the standard output. It can be used with or without the `-t` NLVM option.

### Example

```
more
```



## all

Prints detailed information about each of the devices. This is the same information that is printed for the `nlvm list device <device_name>` command. It can be used with or without the `-t NLVM` option.

### Example

```
all
```

### Command Example

```
nlvm list devices exclude=raid exclude=nonshared exclude=lvm all
```

Print detailed information for all non-LVM shared devices that are not software RAID devices.

### Response Parameters

You can issue the commands with the `--terse NLVM` option to output the same information in a format that is more easily parsed.

### Standard Output

The command returns the following standard information about the devices on the server:

| Label   | Description   |
|---------|---|
| Name    | Device name such as <code>sdb</code> or <code>raid1</code>  |
| Size    | Total amount of space on the device in KB, MB, GB, or TB  |
| Used    | Used space on the device in KB, MB, GB, or TB   |
| Free    | Available space on the device in KB, MB, GB, or TB  |
| Format  | MSDOS, GPT, CSM (legacy EVMS Cluster Segment Manager), LVM (clustered Linux LVM volume), None (not initialized) |
| Shared  | Yes or No; whether this device is marked as Shareable for Clustering  |
| RAID    | Yes or No; whether this is an NSS software RAID device  |
| Enabled | Yes or No; whether the RAID is enabled on this node   |

### More Output

The command returns the following additional information about the devices on the server:

| Label                         | Description   |
|-------------------------------|---|
| RAID                          | Type (0, 1, or 5) or No; type of NSS RAID device, or not a RAID |
| Sync                          | Yes or %; whether the RAID is in sync or percent completed      |
| M:M                           | Major:Minor numbers, such as 8:112                              |
| SectorSize(logical /physical) | Logical or physical sector size of a device                     |
| Provisioning                  | Thick or Thin provisioning type                                 |

## All Output

If the `all` option is used, the command returns the same information about each device as is displayed for the `nlvm list device <device_name>` command. This includes information about its partitions, or about its partitions and segments for RAID devices.

### Sample Command Responses

#### Sample 1: nlvm list devices

| Name    | Size      | Used      | Free     | Format | Shared | RAID | Enabled |
|---------|-----------|-----------|----------|--------|--------|------|---------|
| sda     | 1.63TB    | 1.63TB    | 0KB      | MSDOS  | No     | No   |         |
| sdb     | 931.00GB  | 206.00GB  | 724.99GB | GPT    | No     | No   |         |
| sdc     | 931.00GB  | 699.99GB  | 231.00GB | MSDOS  | No     | No   |         |
| sdd     | 1.63TB    | 1.63TB    | 0KB      | GPT    | No     | No   |         |
| sde     | 1.63TB    | 1.00GB    | 1.63TB   | MSDOS  | Yes    | No   |         |
| sdf     | 1.63TB    | 801.00GB  | 875.12GB | GPT    | Yes    | No   |         |
| MYRAID0 | 15.00GB   | 10.00GB   | 4.99GB   | MSDOS  | Yes    | 0    | No      |
| RD1     | 1023.98MB | 1023.98MB | 0KB      | None   | Yes    | 1    | No      |

#### Sample 2: nlvm list devices --terse

```
Name=sda Size=1.63TB Used=1.63TB Free=0KB Format=MSDOS Shared=No RAID=No
Name=sdb Size=931.00GB Used=206.00GB Free=724.99GB Format=GPT Shared=No
RAID=No
Name=sdc Size=931.00GB Used=699.99GB Free=231.00GB Format=MSDOS Shared=No
RAID=No
Name=sdd Size=1.63TB Used=1.63TB Free=0KB Format=GPT Shared=No RAID=No
Name=sde Size=1.63TB Used=1.00GB Free=1.63TB Format=MSDOS Shared=Yes RAID=No
Name=sdf Size=1.63TB Used=801.00GB Free=875.12GB Format=GPT Shared=Yes
RAID=No
Name=MYRAID0 Size=15.00GB Used=10.00GB Free=4.99GB Format=MSDOS Shared=Yes
RAID=0 Enabled=No
Name=RD1 Size=1023.98MB Used=1023.98MB Free=0KB Format=None Shared=Yes
RAID=1 Enabled=No
```

#### Sample 3: nlvm list devices more

| Name                          | Size      | Used         | Free     | Format | Shared | RAID | Sync  | Maj:Min |
|-------------------------------|-----------|--------------|----------|--------|--------|------|-------|---------|
| SectorSize (logical/physical) |           | Provisioning |          |        |        |      |       |         |
| sda                           | 1.63TB    | 1.63TB       | 0KB      | MSDOS  | No     | No   |       | 8:0     |
| 4096b/4096b                   |           | Thick        |          |        |        |      |       |         |
| sdb                           | 931.00GB  | 206.00GB     | 724.99GB | GPT    | No     | No   |       | 8:16    |
| 512b/512b                     |           | Thick        |          |        |        |      |       |         |
| sdc                           | 931.00GB  | 699.99GB     | 231.00GB | MSDOS  | No     | No   |       | 8:32    |
| 512b/512b                     |           | Thick        |          |        |        |      |       |         |
| sdd                           | 1.63TB    | 1.63TB       | 0KB      | GPT    | No     | No   |       | 8:48    |
| 4096b/4096b                   |           | Thin         |          |        |        |      |       |         |
| sde                           | 1.63TB    | 1.00GB       | 1.63TB   | MSDOS  | Yes    | No   |       | 8:64    |
| 4096b/4096b                   |           | Thin         |          |        |        |      |       |         |
| sdf                           | 1.63TB    | 801.00GB     | 875.12GB | GPT    | Yes    | No   |       | 8:80    |
| 4096b/4096b                   |           | Thin         |          |        |        |      |       |         |
| sdg                           | 7.99GB    | 16KB         | 7.99GB   | MSDOS  | Yes    | No   |       | 8:96    |
| 512b/512b                     |           | Thick        |          |        |        |      |       |         |
| sdh                           | 7.99GB    | 5.00GB       | 2.99GB   | MSDOS  | Yes    | No   |       | 8:112   |
| 512b/512b                     |           | Thin         |          |        |        |      |       |         |
| sdi                           | 7.99GB    | 5.00GB       | 2.99GB   | MSDOS  | Yes    | No   |       | 8:128   |
| 512b/512b                     |           | Thick        |          |        |        |      |       |         |
| sdj                           | 8.40GB    | 5.00GB       | 3.40GB   | MSDOS  | Yes    | No   |       | 8:144   |
| 512b/512b                     |           | Thin         |          |        |        |      |       |         |
| MYRAID0                       | 15.00GB   | 10.00GB      | 4.99GB   | MSDOS  | Yes    | 0    | NoEbl | 254:4   |
| 512b/512b                     |           | Thin         |          |        |        |      |       |         |
| RD1                           | 1023.98MB | 1023.98MB    | 0KB      | None   | Yes    | 1    | NoEbl | 254:3   |
| 4096b/4096b                   |           | Thick        |          |        |        |      |       |         |

#### Sample 4: nlvm list devices all

```
Name=sda
Size=1.63TB(3515088896) Used=1.63TB(3515088896) Free=0KB(0)
Format=MSDOS Shared=No RAID=No M:M=8:0 H:S=255:32 SectorSize(logical/
physical)=4096b/4096b Provisioning=Thick
Partitions on the device:
  Part  Type          Size      Sectors  Pool
  sda1.1 NSS          1.63TB  3515083168  POOL1

Name=sdb
Size=931.00GB(1952448512) Used=206.00GB(432021568)
Free=724.99GB(1520426911)
Format=GPT Shared=No RAID=No M:M=8:16 H:S=255:32 SectorSize(logical/
physical)=512b/512b Provisioning=Thick
Partitions on the device:
  Part  Type          Size      Sectors  Pool
  sdb2  Linux          5.00GB   10491904
  sdb3  Linux          1.00GB    2103296
  sdb1  Linux        199.99GB 419426304

Name=sdc
Size=931.00GB(1952448512) Used=699.99GB(1467987968)
Free=231.00GB(484460544)
Format=MSDOS Shared=No RAID=No M:M=8:32 H:S=255:32 SectorSize(logical/
physical)=512b/512b Provisioning=Thick
Partitions on the device:
  Part  Type          Size      Sectors  Pool
  sdc1  Linux LVM       99.99GB  209711104
  sdc2  Linux LVM      199.99GB 419424256
  sdc3  Linux LVM      199.99GB 419424256
  sdc4  Linux LVM      199.99GB 419426304

Name=sdd
Size=1.63TB(3515088896) Used=1.63TB(3515088856) Free=0KB(0)
Format=GPT Shared=No RAID=No M:M=8:48 H:S=255:32 SectorSize(logical/
physical)=4096b/4096b Provisioning=Thin
Partitions on the device:
  Part  Type          Size      Sectors  Pool
  sdd1  Linux LVM      100.01GB 209741824
  sdd2  Linux LVM      100.01GB 209745920
  sdd3  Linux LVM      100.01GB 209743872
  sdd4  Linux LVM      100.01GB 209745920
  sdd5  Linux LVM      100.01GB 209743872
  sdd6  Linux LVM      100.01GB 209743872
  sdd7  Linux LVM      100.01GB 209745920
  sdd8  Linux LVM      100.01GB 209743872
  sdd9  Linux          876.01GB 1837129728

Name=sde
Size=1.63TB(3515088896) Used=1.00GB(2097184) Free=1.63TB(3512991712)
Format=MSDOS Shared=Yes RAID=No M:M=8:64 H:S=255:32 SectorSize(logical/
physical)=4096b/4096b Provisioning=Thin
Partitions on the device:
  Part  Type          Size      Sectors  Pool
  sde1.1 NSS          1.00GB   2097152  DATA3

Name=sdf
Size=1.63TB(3515088896) Used=801.00GB(1679818816) Free=875.12GB(1835269816)
Format=GPT Shared=Yes RAID=No M:M=8:80 H:S=255:32 SectorSize(logical/
physical)=4096b/4096b Provisioning=Thin
Partitions on the device:
  Part  Type          Size      Sectors  Pool
  sdf1.1 NSS          500.00GB 1048576000  DATA
  sdf1.2 NSS          300.00GB 629145600  WEB
  sdf1.3 NSS          1.00GB   2097152  DATA3

Name=MYRAID0
Size=15.00GB(31457280) Used=10.00GB(20971552) Free=4.99GB(10485728)
Format=MSDOS Shared=Yes RAID=0 Sync=NoEbl M:M=254:4 H:S=255:32
SectorSize(logical/physical)=512b/512b Provisioning=Thin
Segs=3 Enbl=No Missing=None Stripe=64k
Segments of the RAID:
  Segment Name Device      Size      Sectors
  0          sdj1.1 sdj          5.00GB   10485792
```

```

    1   sdi1.1 sdi      5.00GB 10485792
    2   sdh1.1 sdh      5.00GB 10485792
Partitions on the device:
Part   Type              Size   Sectors  Pool
MYRAID0p1.1 NSS              10.00GB 20971520 DATA4

Name=RD1
Size=1023.98MB(2097120) Used=1023.98MB(2097120) Free=0KB(0)
Format=None Shared=Yes RAID=1 Sync=NoEbl M:M=254:3 H:S=255:32
SectorSize(logical/physical)=4096b/4096b Provisioning=Thick
Segs=2 Enbl=No
Segments of the RAID:
Segment Name Device      Size   Sectors  Sync
0   sde1.1 sde      1.00GB 2097152 NoEbl
1   sdf1.3 sdf      1.00GB 2097152 NoEbl

```

## 6.30 List Linux Volume

**list linux volume** <lx\_volume\_name>

Print detailed information about a specified Linux volume.

```
nlvm [nlvm_options] list linux volume <lx_volume_name>
```

### Command Option

**lx\_volume\_name**

Mandatory. Specify the name of the Linux POSIX volume.

Use the name format as it is displayed in NSSMU or with the `nlvm list volumes` command. For a non-LVM that is not NCP-enabled, specify the name as a forward slash with the name of the final directory of the mount point (`/<dir_name>`). For an LVM volume that is not NCP-enabled, specify the volume name of the LVM logical volume. For an NCP-enabled volume, specify the NCP name.

### Examples

```

LV_VOL1           [ex: an LVM volume that is NCP-enabled]
lv_vol1          [ex: an LVM volume that is not NCP-enabled]
MYLVMVOL         [ex: an LVM volume that is not NCP-enabled]
HOME             [ex: a non-LVM volume that is NCP-enabled]
/home           [ex: a non-LVM volume that is not NCP-
enabled]

```

### Command Example

```
nlvm list linux volume MYLVMVOL
```

Print detailed information about the NCP-enabled LVM volume named MYLVMVOL.

### Response Parameters

The Linux volume details include the following. Most labels are self-explanatory.

| Label   | Description  |
|---------|--|
| Name    | Volume name. The format of the name depends on the type of volume and whether it is NCP-enabled. |
| Group   | LVM group name or NA (not applicable) for non-LVM volumes  |
| Mounted | Yes or No; whether the volume is mounted for user access   |

| Label        | Description   |
|--------------|---|
| Size         | Size of the volume in KB, MB, GB, or TB   |
| Shared       | Yes or No; whether volume's device is marked as Shareable for Clustering  |
| Type         | Type of file system (such as btrfs, ext2, ext3, ext4, reiserfs, or xfs)   |
| LVM          | Yes or No; whether the volume is an LVM volume  |
| NCP          | Yes or No; whether the volume is NCP-enabled  |
| Mountpoint   | Full Linux path where the volume is mounted   |
| Path         | Path of the device or partition. For LVM, this is typically <code>/dev/&lt;volume_group_name&gt;/&lt;logical_volume_name&gt;</code> . If it is not LVM, this is the partition path. |
| MountOptions | Defaults or specified mount options, such as <code>rw</code>  |

## Sample Command Responses

### Sample 1: Non-LVM Volume

```
nlvm list linux volume /home
Name=/home
Group=NA Mounted=Yes Size=3.00GB
Shared=No Type=ext4 LVM=No NCP=No
Mountpoint=/home
Path=/dev/sda3
MountOptions=defaults
```

### Sample 2: LVM Volume

```
nlvm list linux volume mylvm
Name=mylvm
Group=ajlvm Mounted=No Size=100.00MB
Shared=No Type=ext4 LVM=Yes NCP=No
Mountpoint=/usr/novell/mylvm
Path=/dev/mylvm/mylvm
MountOptions=rw
```

### Sample 3: NCP-Enabled Non-LVM Volume

```
nlvm list linux volume NCP3
Name=NCP3
Group=NA Mounted=Yes Size=103.59MB
Shared=No Type=ext4 LVM=No NCP=Yes
Mountpoint=/usr/novell/NCP3
Path=/dev/sdc3
MountOptions=rw
```

### Sample 4: NCP-Enabled LVM Volume

```
nlvm list linux volume LVMNCP
Name=LVMNCP
Group=lvmncp Mounted=No Size=100.00MB
Shared=No Type=ext4 LVM=Yes NCP=Yes
Mountpoint=/usr/novell/lvmncp2
Path=/dev/lvmncp/LVMNCP
MountOptions=rw
```

## 6.31 List Linux Volumes

**list [-t] linux volumes [more|all]**

Print a list of Linux POSIX volumes and for each, display its path, mount point, file system type, NCP enabled status, and mount status.

```
nlvm [-t] list linux volumes [more|all]
```

### Command Options

**-t, --terse**

Use this NLVM option to format the output for parsing.

**more**

Prints more information than appears in the standard output. It can be used with or without the `-t` NLVM option.

### Example

```
more
```

**all**

Prints detailed information about each of the Linux volumes. This is the same information that is printed for the `nlvm list linux volume <volume_name>` command. It can be used with or without the `-t` NLVM option.

### Example

```
all
```

### Command Example

```
nlvm list linux volumes
```

Print a list of Linux POSIX volumes and the paths where they are mounted.

### Response Parameters

You can issue the commands with the `--terse` NLVM option to output the same information in a format that is more easily parsed.

## Standard Output

The command returns the following standard information about the Linux volumes on the server:

| Label      | Description  |
|------------|--|
| Name       | Volume name. The format of the name depends on the type of volume and whether it is NCP-enabled. |
| Group      | LVM group name or NA (not applicable) for non-LVM volumes  |
| Mounted    | Yes or No; whether the volume is mounted for user access   |
| Size       | Size of the volume in KB, MB, GB, or TB  |
| Shared     | Yes or No; whether volume's device is marked as Shareable for Clustering                         |
| Type       | Type of file system (such as btrfs, ext2, ext3, reiserfs, or xfs)                                |
| LVM        | Yes or No; whether the volume is an LVM volume   |
| NCP        | Yes or No; whether the volume is NCP-enabled   |
| Mountpoint | Full Linux path where the volume is mounted  |

## More Output

The command returns the following additional information about the Linux volumes on the server:

| Label | Description   |
|-------|---|
| Path  | Path of the device or partition. For LVM, this is typically <code>/dev/&lt;volume_group_name&gt;/&lt;logical_volume_name&gt;</code> . If it is not LVM, this is the partition path. |

## All Output

If the `all` option is used, the command returns the same information about each Linux volume as is displayed for the `nlvm list volume <volume_name>` command.

## Sample Command Responses

### Sample 1: nlvm list linux volumes

```
Name      Group  Mounted  Size Shared  Type LVM NCP Mountpoint
/         /      Yes      15.98GB No      ext3 No  No  /
/home    /home  Yes      3.00GB  No      ext3 No  No  /home
mylvm    mylvm  No       100.00MB No      ext3 Yes No  /usr/novell/mylvm
LVMNCP   lvmncp No       100.00MB No      ext3 Yes Yes /usr/novell/lvmncp2
NCP3     /      Yes      103.59MB No      ext3 No  Yes /usr/novell/NCP3
```

### Sample 2: nlvm list linux volumes --terse

```
Name=/ Group=NA Mounted=Yes Size=15.98GB Shared=No Type=ext3 LVM=No NCP=No
Mountpoint=/
Name=/home Group=NA Mounted=Yes Size=3.00GB Shared=No Type=ext3 LVM=No
NCP=No Mountpoint=/home
Name=mylvm Group=mylvm Mounted=No Size=100.00MB Shared=No Type=ext3 LVM=Yes
NCP=No Mountpoint=/usr/novell/mylvm
Name=LVMNCP Group=lvmnpc Mounted=No Size=100.00MB Shared=No Type=ext3
LVM=Yes NCP=Yes Mountpoint=/usr/novell/lvmncp2
Name=NCP3 Group=NA Mounted=Yes Size=103.59MB Shared=No Type=ext3 LVM=No
NCP=Yes Mountpoint=/usr/novell/NCP3
```

### Sample 3: nlvm list linux volumes more

| Name              | Group  | Mounted | Size     | Shared | Type | LVM | NCP | Mountpoint          | Path |
|-------------------|--------|---------|----------|--------|------|-----|-----|---------------------|------|
| /                 |        | Yes     | 15.98GB  | No     | ext3 | No  | No  | /                   | /    |
| dev/sda2          |        |         |          |        |      |     |     |                     |      |
| /home             |        | Yes     | 3.00GB   | No     | ext3 | No  | No  | /home               | /    |
| dev/sda3          |        |         |          |        |      |     |     |                     |      |
| mylvm             | mylvm  | No      | 100.00MB | No     | ext3 | Yes | No  | /usr/novell/mylvm   | /    |
| dev/mylvm/mylvm   |        |         |          |        |      |     |     |                     |      |
| LVMNCP            | lvmnpc | No      | 100.00MB | No     | ext3 | Yes | Yes | /usr/novell/lvmncp2 | /    |
| dev/lvmncp/LVMNCP |        |         |          |        |      |     |     |                     |      |
| NCP2              |        | Yes     | 103.59MB | No     | ext3 | No  | Yes | /usr/novell/NCP3    | /    |
| dev/sdc3          |        |         |          |        |      |     |     |                     |      |

### Sample 4: nlvm list linux volumes all

```
Name=/
Group=NA Mounted=Yes Size=15.98GB
Shared=No Type=ext3 LVM=No NCP=No
Mountpoint=/
Path=/dev/sda2
MountOptions=acl,user_xattr

Name=/home
Group=NA Mounted=Yes Size=3.00GB
Shared=No Type=ext3 LVM=No NCP=No
Mountpoint=/home
Path=/dev/sda3
MountOptions=defaults

Name=mylvm
Group=mylvm Mounted=No Size=100.00MB
Shared=No Type=ext3 LVM=Yes NCP=No
Mountpoint=/usr/novell/mylvm
Path=/dev/mylvm/mylvm
MountOptions=rw

Name=LVMNCP
Group=lvmnpc Mounted=No Size=100.00MB
Shared=No Type=ext3 LVM=Yes NCP=Yes
Mountpoint=/usr/novell/lvmncp2
Path=/dev/lvmncp/LVMNCP
MountOptions=rw

Name=NCP3
Group=NA Mounted=Yes Size=103.59MB
Shared=No Type=ext3 LVM=No NCP=Yes
Mountpoint=/usr/novell/NCP3
Path=/dev/sdc3
MountOptions=rw
```

## 6.32 List Move

**list move <<move\_name>|<pool\_name>>**

Print detailed information about a specified NSS pool move. It lists the devices you are moving from and the devices you are moving to, such as



```
from=sdc,sdd,sde to=sdg
```

If a pool is cluster-enabled, the pool move is enabled and active only on the node where the pool cluster resource is currently online. On other nodes in the cluster, the pool move is not enabled.

```
nlvm [nlvm_options] list move <<move_name>|<pool_name>>
```

The move occurs as a low-level block mirror between the original location and the new location. The entire pool area is mirrored. The response reports the number of mirror regions to be moved for the pool relative to the maximum source pool size, which is unrelated to the NSS blocks in use. The region count for the old pool location does not change during the move. The complete parameter indicates the number of regions that have been moved so far and the percentage that it represents of the total number of regions to be moved.

The size of a mirror region is determined internally based on the total size of the mirror. One sector is used to track the number of mirror regions that are currently synchronized. A bit represents a mirror region, and there are 4096 bits total (512 \* 8) to track. A shift technique is used so that the mirror region size is always a power of 2 (128, 256, 512, and so on) and the total number of regions to move is less than or equal to 4096. Except for very small mirrors, the number of mirror regions is usually between 2048 and 4096. The minimum mirror region size used is 64 sectors (32 KB). There is no maximum. For an 8 TB pool, the mirror region size is 2 GB. When a complete region is mirrored, the bit is set. If a region is partially mirrored during a system failure or cluster resource migration, the entire region is remirrored when mirroring resumes.

The response lists the set of devices that are being used for the original location (from) and the new location (to).

### Command Option

#### *move\_name* or *pool\_name*

Mandatory. Specify the name of the move, such as `POOLNAME_move`. You can alternatively specify the pool name.

#### Example

```
MYPOOL_move
```

### Command Example

```
nlvm list move MYPOOL_move
```

Print detailed information about the `MYPOOL_move` move.

### Response Parameters

The command returns the following information about the specified pool move:

| Label    | Description  |
|----------|--|
| Name     | Name of the move. Typically, <code>&lt;pool_name&gt;_move</code> .   |
| Pool     | Name of the pool being moved   |
| FromStat | Status of the “from” group of devices that make up the source pool (Active, ReadError, WriteError, Missing, NotEnabled)            |
| ToStat   | Status of the “to” group of devices that make up the new instance of the pool (Active, ReadError, WriteError, Missing, NotEnabled) |

| Label           | Description   |
|-----------------|---|
| Complete        | Percent complete  |
| OldSize         | Size of the old/source pool in MB, GB, or TB, and the number of whole sectors in that space |
| From            | From set of devices for the pool being moved  |
| To              | To set of devices for the pool in its new location  |
| M:M             | Major:minor numbers of the move object  |
| Regions         | Total number of mirror regions to be moved  |
| RegionsComplete | Number of mirror regions that are complete  |

### Sample Command Responses

#### Sample 1: nlvm list move MYPOOL\_move

```
Name=MYPOOL_move
Pool=MYPOOL FromStat=Active ToStat=Active Complete=100%
OldSize=99.00MB(202752) From=sdb To=sdc
M:M=253:21 Regions=3168 RegionsComplete=3168
```

#### Sample 2: Cluster Node where the Pool Cluster Resource Is Active

```
Name=CLUSPOOL_move
Pool=CLUSPOOL FromStat=Active ToStat=Active Complete=33%
OldSize=7.19GB(15087616) From=sdc,sdd,sde,sdf To=sdh
M:M=253:21 Regions=3684 RegionsComplete=1245
```

#### Sample 3: Any Cluster Node where the Pool Cluster Resource Is Not Active

```
Name=CLUSPOOL_move
Pool=CLUSPOOL FromStat=NotEnabled ToStat=NotEnabled
From=sdc,sdd,sde,sdf To=sdh
Move is not enabled on this node.
```

## 6.33 List Moves

### list moves [more|all]

Print a list of current NSS pool moves.

If a pool is cluster-enabled, the pool move is enabled and active only on the node where the pool cluster resource is currently online. On other nodes in the cluster, the pool move is not enabled.

```
nlvm [-t] list moves [more|all]
```

#### Command Options

**-t, --terse**

Use this NLVM option to format the output for parsing.

**more**

Prints more information than appears in the standard output. It can be used with or without the `-t` NLVM option.

## Example

```
more
```

## all

Prints detailed information about each of the pool moves. This is the same information that is printed for the `nlvm list moves <move_name>` command. It can be used with or without the `-t` NLVM option.

## Example

```
all
```

## Command Example

```
nlvm list moves
```

Print a list of NSS pool moves that are in progress now.

## Response Parameters

You can issue the commands with the `--terse` NLVM option to output the same information in a format that is more easily parsed.

## Standard Output

The command returns the following standard information about the pool moves on the server:

| Label    | Description  |
|----------|--|
| Name     | Name of the move. Typically, <code>&lt;pool_name&gt;_move</code> .   |
| Pool     | Name of the pool being moved   |
| FromStat | Status of the “from” group of devices that make up the source pool (Active, ReadError, WriteError, Missing, NotEnabled)            |
| ToStat   | Status of the “to” group of devices that make up the new instance of the pool (Active, ReadError, WriteError, Missing, NotEnabled) |
| Complete | Percent complete   |

## More Output

The command returns the following additional information about the pool moves on the server:

| Label   | Description   |
|---------|---|
| OldSize | Size of the old/source pool in MB, GB, or TB, and the number of whole sectors in that space |
| From    | From set of devices for the pool being moved  |
| To      | To set of devices for the pool in its new location  |

## All Output

If the `all` option is used, the command returns the same information about each pool move as is displayed for the `nlvm list move <move_name>` command.

### Sample Command Responses

#### Sample 1: Server with No Active Moves

```
nlvm list moves
No moves
```

#### Sample 2: nlvm list moves

```
Name          Pool    FromStat  ToStat  Complete
MYPOOL_move  MYPOOL  Active    Active  100%
```

#### Sample 3: nlvm list moves --terse

```
Name=PMOVE_move Pool=PMOVE FromStat=Active ToStat=Active Complete=100%
```

#### Sample 4: nlvm list moves more

```
Name          Pool    FromStat  ToStat  Complete  OldSize  From  To
MYPOOL_move  MYPOOL  Active    Active  100%     99.00MB  sdb   sdc
```

#### Sample 5: nlvm list moves all

```
Name=MYPOOL_move
Pool=MYPOOL FromStat=Active ToStat=Active Complete=100%
OldSize=99.00MB(202752) From=sdb To=sdc
M:M=253:21 Regions=3168 RegionsComplete=3168
```

#### Sample 6: Cluster Node where the Pool Cluster Resource Is Active

```
Name          Pool    FromStat  ToStat  Complete
MYPOOL_move  MYPOOL  Active    Active  71%
```

#### Sample 7: Any Cluster Node where the Pool Cluster Resource Is Not Active

```
Name          Pool    FromStat  ToStat  Complete
MYPOOL_move  MYPOOL  NotEnabled NotEnabled 0%
Move is not enabled on this node.
```

## 6.34 List Partition

### `list partition <partition_name>`

Print detailed information about a specified partition.

```
nlvm [nlvm_options] list partition <partition_name>
```

#### Command Option

##### *partition\_name*

Mandatory. Specify the node name (such as `sdcl.1`) for the partition.

#### Example

```
sdcl.1
```

#### Command Example

```
nlvm list partition sdcl.1
```

Print detailed information about the `/dev/sdcl.1` partition.

## Response Parameters

The command returns the following information about the specified partition:

| Label  | Description  |
|--------|--|
| Name   | Name of the partition  |
| Type   | Partition type in both hex and type name if known  |
| Start  | Starting sector of the partition   |
| Size   | Size of the partition in MB, GB, or TB, and the number of whole sectors that consist in that space |
| Device | Device the partition is on, such as <code>sda</code> or <code>raid1</code>                         |
| Shared | Whether the partition is marked Shareable for Clustering (1, 0)                                    |
| M:M    | Major:minor numbers of the partition (if applicable)   |
| Pool   | Name of the NSS pool using this partition (if applicable)  |
| Label  | Label for SBD partition (if applicable). Typically, the same as the cluster name.                  |

## Sample Command Responses

### Sample 1: Linux Swap Partition

```
nlvm list partition sda1  
  
Name=sda1  
Type=82(Linux Swap) Start=2048 Size=1.00GB(2103296)  
Device=sda Shared=No M:M=8:1 Pool=None
```

### Sample 2: NSS Pool Partition

```
nlvm list partition sdd1.1  
  
Name=sdd1.1  
Type=169(NSS) Start=32 Size=1023.96MB(2097088)  
Device=sdd Shared=No M:M=0:0 Pool=TEST2
```

### Sample 3: NSS RAID Partition

```
nlvm list partition sde1.2  
  
Name=sde1.2  
Type=1CF(NSS_Raid) Start=204832 Size=100.01MB(204832)  
Device=sde Shared=No M:M=0:0 Pool=None
```

### Sample 4: NSS Pool Snapshot Partition

```
nlvm list partition sdi6.1  
  
Name=sdi6.1  
Type=1AC(Snapshot) Start=206880 Size=75.00MB(153600)  
Device=sdi Shared=No M:M=253:17 Pool=SNAP1
```

### Sample 5: OES Cluster Services SBD Partition

```
nlvm list partition clstr.sbd

Name=clstr.sbd
  Type=1AD(Cluster) Start=32 Size=100.00MB(204800)
  Device=sde Shared=No M:M=253:4 Pool=None
  Label: clstr
```

### Sample 6: Linux Partition

```
nlvm list partition sdc1

Name=sdc1
  Type=83(Linux) Start=32 Size=103.57MB(212128)
  Device=sdc Shared=No M:M=8:33 Pool=None
```

### Sample 7: Linux LVM Partition

```
nlvm list partition sdc2

Name=sdc2
  Type=8E(Linux_LVM) Start=212160 Size=103.59MB(212160)
  Device=sdc Shared=No M:M=8:34 Pool=None
```

### Sample 8: DOS Extended Partition

```
nlvm list partition sdc4

Name=sdc4
  Type=5(DOS_Extended) Start=530400 Size=765.00MB(1566720)
  Device=sdc Shared=No M:M=8:36 Pool=None
```

## 6.35 List Partitions

### **list partitions [device] [mask] [more|all]**

Print a list of partitions based on the options. If no command options are specified, all data partitions are listed.

```
nlvm [-t] list partitions [device] [mask] [more|all]
```

#### **Command Options**

##### **device=device\_name**

Print a list of the partitions on the specified device.

##### **Example**

```
device=sdb
```

##### **mask=<free|all|nss|nssfree>**

Print a list of the partitions that meet the specified mask option.

##### **Mask Options**

##### **free**

Print a list of only the free space partitions.

**all**

Print a list of both data and free space partitions.

**nss**

Print a list of only NSS type partitions.

**nssfree**

Print a list of free space that can be used to create NSS partitions. This option combines contiguous free space together to give a true view of available space.

**Example**

```
mask=nss
```

**-t, --terse**

Use this NLVM option to format the output for parsing.

**more**

Prints more information than appears in the standard output. It can be used with or without the `-t` NLVM option.

**Example**

```
more
```

**all**

Prints detailed information about each of the partitions. This is the same information that is printed for the `nlvm list partition <partition_name>` command. It can be used with or without the `-t` NLVM option.

**Example**

```
all
```

**Command Example**

```
nlvm list partitions device=sdb mask=nss
```

Print a list of partitions of type nss on the `/dev/sdb` device.

**Response Parameters**

You can issue the commands with the `--terse` NLVM option to output the same information in a format that is more easily parsed.

**Standard Output**

The command returns the following information about the partitions on the server:

| Label  | Description  |
|--------|--|
| Name   | Name of the partition  |
| Type   | Partition type in both hex and type name if known                          |
| Start  | Starting sector of the partition   |
| Size   | Size of the partition in MB, GB, or TB                                     |
| Device | Device the partition is on, such as <code>sda</code> or <code>raid1</code> |

## More Output

The command returns the following additional information about the partitions on the server:

| Label  | Description   |
|--------|---|
| Shared | Whether the partition is marked Shareable for Clustering (1, 0) |
| M:M    | Major:minor numbers of the partition (if applicable)            |
| Pool   | Name of the NSS pool using this partition (if applicable)       |

## All Output

If the `all` option is used, the command returns the same information about each partition as is displayed for the `nlvm list partition <partition_name>` command.

## Sample Command Responses

### Sample 1: nlvm list partitions

```
nlvm list partitions
```

| Name        | Type             | Start    | Size     | Device |
|-------------|------------------|----------|----------|--------|
| sda1        | 82 (Linux_Swap)  | 2048     | 1.00GB   | sda    |
| sda2        | 83 (Linux)       | 2105344  | 15.98GB  | sda    |
| sda3        | 83 (Linux)       | 35633152 | 3.00GB   | sda    |
| sdb1.1      | 169 (NSS)        | 32       | 100.00MB | sdb    |
| sdb1.2      | 169 (NSS)        | 204832   | 100.00MB | sdb    |
| sdb1.3      | 169 (NSS)        | 409632   | 100.00MB | sdb    |
| sdb1.4      | 1CF (NSS_Raid)   | 614432   | 100.00MB | sdb    |
| sdcl        | 8E (Linux_LVM)   | 32       | 103.57MB | sdcl   |
| sdcl2       | 8E (Linux_LVM)   | 212160   | 103.59MB | sdcl   |
| sdcl3       | 83 (Linux)       | 424320   | 103.59MB | sdcl   |
| sdcl4       | 5 (DOS_Extended) | 636480   | 713.20MB | sdcl   |
| sdcl5.1     | 169 (NSS)        | 636512   | 100.00MB | sdcl   |
| sdcl7.1     | 169 (NSS)        | 1060832  | 100.00MB | sdcl   |
| sdcl8.1     | 1CF (NSS_Raid)   | 1272992  | 100.00MB | sdcl   |
| sdcl9.1     | 169 (NSS)        | 1485152  | 100.00MB | sdcl   |
| cluster.sbd | 1AD (Cluster)    | 409664   | 50.00MB  | sdd    |
| sde1        | 83 (Linux)       | 32       | 103.57MB | sde    |
| sde2        | 83 (Linux)       | 212160   | 103.59MB | sde    |
| sde3        | 8E (Linux_LVM)   | 424320   | 103.59MB | sde    |
| sde4        | 5 (DOS_Extended) | 636480   | 713.20MB | sde    |



```

sde5                83 (Linux)    636512  103.57MB sde
sde6                83 (Linux)    848672  103.57MB sde
sde7.1              1AC (Snapshot) 1060832  50.00MB sde
sdf1.1              169 (NSS)      32      100.00MB sdf
sdf2                83 (Linux)    212992  95.00MB  sdf
sdf3.1              169 (NSS)    407552  100.00MB sdf
sdh1                5 (DOS_Extended) 204800  919.75MB sdh
sdh5                7 (NTFS/HPFS) 409600  200.05MB sdh
sdh6                7 (NTFS/HPFS) 819347  200.00MB sdh
sdi1                5 (DOS_Extended) 32      499.98MB sdi
sdi5                7 (NTFS/HPFS) 64      100.96MB sdi
sdi6.1              1AC (Snapshot) 206880  75.00MB sdi
rr0p1.1             169 (NSS)      32      50.00MB rr0
rr0p1.2             169 (NSS)    102432  50.00MB rr0

```

## Sample 2: nlvm list partitions more

```
nlvm list partitions more
```

```

Name                Type          Start      Size Device Shared Maj:Min Pool
sda1                82 (Linux_Swap) 2048       1.00GB sda   No    8:1
sda2                83 (Linux)    2105344   15.98GB sda   No    8:2
sda3                83 (Linux)    35633152  3.00GB  sda   No    8:3
sdb1.1              169 (NSS)      32      100.00MB sdb   No    0:0  PMOVE
sdb1.2              169 (NSS)    204832   100.00MB sdb   No    0:0
sdb1.3              169 (NSS)    409632   100.00MB sdb   No    0:0
BIGLONGPPOOLNAME
sdb1.4              1CF (NSS_Raid) 614432   100.00MB sdb   No    0:0
sdc1                8E (Linux_LVM) 32      103.57MB sdc   No    8:33
sdc2                8E (Linux_LVM) 212160   103.59MB sdc   No    8:34
sdc3                83 (Linux)    424320   103.59MB sdc   No    8:35
sdc4                5 (DOS_Extended) 636480   713.20MB sdc   No    8:36
sdc5.1              169 (NSS)    636512   100.00MB sdc   No    0:0
sdc7.1              169 (NSS)    1060832  100.00MB sdc   No    0:0  POOL1
sdc8.1              1CF (NSS_Raid) 1272992  100.00MB sdc   No    0:0
sdc9.1              169 (NSS)    1485152  100.00MB sdc   No    0:0  PMOVE_move
cluster.sbd        1AD (Cluster) 409664    50.00MB sdd   Yes   253:20
sde1                83 (Linux)      32      103.57MB sde   No    8:65
sde2                83 (Linux)    212160   103.59MB sde   No    8:66
sde3                8E (Linux_LVM) 424320   103.59MB sde   No    8:67
sde4                5 (DOS_Extended) 636480   713.20MB sde   No    8:68
sde5                83 (Linux)    636512   103.57MB sde   No    8:69
sde6                83 (Linux)    848672   103.57MB sde   No    8:70
sde7.1              1AC (Snapshot) 1060832  50.00MB sde   No    253:6  SNAPSHOT1
sdf1.1              169 (NSS)      32      100.00MB sdf   No    0:0  T1
sdf2                83 (Linux)    212992   95.00MB sdf   No    8:82
sdf3.1              169 (NSS)    407552   100.00MB sdf   No    0:0  T2
sdh1                5 (DOS_Extended) 204800  919.75MB sdh   No    8:113
sdh5                7 (NTFS/HPFS) 409600  200.05MB sdh   No    8:117
sdh6                7 (NTFS/HPFS) 819347  200.00MB sdh   No    8:118
sdi1                5 (DOS_Extended) 32      499.98MB sdi   No    8:129
sdi5                7 (NTFS/HPFS) 64      100.96MB sdi   No    8:133
sdi6.1              1AC (Snapshot) 206880  75.00MB sdi   No    253:17  SNAP1
rr0p1.1             169 (NSS)      32      50.00MB rr0   No    0:0
rr0p1.2             169 (NSS)    102432  50.00MB rr0   No    0:0  RRPOOL

```

### Sample 3: nlvm list partitions all

```
nlvm list partitions all

Name=sda1
  Type=82(Linux_Swap) Start=2048 Size=1.00GB(2103296)
  Device=sda Shared=No M:M=8:1 Pool=None

Name=sda2
  Type=83(Linux) Start=2105344 Size=15.98GB(33527808)
  Device=sda Shared=No M:M=8:2 Pool=None

Name=sda3
  Type=83(Linux) Start=35633152 Size=3.00GB(6309888)
  Device=sda Shared=No M:M=8:3 Pool=None

Name=sdb1.1
  Type=169(NSS) Start=32 Size=100.00MB(204800)
  Device=sdb Shared=No M:M=0:0 Pool=PMOVE

Name=sdb1.2
  Type=169(NSS) Start=204832 Size=100.00MB(204800)
  Device=sdb Shared=No M:M=0:0 Pool=None

Name=sdb1.3
  Type=169(NSS) Start=409632 Size=100.00MB(204800)
  Device=sdb Shared=No M:M=0:0 Pool=BIGLONGPOOLNAME
  Label: This partition belongs to big long pool name.

Name=sdb1.4
  Type=1CF(NSS_Raid) Start=614432 Size=100.00MB(204800)
  Device=sdb Shared=No M:M=0:0 Pool=None

Name=sdcl
  Type=8E(Linux_LVM) Start=32 Size=103.57MB(212128)
  Device=sdcl Shared=No M:M=8:33 Pool=None

Name=sdcl2
  Type=8E(Linux_LVM) Start=212160 Size=103.59MB(212160)
  Device=sdcl Shared=No M:M=8:34 Pool=None

Name=sdcl3
  Type=83(Linux) Start=424320 Size=103.59MB(212160)
  Device=sdcl Shared=No M:M=8:35 Pool=None

Name=sdcl4
  Type=5(DOS_Extended) Start=636480 Size=713.20MB(1460640)
  Device=sdcl Shared=No M:M=8:36 Pool=None

Name=sdcl5.1
  Type=169(NSS) Start=636512 Size=100.00MB(204800)
  Device=sdcl Shared=No M:M=0:0 Pool=None

Name=sdcl7.1
  Type=169(NSS) Start=1060832 Size=100.00MB(204800)
  Device=sdcl Shared=No M:M=0:0 Pool=POOL1

Name=sdcl8.1
  Type=1CF(NSS_Raid) Start=1272992 Size=100.00MB(204800)
  Device=sdcl Shared=No M:M=0:0 Pool=None

Name=sdcl9.1
  Type=169(NSS) Start=1485152 Size=100.00MB(204800)
  Device=sdcl Shared=No M:M=0:0 Pool=PMOVE_move

Name=cluster.sdb
  Type=1AD(Cluster) Start=409664 Size=50.00MB(102400)
  Device=sdd Shared=Yes M:M=253:20 Pool=None
  Label: cluster

Name=sdel
  Type=83(Linux) Start=32 Size=103.57MB(212128)
  Device=sde Shared=No M:M=8:65 Pool=None

Name=sde2
```

```

Type=83(Linux) Start=212160 Size=103.59MB(212160)
Device=sde Shared=No M:M=8:66 Pool=None

Name=sde3
Type=8E(Linux_LVM) Start=424320 Size=103.59MB(212160)
Device=sde Shared=No M:M=8:67 Pool=None

Name=sde4
Type=5(DOS_Extended) Start=636480 Size=713.20MB(1460640)
Device=sde Shared=No M:M=8:68 Pool=None

Name=sde5
Type=83(Linux) Start=636512 Size=103.57MB(212128)
Device=sde Shared=No M:M=8:69 Pool=None

Name=sde6
Type=83(Linux) Start=848672 Size=103.57MB(212128)
Device=sde Shared=No M:M=8:70 Pool=None

Name=sde7.1
Type=1AC(Snapshot) Start=1060832 Size=50.00MB(102400)
Device=sde Shared=No M:M=253:6 Pool=SNAPSHOT1

Name=sdf1.1
Type=169(NSS) Start=32 Size=100.00MB(204800)
Device=sdf Shared=No M:M=0:0 Pool=T1

Name=sdf2
Type=83(Linux) Start=212992 Size=95.00MB(194560)
Device=sdf Shared=No M:M=8:82 Pool=None

Name=sdf3.1
Type=169(NSS) Start=407552 Size=100.00MB(204800)
Device=sdf Shared=No M:M=0:0 Pool=T2

Name=sdh1
Type=5(DOS_Extended) Start=204800 Size=919.75MB(1883650)
Device=sdh Shared=No M:M=8:113 Pool=None

Name=sdh5
Type=7(NTFS/HPFS) Start=409600 Size=200.05MB(409715)
Device=sdh Shared=No M:M=8:117 Pool=None

Name=sdh6
Type=7(NTFS/HPFS) Start=819347 Size=200.00MB(409600)
Device=sdh Shared=No M:M=8:118 Pool=None

Name=sdi1
Type=5(DOS_Extended) Start=32 Size=499.98MB(1023968)
Device=sdi Shared=No M:M=8:129 Pool=None

Name=sdi5
Type=7(NTFS/HPFS) Start=64 Size=100.96MB(206784)
Device=sdi Shared=No M:M=8:133 Pool=None

Name=sdi6.1
Type=1AC(Snapshot) Start=206880 Size=75.00MB(153600)
Device=sdi Shared=No M:M=253:17 Pool=SNAP1

Name=rr0p1.1
Type=169(NSS) Start=32 Size=50.00MB(102400)
Device=rr0 Shared=No M:M=0:0 Pool=None

Name=rr0p1.2
Type=169(NSS) Start=102432 Size=50.00MB(102400)
Device=rr0 Shared=No M:M=0:0 Pool=RRPOOL

```

#### Sample 4: Partitions that Contain a Specified String in the Name

```
nlvm list partitions | grep LH-DFS01-
```

| Name                  | Type      | Start | Size    | Device        |
|-----------------------|-----------|-------|---------|---------------|
| D1_LH-DFS01-1_part1.1 | 169 (NSS) | 32    | 24.99GB | D1_LH-DFS01-1 |

## 6.36 List Pool

### **list pool** <pool\_name>

Print detailed information about a specified NSS pool including its pool type (NSS64 or NSS32).

```
nlvm [nlvm_options] list pool <pool_name>
```

#### **Command Option**

##### **pool\_name**

Mandatory. Specify the name of the NSS pool.

##### **Example**

```
MYPOOL1
```

#### **Command Example**

```
nlvm list pool MYPOOL1
```

Print detailed information about the pool MYPOOL1.

#### **Response Parameters**

The command returns the following information about the specified pool:

| Label     | Description   |
|-----------|---|
| Name      | Name of the pool  |
| State     | State of the pool (Active, Deactive, Maintenance, Unknown, Not Mounted (for snapshot pool)) |
| Type      | Type of the pool (NSS64 or NSS32-bit pool)  |
| Size      | Size of the pool in MB, GB, or TB   |
| Shared    | Yes or No; whether the pool's device is marked as Shareable for Clustering                  |
| IsSnap    | Yes or No; whether the pool is a snapshot   |
| Used      | Used space in the pool in KB, MB, GB, or TB   |
| Free      | Free space in the pool in KB, MB, GB, or TB   |
| Segs      | Number of segments in the pool  |
| Volumes   | Number of volumes in the pool   |
| Snapshots | Number of snapshots, or No  |
| Move      | Name of the pool move (if applicable), or No  |

| Label         | Description  |
|---------------|--|
| Status        | Status of the pool move (if applicable)                        |
| Complete      | Percent complete for the pool move (if applicable)             |
| SnapshotNames | Names of the pool snapshots (if applicable)                    |
| Created       | If the pool is mounted, the date and time the pool was created |

The command returns the following information about the pool's segments:

| Label     | Description                          |
|-----------|--------------------------------------|
| Index     | Index number of the segment          |
| Start     | Starting offset in the pool          |
| Next      | Next offset in the pool              |
| Size      | Size of the segment in MB, GB, or TB |
| Partition | Partition name for this segment      |

The command returns the following information about each of the pool's volumes if the pool is active and it has volumes:

| Label   | Description   |
|---------|---|
| Volume  | Volume name   |
| State   | Volume state (Active, Deactive)   |
| Mounted | Yes or No; whether the volume is mounted for user access                                |
| Quota   | Volume quota in MB, GB, or TB, or None (if the volume can grow to the size of the pool) |
| Used    | Used size of the volume in KB, MB, GB, or TB  |
| Free    | Free size of the volume in KB, MB, GB, or TB  |

## Sample Command Responses

### Sample 1: Pool with 3 Volumes and 1 Snapshot

```
nlvm list pool POOL1

Name=POOL1
State=Active Type=NSS32 Size=99.00MB Shared=No IsSnap=No
Used=11.75MB Free=87.24MB Segs=1 Volumes=3 Snapshots=1 Move=No
SnapNames=SNAP1
Created=Wed May 22 16:03:26 2013
Pool segments:
  Index      Start      Next      Size      Partition
  1          0          204768    99.98MB   sdc7.1
Volumes on this pool:
  Volume     State  Mounted  Quota    Used      Free
  NSS1      Active Yes       None     1.28MB   87.27MB
  TESTVOL   Active Yes       None     564KB    87.27MB
nl VOL1     Active Yes       None     600KB    87.27MB
```

## Sample 2: Snapshot Pool for POOL1, Active with 2 of 3 Snap Volumes Mounted

```
nlvm list pool SNAP1

Name=SNAP1
State=Active Size=99.00MB Type=NSS32 Shared=No IsSnap=Yes
Used=10.96MB Free=88.03MB Segs=1 Volumes=2 Snapshots=0 Move=No
Created=Wed Jun 5 16:57:21 2013
Pool segments:
  Index      Start          Next      Size  Partition
  1          0            202752   99.00MB sdi6.1
Volumes on this pool:
  Volume    State Mounted  Quota   Used    Free
  NSS1_SV   Active  Yes      None    572MB   88.05MB
  VOL1_SV   Active  Yes      None    600KB   88.05MB
```

## Sample 3: Pool with a No Volumes and 1 Snapshot

```
nlvm list pool POOL2

Name=POOL2
State=Active Size=99.00MB Type=NSS32 Shared=No IsSnap=No
Used=10.78MB Free=88.21MB Segs=1 Volumes=0 Snapshots=1 Move=No
SnapNames=SNAPSHOT1
Created=Wed May 22 16:03:27 2013
Pool segments:
  Index      Start          Next      Size  Partition
  1          0            204768   99.98MB sdb1.3
```

## Sample 4: Snapshot Pool for POOL2, Not Mounted

```
nlvm list pool SNAPSHOT1

Name=SNAPSHOT1
State=NotMounted Size=99.00MB Type=NSS32 Shared=No IsSnap=Yes
Used=NA Free=NA Segs=1 Volumes=NA Snapshots=0 Move=No
Pool segments:
  Index      Start          Next      Size  Partition
  1          0            204768   99.98MB sde1.1
```

## Sample 5: Pool with a Pool Move at 100% Complete but before a Complete Move

```
nlvm list pool TEST

Name=TEST
State=Active Size=99.00MB Type=NSS32 Shared=No IsSnap=No
Used=10.78MB Free=88.21MB Segs=1 Volumes=0 Snapshots=0
Move=TEST_move Status=Active:Active Complete=100%
Created=Tue Jun 11 17:18:08 2013
Pool segments:
  Index      Start          Next      Size  Partition
  1          0            204768   99.98MB sdb1.1
Volumes on this pool:
  Volume    State Mounted  Quota   Used    Free
  VOL2     Active  Yes      None    572MB   88.05MB
  VOL3     Active  Yes      None    600KB   88.05MB
```

## Sample 6: Pool Is Deactive

```
nlvm list pool TEST2

NAME=TEST2
State=Deactive Size=1019.00MB Type=NSS32 Shared=No IsSnap=No
Used=NA Free=NA Segs=2 Volumes=NA Snapshots=0 Move=No
Created: Mon Sep 23 16:33:20 2013
Pool segments:
  Index      Start          Next      Size  Partition
  1          0            1044416  509.96MB sdf1.1
  2        1044416        2088832  509.96MB sde1.1
```

## 6.37 List Pools

### **list pools [exclude] [more|all]**

Print a list of all NSS pools.

```
nlvm [-t] list pools [exclude] [more|all]
```

#### **Command Options**

##### **exclude=<nss|shared|nonshared|snap|snapnomount>**

Specify types of pools to exclude from the list. The `exclude` option can be used multiple times to add exclusions.

##### **Exclude Options**

###### **nss**

NSS pools

###### **shared**

Shared pools

###### **nonshared**

Pools that are not shared

###### **snap**

Snapshot pools that are mounted

###### **snapnomount**

Snapshot pools that are not mounted

##### **Example**

```
exclude=snap exclude=snapnomount
```

##### **-t, --terse**

Use this NLVM option to format the output for parsing.

##### **more**

Prints more information than appears in the standard output. It can be used with or without the `-t` NLVM option.

##### **Example**

```
more
```

##### **all**

Prints detailed information about each of the NSS pools. This is the same information that is printed for the `nlvm list pool <pool_name>` command. It can be used with or without the `-t` NLVM option.

##### **Example**

```
all
```

#### **Command Example**

```
nlvm list pools more exclude=shared
```

Print detailed information about each of the pools, but exclude shared pools.

## Response Parameters

You can issue the commands with the `--terse` NLVM option to output the same information in a format that is more easily parsed.

### Standard Output

The command returns the following information about the pools on the server:

| Label  | Description   |
|--------|---|
| Name   | Name of the pool  |
| State  | State of the pool (Active, Deactive, Maintenance, Unknown, Not Mounted (for snapshot pool)) |
| Type   | Denotes the type of pool (NSS64 or NSS32 pool)  |
| Size   | Size of the pool in MB, GB, or TB   |
| Shared | Yes or No; whether the pool's device is marked as Shareable for Clustering                  |
| IsSnap | Yes or No; whether the pool is a snapshot   |

### More Output

The command returns the following additional information about the pools on the server:

| Label   | Description  |
|---------|--|
| Used    | Used space in the pool in KB, MB, GB, or TB          |
| Free    | Free space in the pool in KB, MB, GB, or TB          |
| Segs    | Number of segments in the pool                       |
| Volumes | Number of volumes in the pool                        |
| Move    | If there is a pool move, its percent complete; or No |

### All Output

If the `all` option is used, the command returns the same information about each pool as is displayed for the `nlvm list pool <pool_name>` command.

## Sample Command Response

### Sample 1: nlvm list pools

```
nlvm list pools

Name          State   Type   Size Shared IsSnap
POOL1        Active NSS32  99.00MB No     No
PMOVE        Active NSS32  99.00MB No     No
BIGLONGPOOLNAME Active NSS32  99.00MB No     No
RRPOOL       Active NSS64  49.00MB No     No
SNAP1        Active NSS32  99.00MB No     Yes
SNAPSHOT1    NotMounted NSS32  99.00MB No     Yes
T1           Active NSS32  99.00MB No     No
T2           Active NSS32  99.00MB No     No
```



## Sample 2: nlvms list pools more

```
nlvm list pools more
```

| Name                    | State      | Size    | Shared | IsSnap | Used    | Free    | Segs | Vols |
|-------------------------|------------|---------|--------|--------|---------|---------|------|------|
| Move<br>POOL1           | Active     | 99.00MB | No     | No     | 11.75MB | 87.24MB | 1    | 3    |
| No<br>PMOVE             | Active     | 99.00MB | No     | No     | 10.78MB | 88.21MB | 1    | 0    |
| 100%<br>BIGLONGPOOLNAME | Active     | 99.00MB | No     | No     | 10.78MB | 88.21MB | 1    | 0    |
| No<br>RRPOOL            | Active     | 49.00MB | No     | No     | 10.78MB | 38.21MB | 1    | 0    |
| No<br>SNAP1             | Active     | 99.00MB | No     | Yes    | 10.96MB | 88.03MB | 1    | 2    |
| No<br>SNAPSHOT1         | NotMounted | 99.00MB | No     | Yes    | NA      | NA      | 1    | NA   |
| No<br>T1                | Active     | 99.00MB | No     | No     | 10.78MB | 88.21MB | 1    | 0 No |
| T2                      | Active     | 99.00MB | No     | No     | 10.78MB | 88.21MB | 1    | 0 No |

## Sample 3: nlvms list pools all

```
nlvm list pools all
```

```
Name=POOL1
```

```
State=Active Size=99.00MB Shared=No IsSnap=No  
Used=11.75MB Free=87.24MB Segs=1 Volumes=3 Snapshots=1 Move=No  
SnapNames=SNAP1  
Created: Wed May 22 16:03:26 2013
```

```
Pool segments:
```

| Index | Start | Next   | Size    | Partition |
|-------|-------|--------|---------|-----------|
| 1     | 0     | 204768 | 99.98MB | sdc7.1    |

```
Volumes on this pool:
```

| Volume  | State  | Mounted | Quota | Used   | Free    |
|---------|--------|---------|-------|--------|---------|
| NSS1    | Active | Yes     | None  | 1.28MB | 87.27MB |
| TESTVOL | Active | Yes     | None  | 564KB  | 87.27MB |
| VOL1    | Active | Yes     | None  | 600KB  | 87.27MB |

```
Name=PMOVE
```

```
State=Active Size=99.00MB Shared=No IsSnap=No  
Used=10.78MB Free=88.21MB Segs=1 Volumes=0 Snapshots=0  
Move=PMOVE_move Status=Active:Active Complete=100%  
Created: Tue Jun 11 17:18:08 2013
```

```
Pool segments:
```

| Index | Start | Next   | Size    | Partition |
|-------|-------|--------|---------|-----------|
| 1     | 0     | 204768 | 99.98MB | sdb1.1    |

```
Name=BIGLONGPOOLNAME
```

```
State=Active Size=99.00MB Shared=No IsSnap=No  
Used=10.78MB Free=88.21MB Segs=1 Volumes=0 Snapshots=1 Move=No  
SnapNames=SNAPSHOT1  
Created: Wed May 22 16:03:27 2013
```

```
Pool segments:
```

| Index | Start | Next   | Size    | Partition |
|-------|-------|--------|---------|-----------|
| 1     | 0     | 204768 | 99.98MB | sdb1.3    |

```
Name=RRPOOL
```

```
State=Active Size=49.00MB Shared=No IsSnap=No  
Used=10.78MB Free=38.21MB Segs=1 Volumes=0 Snapshots=0 Move=No  
Created: Wed Jun 12 17:30:53 2013
```

```
Pool segments:
```

| Index | Start | Next   | Size    | Partition |
|-------|-------|--------|---------|-----------|
| 1     | 0     | 102368 | 49.98MB | rr0pl.2   |

```
Name=SNAP1
```

```
State=Active Size=99.00MB Shared=No IsSnap=Yes  
Used=10.96MB Free=88.03MB Segs=1 Volumes=2 Snapshots=0 Move=No  
Created: Wed Jun 5 16:57:21 2013
```

```
Pool segments:
```

| Index | Start | Next   | Size    | Partition |
|-------|-------|--------|---------|-----------|
| 1     | 0     | 202752 | 99.00MB | sdi6.1    |

```
Volumes on this pool:
```

| Volume  | State  | Mounted | Quota | Used  | Free    |
|---------|--------|---------|-------|-------|---------|
| NSS1_SV | Active | Yes     | None  | 572KB | 88.05MB |

```

VOL1_SV    Active    Yes        None       600KB    88.05MB

Name=SNAPSHOT1
State=NotMounted Size=99.00MB Shared=No IsSnap=Yes
Used=NA Free=NA Segs=1 Volumes=NA Snapshots=0 Move=No
Pool segments:
  Index      Start      Next      Size Partition
    1         0        204768   99.98MB sde7.1

Name=T1
State=Active Size=99.00MB Shared=No IsSnap=No
Used=10.78MB Free=88.21MB Segs=1 Volumes=0 Snapshots=0 Move=No
Created: Tue Jun 25 17:33:25 2013
Pool segments:
  Index      Start      Next      Size Partition
    1         0        204768   99.98MB sdf1.1

Name=T2
State=Active Size=99.00MB Shared=No IsSnap=No
Used=10.78MB Free=88.21MB Segs=1 Volumes=0 Snapshots=0 Move=No
Created: Fri Jun 28 10:25:43 2013
Pool segments:
  Index      Start      Next      Size Partition
    1         0        204768   99.98MB sdf3.1

```

#### Sample 4: nlvm list pools all --terse

```

nlvm list pools all --terse

Name=POOL1 State=Active Size=99.00MB Shared=No IsSnap=No Used=11.75MB
Free=87.24MB Segs=1 Volumes=3 Snapshots=1 Move=No SnapNames=SNAP1
Created=Wed May 22 16:03:26 2013
Index=1 Start=0 Next=204768 Size=99.98MB Part=sdc7.1
Volume=NSS1 State=Active Mounted=Yes Quota=None Used=1.28MB Free=87.27MB
Volume=TESTVOL State=Active Mounted=Yes Quota=None Used=564KB Free=87.27MB
Volume=VOL1 State=Active Mounted=Yes Quota=None Used=600KB Free=87.27MB

Name=PMOVE State=Active Size=99.00MB Shared=No IsSnap=No Used=10.78MB
Free=88.21MB Segs=1 Volumes=0 Snapshots=0 Move=PMOVE_move
Status=Active:Active Complete=100%
Created=Tue Jun 11 17:18:08 2013
Index=1 Start=0 Next=204768 Size=99.98MB Part=sdb1.1

Name=BIGLONGPOOLNAME State=Active Size=99.00MB Shared=No IsSnap=No
Used=10.78MB Free=88.21MB Segs=1 Volumes=0 Snapshots=1 Move=No
SnapNames=SNAPSHOT1
Created=Wed May 22 16:03:27 2013
Index=1 Start=0 Next=204768 Size=99.98MB Part=sdb1.3

Name=RRPOOL State=Active Size=49.00MB Shared=No IsSnap=No Used=10.78MB
Free=38.21MB Segs=1 Volumes=0 Snapshots=0 Move=No
Created=Wed Jun 12 17:30:53 2013
Index=1 Start=0 Next=102368 Size=49.98MB Part=rr0p1.2

Name=SNAP1 State=Active Size=99.00MB Shared=No IsSnap=Yes Used=10.96MB
Free=88.03MB Segs=1 Volumes=2 Snapshots=0 Move=No
Created=Wed Jun 5 16:57:21 2013
Index=1 Start=0 Next=202752 Size=99.00MB Part=sdi6.1
Volume=NSS1_SV State=Active Mounted=Yes Quota=None Used=572KB Free=88.05MB
Volume=VOL1_SV State=Active Mounted=Yes Quota=None Used=600KB Free=88.05MB

Name=SNAPSHOT1 State=NotMounted Size=99.00MB Shared=No IsSnap=Yes Used=NA
Free=NA Segs=1 Volumes=NA Snapshots=0 Move=No
Index=1 Start=0 Next=204768 Size=99.98MB Part=sde7.1

Name=T1 State=Active Size=99.00MB Shared=No IsSnap=No Used=10.78MB
Free=88.21MB Segs=1 Volumes=0 Snapshots=0 Move=No
Created=Tue Jun 25 17:33:25 2013
Index=1 Start=0 Next=204768 Size=99.98MB Part=sdf1.1

Name=T2 State=Active Size=99.00MB Shared=No IsSnap=No Used=10.78MB
Free=88.21MB Segs=1 Volumes=0 Snapshots=0 Move=No
Created=Fri Jun 28 10:25:43 2013
Index=1 Start=0 Next=204768 Size=99.98MB Part=sdf3.1

```

## 6.38 List Snap

**list snap** <snap\_name>

Print detailed information about a specified snapshot.

```
nlvm [nlvm_options] list snap <snap_name>
```

### Command Example

```
nlvm list SNAP1
```

Print detailed information about SNAP1.

### Response Parameters

The command returns the following information about the specified pool snapshot:

| Label     | Description  |
|-----------|--|
| Name      | Name of the snapshot   |
| Pool      | Name of the pool being snapped   |
| Mounted   | Yes or No; whether the snapshot is mounted as a pool   |
| Size      | Size of the pool in MB, GB, or TB  |
| Shared    | Yes or No; whether the snapshot's device is marked as Shareable for Clustering. Typically, No, because snapshots are not supported for clustered pools at this time. |
| PoolSize  | Source pool in KB, MB, GB, or TB   |
| Chunk     | Snapshot chunk size in KB (ex: 128)  |
| Full      | Percent of space on the partition that is used for copy-on-write-data  |
| PartSize  | Partition size in MB, GB, or TB, and the number of whole sectors in that space   |
| Partition | Name of the partition for the snapshot   |
| M:M       | Major:minor of the snapshot object   |
| Writeable | Yes or No; whether the snapshot is writeable   |

### Sample Command Responses

#### Sample 1: Snap Is Not Mounted

```
nlvm list snap SNAPSHOT1
```

```
Name=SNAPSHOT1  
Pool=BIGLONGPOOLNAME Mounted=No Shared=No  
PoolSize=99.98MB Chunk=128 Full=1%  
PartSize=50.00MB(102400) Partition=sde7.1  
M:M=253:8 Writeable=Yes
```

## Sample 2: Snap Is Mounted

```
nlvm list snap SNAP1

Name=SNAP1
Pool=POOL1 Mounted=Yes Shared=No
PoolSize=99.00MB Chunk=128 Full=12%
PartSize=75.00MB(153600) Partition=sdi6.1
M:M=253:19 Writeable=Yes
```

## 6.39 List Snaps

### list snaps [more|all]

Print a list of all NSS pool snapshots. For each, display its pool name and its mount state.

```
nlvm [-t] list snaps [more|all]
```

#### Command Options

**-t, --terse**

Use this NLVM option to format the output for parsing.

#### more

Prints more information than appears in the standard output. It can be used with or without the `-t` NLVM option.

#### Example

```
more
```

#### all

Prints detailed information about each of the snapshots. This is the same information that is printed for the `nlvm list snapshot <snap_name>` command. It can be used with or without the `-t` NLVM option.

#### Example

```
all
```

#### Command Example

```
nlvm list snaps more
```

Print a list of all snapshots and detailed information about each one.

#### Response Parameters

You can issue the commands with the `--terse` NLVM option to output the same information in a format that is more easily parsed.

#### Standard Output

The command returns the following information about the pool snapshots on the server:

| Label   | Description  |
|---------|--|
| Name    | Name of the snapshot   |
| Pool    | Name of the pool being snapped   |
| Mounted | Yes or No; whether the snapshot is mounted as a pool   |
| Size    | Size of the pool in MB, GB, or TB  |
| Shared  | Yes or No; whether the snapshot's device is marked as Shareable for Clustering. Typically, No, because snapshots are not supported for clustered pools at this time. |

### More Output

The command returns the following additional information about the pool snapshots on the server:

| Label     | Description  |
|-----------|--|
| PoolSize  | Source pool in KB, MB, GB, or TB   |
| Chunk     | Snapshot chunk size in KB (ex: 128)  |
| Full      | Percent of space on the partition that is used for copy-on-write-data          |
| PartSize  | Partition size in MB, GB, or TB, and the number of whole sectors in that space |
| Partition | Name of the partition for the snapshot   |

### All Output

If the `all` option is specified, the information returned for each pool snapshot is the same as for the `nlvm list snap <snap_name>` command.

### Sample Command Response

#### Sample 1: nlvm list snaps

```
nlvm list snaps

      Name      Pool           Mounted Shared
  SNAP1      POOL1             Yes     No
  SNAPSHOT1  BIGLONGPOOLNAME    No      No
```

#### Sample 2: nlvm list snaps more

```
nlvm list snaps more

      Name      Pool           Mounted Shared PoolSize Chunk Full  PartSize
Partition
  SNAP1      POOL1             Yes     No   99.00MB 128 12% 75.00MB sdi6.1
  SNAPSHOT1  BIGLONGPOOLNAME No      No   99.98MB 128 1%  50.00MB sde7.1
```

### Sample 3: nlvm list snaps all

```
nlvm list snaps all

Name=SNAP1
Pool=POOL1 Mounted=Yes Shared=No
PoolSize=99.00MB Chunk=128 Full=12%
PartSize=75.00MB(153600) Partition=sdi6.1
M:M=253:19 Writeable=Yes

Name=SNAPSHOT1
Pool=BIGLONGPOOLNAME Mounted=No Shared=No
PoolSize=99.98MB Chunk=128 Full=1%
PartSize=50.00MB(102400) Partition=sde7.1
M:M=253:8 Writeable=Yes
```

### Sample 4: nlvm list snaps all --terse

```
nlvm list snaps all --terse

Name=SNAP1 Pool=POOL1 Mounted=Yes Shared=No PoolSize=99.00MB Chunk=128
Full=12% PartSize=75.00MB(153600) Partition=sdi6.1 M:M=253:19 Writeable=Yes

Name=SNAPSHOT1 Pool=BIGLONGPOOLNAME Mounted=No Shared=No PoolSize=99.98MB
Chunk=128 Full=1% PartSize=50.00MB(102400) Partition=sde7.1 M:M=253:8
```

## 6.40 List Volume

### list volume <volume\_name>

Print detailed information about a specified NSS volume.

```
nlvm [nlvm_options] list volume <volume_name>
```

#### Command Option

##### **volume\_name**

Mandatory. Specify the name of the NSS volume.

##### **Example**

```
MYVOL1
```

#### Command Example

```
nlvm list volume MYVOL1
```

Print detailed information about the volume MYVOL1.

#### Response Parameters

The command returns the following information about the specified NSS volume:

| Label   | Description   |
|---------|---|
| Name    | Name of the volume  |
| Pool    | Name of the pool  |
| State   | Volume state (Active, Deactive)   |
| Mounted | Yes or No; whether the volume is mounted for user access                  |
| Shared  | Yes or No; whether the volume's device is marked Shareable for Clustering |

| Label               | Description  |
|---------------------|--|
| Mountpoint          | Full Linux path where the volume is mounted; typically, /media/nss/<volume_name>                 |
| Used                | Amount of used space in KB, MB, GB, or TB  |
| Avail               | Amount of available space (free space plus purgeable space) in KB, MB, GB, or TB                 |
| Quota               | None, or amount of the volume quota in MB, GB, or TB   |
| Purgeable           | Amount of purgeable space in KB, MB, GB, or TB   |
| Attributes          | Volume attributes (such as Salvage, Compression, User Space Quotas, Directory Quotas, Encrypted) |
| ReadAheadBlocks     | Setting for the Read Ahead Blocks parameter  |
| PrimaryNameSpace    | Primary lookup name space; the default is LONG   |
| Objects             | Number of objects  |
| Files               | Number of files  |
| BlockSize           | Block size; typically, 4096 bytes  |
| ShredCount          | Number of shredding cycles (1 to 7), where 0 is no shredding                                     |
| AuthModelID         | Authentication model ID (1, 0)   |
| SupportedNameSpaces | Supported name spaces (DOS, MAC, UNIX, LONG)   |
| CreateTime          | Date created (Day Month dd hh:mm:ss yyyy)  |
| ArchiveTime         | Date last archived (Never , Day Month dd hh:mm:ss yyyy), or Never                                |

## Sample Command Response

### Sample 1: NSS Volume, Mounted, Unshared

```
nlvm list volume NSS1

Name=NSS1
Pool=POOL1 State=Active Mounted=Yes Shared=No
Mountpoint=/media/nss/NSS1
Used=1.28MB Avail=87.27MB Quota=None Purgeable=12KB
Attributes=Salvage,Compression,Encrypted[AES128]
ReadAheadBlocks=16 PrimaryNameSpace=LONG
Objects=28 Files=23
BlockSize=4096 ShredCount=1 AuthModelID=1
SupportedNameSpaces=DOS,MAC,UNIX,LONG
CreateTime: Wed May 22 16:03:26 2013
ArchiveTime: Never
```

### Sample 2: NSS Volume, Not Mounted, Unshared

```
nlvm list volume NSS1

Name=NSS1
Pool=POOL1 State=Dective Mounted=No Shared=No
Mountpoint=/media/nss/NSS1
CreateTime: Wed May 22 16:03:26 2013
ArchiveTime: Never
```

### Sample 3: NSS Snapshot Volume, Mounted, Unshared

```
nlvm list volume NSS1_SV

Name=NSS1_SV
Pool=SNAP1 State=Active Mounted=Yes Shared=No
Mountpoint=/media/nss/NSS1_SV
Used=572KB Avail=88.05MB Quota=None Purgeable=12KB
Attributes=Salvage,Compression,Encrypted[AES128]
ReadAheadBlocks=16 PrimaryNameSpace=LONG
Objects=15 Files=15
BlockSize=4096 ShredCount=1 AuthModelID=1
SupportedNameSpaces=DOS,MAC,UNIX,LONG
CreateTime: Wed Jun 5 16:57:21 2013
ArchiveTime: Never
```

## 6.41 List Volumes

### list volumes [more|all]

Print a list of NSS volumes on the system. For each, display its pool name and volume state (active or deactive).

```
nlvm [-t] list volumes [more|all]
```

#### Command Options

##### -t, --terse

Use this NLVM option to format the output for parsing.

##### more

Prints more information than appears in the standard output. It can be used with or without the -t NLVM option.

##### Example

```
more
```

##### all

Prints detailed information about each of the NSS volumes. This is the same information that is printed for the `nlvm list volume <nss_volume_name>` command. It can be used with or without the -t NLVM option.

##### Example

```
all
```

#### Command Example

```
nlvm list volumes [more]
```

Print a list of NSS volumes, and display detailed information about each volume.

#### Response Parameters

You can issue the commands with the `--terse` NLVM option to output the same information in a format that is more easily parsed.

#### Standard Output

The command returns the following information about the NSS volumes on the server:



| Label   | Description   |
|---------|---|
| Name    | Name of the volume  |
| Pool    | Name of the pool  |
| State   | Volume state (Active, Deactive)   |
| Mounted | Yes or No; whether the volume is mounted for user access                  |
| Shared  | Yes or No; whether the volume's device is marked Shareable for Clustering |

### More Output

The command returns the following additional information about the NSS volumes on the server:

| Label      | Description  |
|------------|--|
| Used       | Amount of used space in KB, MB, GB, or TB  |
| Avail      | Amount of available space (free space plus purgeable space) in KB, MB, GB, or TB                 |
| Quota      | None, or amount of the volume quota in MB, GB, or TB   |
| Attributes | Volume attributes (such as Salvage, Compression, User Space Quotas, Directory Quotas, Encrypted) |

### All Output

If the `all` option is specified, the information returned for each volume is the same as for the `list volume <volume_name>` command.

### Sample Command Response

#### Sample 1: `nlvm list volumes`

```
nlvm list volumes

Name      Pool      State Mounted Shared
NSS1     POOL1     Active Yes     No
NSS1_SV  SNAP1     Active Yes     No
TESTVOL  POOL1     Active Yes     No
VOL1     POOL1     Active Yes     No
VOL1_SV  SNAP1     Active Yes     No
```

#### Sample 2: `nlvm list volumes more`

```
nlvm list volumes more

Name      Pool      State Mounted Shared      Used      Avail      Quota  Attributes
NSS1     POOL1     Active Yes     No      1.28MB    87.27MB    None
Salvage,Compression,Encrypted[AES128]
NSS1_SV  SNAP1     Active Yes     No      572KB     88.05MB    None
Salvage,Compression,Encrypted[AES128]
TESTVOL  POOL1     Active Yes     No      564KB     87.27MB    None  Salvage
VOL1     POOL1     Active Yes     No      600KB     87.27MB    None  Salvage
VOL1_SV  SNAP1     Active Yes     No      600KB     88.05MB    None  Salvage
```

### Sample 3: nlvm list volumes more --terse

```
nlvm list volumes more --terse

Name=NSS1 Pool=POOL1 State=Active Mounted=Yes Shared=No Used=1.28MB
Avail=87.27MB Quota=None Attributes=Salvage,Compression,Encrypted[AES128]
Name=NSS1_SV Pool=SNAP1 State=Active Mounted=Yes Shared=No Used=572KB
Avail=88.05MB Quota=None Attributes=Salvage,Compression,Encrypted[AES128]
Name=TESTVOL Pool=POOL1 State=Active Mounted=Yes Shared=No Used=564KB
Avail=87.27MB Quota=None Attributes=Salvage
Name=VOL1 Pool=POOL1 State=Active Mounted=Yes Shared=No Used=600KB
Avail=87.27MB Quota=None Attributes=Salvage
Name=VOL1_SV Pool=SNAP1 State=Active Mounted=Yes Shared=No Used=600KB
Avail=88.05MB Quota=None Attributes=Salvage
```

### Sample 4: nlvm list volumes all

```
nlvm list volumes all

Name=NSS1
Pool=POOL1 State=Active Mounted=Yes Shared=No
Used=1.28MB Avail=87.27MB Quota=None Purgeable=12KB
Attributes=Salvage,Compression,Encrypted[AES128]
ReadAheadBlocks=16 PrimaryNameSpace=LONG
Mountpoint=/media/nss/NSS1
Objects=28 Files=23
BlockSize=4096 ShredCount=1 AuthModelID=1
SupportedNameSpaces=DOS,MAC,UNIX,LONG
CreateTime: Wed May 22 16:03:26 2013
ArchiveTime: Never

Name=NSS1_SV
Pool=SNAP1 State=Active Mounted=Yes Shared=No
Used=572KB Avail=88.05MB Quota=None Purgeable=12KB
Attributes=Salvage,Compression,Encrypted[AES128]
ReadAheadBlocks=16 PrimaryNameSpace=LONG
Mountpoint=/media/nss/NSS1_SV
Objects=15 Files=15
BlockSize=4096 ShredCount=1 AuthModelID=1
SupportedNameSpaces=DOS,MAC,UNIX,LONG
CreateTime: Wed Jun 5 16:57:21 2013
ArchiveTime: Never

Name=TESTVOL
Pool=POOL1 State=Active Mounted=Yes Shared=No
Used=564KB Avail=87.27MB Quota=None Purgeable=8KB
Attributes=Salvage
ReadAheadBlocks=16 PrimaryNameSpace=LONG
Mountpoint=/media/nss/TESTVOL
Objects=14 Files=14
BlockSize=4096 ShredCount=1 AuthModelID=1
SupportedNameSpaces=DOS,MAC,UNIX,LONG
CreateTime: Mon Jun 17 15:21:02 2013
ArchiveTime: Never

Name=VOL1
Pool=POOL1 State=Active Mounted=Yes Shared=No
Used=600KB Avail=87.27MB Quota=None Purgeable=12KB
Attributes=Salvage
ReadAheadBlocks=16 PrimaryNameSpace=LONG
Mountpoint=/media/nss/VOL1
Objects=15 Files=15
BlockSize=4096 ShredCount=1 AuthModelID=1
SupportedNameSpaces=DOS,MAC,UNIX,LONG
```

```

CreateTime: Wed May 22 16:03:26 2013
ArchiveTime: Never

Name=VOL1_SV
Pool=SNAP1 State=Active Mounted=Yes Shared=No
Used=600KB Avail=88.05MB Quota=None Purgeable=12KB
Attributes=Salvage
ReadAheadBlocks=16 PrimaryNameSpace=LONG
Mountpoint=/media/nss/VOL1_SV
Objects=15 Files=15
BlockSize=4096 ShredCount=1 AuthModelID=1
SupportedNameSpaces=DOS,MAC,UNIX,LONG
CreateTime: Wed Jun 5 16:57:21 2013
ArchiveTime: Never

```

## 6.42 Mount

**mount** *<pool\_name>*

Mount a specified NSS pool.

```
nlvm [nlvm_options] mount <pool_name>
```

### Command Option

#### *pool\_name*

Mandatory. Specify the name of the NSS pool to mount.

The `nlvm mount` command internally sets the `-m` flag, so only the specified pool is mounted.

#### Example

```
MYPOOL1
```

### Command Example

```
nlvm mount MYPOOL1
```

Mount the pool MYPOOL1.

## 6.43 Move

**move** *<pool\_name>* *<device>* *<size>* [*<device>* *<size>*...]

Move an NSS pool from one location to another on the same system. If the new location is larger than the original location, the pool is automatically expanded after the move is complete.

```
nlvm [nlvm_options] move <pool_name> <device> <size> [<device>
<size>...]
```

You can use the `device` and `size` combination multiple times to create a move target comprised of multiple segments. You must specify a size for each device. The `device` and `size` options can be used in any order. The first device instance is matched to the first size instance, and so on. The move target's size is the sum of the space contributed from the specified segments.

The total size of the target must be at least as big as the pool. You cannot shrink a pool by using the move command. If the size is larger, the pool size is expanded when the move is complete.

If a pool is cluster-enabled, issue the command on the node where the pool cluster resource is currently online. The move advances only when the resource is online. If the pool cluster resource is cluster migrated to another node, the move is enabled and active on the new node when the resource is brought online, and then the pool move continues. The status of the pool move can be reported only on the node where the resource is online.

The `move` command uses NSS software RAID mirroring underneath to copy the data to the target location. If server performance is too slow during a move, you can temporarily pause the mirroring with the `nlvm pause move` command. While the move is paused, the pool move status is reported as `Not Enabled`. Resume the mirroring with the `nlvm resume move` command. The pool move continues from where it was paused. The move will automatically resume in a cluster setup under certain conditions. See the `nlvm resume move` command for details.

You can check the status of a pool move by using the `nlvm list move` command.

When the move status is 100% complete, it is not yet final. You can issue the `nlvm complete move <move_name>` command to finalize the move. This sets the pool to the new location and removes the original location. Other NSS utilities might also complete the move. For information, see “Moving a Pool” in the *OES 2023: NSS File System Administration Guide for Linux*.

You can delete a pool move by using the `nlvm delete move` command with the `abort` option. This sets the pool back to the original location and removes the new location.

In a cluster, issue the commands to complete, delete, or list the pool move from the node where the pool cluster resource is currently online.

### Command Options

#### ***pool\_name***

Mandatory. Specify the name of the NSS pool to be moved. This must be the first command option.

#### **Example**

```
MYPOOL1
```

#### ***device=device\_name***

Mandatory. Specify the target device where the pool will be relocated.

You can specify multiple device instances to create a move target comprised of multiple segments. Each device instance must have a matching `size` instance. The first device instance is matched to the first size instance, and so on.

#### **Example**

```
device=sdg
```

#### ***size=<value[K|M|G|T]|max>***

Mandatory. Specify the size of the target partition. The size must be the same size or larger than the source pool.

If multiple devices are specified, each device instance must have a matching `size` instance. The first size instance is matched to the first device instance, and so on.

#### **Example**

```
size=200G  
size=3.98T
```

### Command Examples

```
nlvm move MYPOOL1 device=sdg size=200G
```

Move the NSS pool named MYPOOL1 to the `/dev/sdg` device and allocate 200 GB to the partition.

```
nlvm move MYPOOL1 device=sdg size=200G device=sdh size=500G
```

Move the NSS pool named MYPOOL1 to a 700 GB space comprised of 200 GB of free space from device `sdg` and 500 GB of free space from device `sdh`.

## 6.44 Pause Move

```
pause move <move_name | pool_name>
```

Temporarily pause the mirroring for a specified pool move. While the move is paused, the pool move status is reported as `Not Enabled`.

```
nlvm [nlvm_options] pause move <move_name | pool_name>
```

The `move` command uses NSS software RAID mirroring underneath to copy the data to the target location. If server performance is too slow during a move, you can use this command to temporarily pause the mirroring. For example, you can pause the move during the server's peak usage times, and resume the move during the server's off-peak usage times.

Use the `nlvm resume move` command to resume mirroring for the pool move.

### Command Option

#### **move\_name**

Mandatory. Specify the name of the pool move that you want to pause. The move name typically looks like `POOLNAME_move`.

#### **pool\_name**

Mandatory. Specify the name of the NSS pool that is being moved.

#### **Example**

```
MYPOOL1
```

### Command Example

```
nlvm pause move MYPOOL1
```

Temporarily pause the mirroring for the pool move for the pool `MYPOOL1`.

## 6.45 Pool Activate

```
pool activate <pool_name>
```

Activate a specified NSS pool.

```
nlvm [nlvm_options] pool activate <pool_name>
```

### Command Option

#### **pool\_name**

Mandatory. Specify the name of the NSS pool to activate.

### Example

```
MYPOOL1
```

### Command Example

```
nlvm pool activate MYPOOL1
```

Activate the pool MYPOOL1.

## 6.46 Pool Deactivate

### **pool deactivate** <pool\_name>

Deactivate a specified NSS pool.

```
nlvm [nlvm_options] pool deactivate <pool_name>
```

### Command Option

#### **pool\_name**

Mandatory. Specify the name of the NSS pool to deactivate.

### Example

```
MYPOOL1
```

### Command Example

```
nlvm pool deactivate MYPOOL1
```

Deactivate the pool MYPOOL1.

## 6.47 RAID

### **raid** <raid\_action> <raid\_name>

Perform actions on an NSS software RAID device.

```
nlvm [nlvm_options] raid <raid_action> <raid_name>
```

### RAID Actions

#### **abort** <raid\_name>

Abort the restripe or remirror currently in progress on the specified NSS software RAID. If the restripe/remirror is complete, the command has no effect.

### Example

```
nlvm raid abort MYRAID1
```

#### **delete** <raid\_name>

Delete a single element mirror from a pool, and leave the pool on the corresponding partition. This applies for RAID 1 (mirror) objects only.

This is a duplicate of the `nlvm delete raid` command, but it is added here for support reasons. This command removes only a single element mirror object.

### Example

```
nlvm raid delete MYRAID1
```

### disable <raid\_name>

Disable an NSS software RAID device from remirroring or restriping on this server, and do not allow stamp updates to occur.

This command is used in OES Cluster Services clusters to disable an NSS software RAID device that is active on another node.

### Example

```
nlvm raid disable MYRAID1
```

### enable <raid\_name>

Enable a RAID device to remirror or restripe on this server. This enables an NSS software RAID device that was disabled by using the `nlvm raid disable` command.

This command is used in OES Cluster Services clusters to enable an NSS software RAID device for this node. It is important that the RAID device be enabled on only one node at a time.

---

**WARNING:** Use caution when in a cluster configuration to avoid possible corruption that can occur if the RAID is enabled on multiple nodes at the same time.

---

### Example

```
nlvm raid enable MYRAID1
```

### force <raid\_name>

Force a single element mirror to be in sync. This condition can occur if a mirror element was removed, and the last element shows that it is not in sync due to a crash after a successful remirror. This command is only valid on NSS software RAID 1 (mirror) devices.

If you have a single element RAID 1 where the element shows out of sync, you can alternatively put it into sync (if you feel that it has all of the data) by selecting the Restripe (F6) function on the Software RAID page in NSSMU.

---

**WARNING:** If a remirror has not completed successfully on this element, using the `nlvm raid force` command causes the element to look in sync, but the data is not there, and is corrupt. Use this command only if you know that a remirror was completed successfully on this element.

---

### Example

```
nlvm raid force MYRAID1
```

### pause <raid\_name>

Pause a remirror process to allow other I/O to happen during a heavy I/O process. This command is valid only on NSS software RAID 1 (mirror) devices. Because remirroring can cause many I/Os to the devices, a pause allows other I/Os to happen more quickly.

The device must be resumed again by using the `nlvm raid resume` command. The pause is intended to be used only for a short time.

### Example

```
nlvm raid pause MYRAID1
```

### remirror <raid\_name>

Restart a remirror or restripe process on the specified NSS software RAID device that has either been aborted or has failed.

### Examples

```
nlvm raid remirror MYRAID1
nlvm raid remirror MYRAID5
```

### resume <raid\_name>

Resume a remirror process that was paused by using the `nlvm raid pause` command. This command is valid only on NSS software RAID 1 (mirror) devices.

### Example

```
nlvm raid resume MYRAID1
```

### status [raid\_name]

Check the status on one or all NSS software RAID devices. The name is optional. If a name is specified, it returns detailed status for the given RAID device. If the name is omitted, it returns the status for all the NSS software RAID devices on the server.

### Examples

```
nlvm raid status MYRAID1
nlvm raid status
```

### Command Option

#### *raid\_name*

Mandatory where specified. Specify the name of the NSS software RAID device to be acted upon.

### Example

```
MYRAID1
```

### Sample Command Responses

#### Sample 1: RAID Status During a Remirror

```
nlvm raid status LH_DFS01_01_R1

LH_DFS01_01_R1 is remirroring at 9%
--> D1_LH-DFS01-1_part1.1 (100%) In Sync
--> D2_LH-DFS01-1_part1.1 (9%) Out of Sync
```

#### Sample 2: RAID Status During a Remirror on a Cluster Node where the RAID Is Not Active

```
nlvm raid status

LH_DFS01_01_R1 is remirroring at 5%
LH_DFS02_R1 is Synchronized
tst-nda04150cl.sbd is not active on this node
```



## 6.48 Rename Pool

**rename pool** *<pool\_name>* *<new\_pool\_name>*

Rename a specified NSS pool.

```
nlvm [nlvm_options] rename pool <pool_name> <new_pool_name>
```

### Command Option

#### *pool\_name*

Mandatory. Specify the name of the NSS pool to rename.

#### Example

```
MYPOOL1
```

#### *new\_pool\_name*

Mandatory. Specify the new name of the NSS pool.

For pool naming conventions, see the [create pool](#) command.

#### Example

```
P_SALES
```

### Command Example

```
nlvm rename pool MYPOOL1 P_SALES
```

Rename the pool MYPOOL1 as P\_SALES.

## 6.49 Rename RAID

**rename raid** *<raid\_name>* *<new\_raid\_name>*

Rename a specified NSS software RAID device.

If the RAID device is shared, issue the command on the node where the device is currently online.

```
nlvm [nlvm_options] rename raid <raid_name> <new_raid_name>
```

### Command Option

#### *raid\_name*

Mandatory. Specify the name of the NSS software RAID device to rename.

#### Example

```
MYRAID1
```

#### *new\_raid\_name*

Mandatory. Specify the new name of the NSS software RAID device. See the [create raid](#) command for RAID naming conventions.

#### Example

```
R1_SALES
```

### Command Example

```
nlvm rename raid MYRAID1 R1_SALES
```

Rename the NSS software RAID device MYRAID1 as R1\_SALES.

## 6.50 Rename Volume

```
rename volume <volume_name> <new_volume_name> [encryption_password]
```

Rename a specified NSS volume. If the volume is encrypted, you might also need to provide its encryption password.

If a volume's pool is cluster-enabled, issue the command on the node where the pool cluster resource is currently online.

```
nlvm [nlvm_options] rename volume <volume_name> <new_volume_name>
[encryption_password]
```

### Command Option

#### **volume\_name**

Mandatory. Specify the name of the NSS volume to rename.

#### **Example**

```
MYVOL1
```

#### **new\_volume\_name**

Mandatory. Specify the new name of the NSS volume.

Volume names are 2 to 15 characters. The naming conventions are the same as for pools. See the `create pool` command for naming conventions.

#### **Example**

```
V_SALES
```

#### **encryption\_password**

Optional. If the volume is encrypted, the volume's encryption password might be needed. You can try the command without the password. If the password is needed, you are prompted to enter it.

#### **Example**

```
novell
```

### Command Example

```
nlvm rename volume MYVOL1 V_SALES
```

Rename the NSS volume MYVOL1 as V\_SALES.

```
nlvm rename volume MYVOL2 V_FINANCE novell
```

Rename the encrypted NSS volume MYVOL2 as V\_FINANCE. In this example, the encryption password is novell.

## 6.51 Rescan

### rescan

Performs a rescan of the storage objects (such as partitions, NSS pools, and NSS software RAIDs) on known devices, and creates any Device Mapper device or partition objects, or updates them as needed. It also mounts all pools that are not mounted unless you use the `-m` option. There are no command options.

```
nlvm [nlvm_options] rescan
```

#### Command Example

```
nlvm rescan
```

Scans for storage objects, creates and updates Device Mapper objects, and mounts pools as needed.

## 6.52 Resume Move

### resume move <move\_name | pool\_name>

Resume the mirroring for a specified pool move that has been paused with the [nlvm pause move](#) command. The pool move continues from where it was paused.

```
nlvm [nlvm_options] resume move <move_name | pool_name>
```

If the pool is cluster-enabled, you must issue the command on the node where the pool is currently active. You cannot resume a paused pool move while the pool cluster resource is offline. A paused pool move for a clustered pool will resume automatically:

- ♦ If the pool cluster resource fails over to a different node
- ♦ If you cluster migrate the pool cluster resource to a different node
- ♦ If you take the pool cluster resource offline and bring it online again

To re-pause the pool move, use the `nlvm pause move` command.

#### Command Option

##### `move_name`

Mandatory. Specify the name of the paused pool move that you want to resume. The move name typically looks like `POOLNAME_move`.

##### `pool_name`

Mandatory. Specify the name of the NSS pool that is being moved.

#### Example

```
MYPOOL1
```

#### Command Example

```
nlvm resume move MYPOOL1
```

Resume mirroring for the pool move for the pool `MYPOOL1`.

## 6.53 Share

**share** <device\_name>

Set the specified device as shared.

```
nlvm [nlvm_options] share <device_name>
```

### Command Option

**device\_name**

Mandatory. Specify the device to be shared.

You can enter multiple devices by separating the device names with a comma and no spaces.

### Examples

```
sdb  
sde, sdf, sdg
```

### Command Example

```
nlvm share sdb
```

Sets the /dev/sdb device as shared.

```
nlvm share sde, sdf, sdg
```

Sets the /dev/sde, /dev/sdf, and /dev/sdg devices as shared.

## 6.54 Unmount

**unmount** <pool\_name>

Unmount a specified NSS pool. This removes the pool from NSS and causes any open files to be closed and any volumes to be deactivated. It also removes the Device Mapper object for the pool, the link to the Device Mapper object, and the mount point for the pool. This allows you to gracefully log out the server from an iSCSI device that contains a pool. Use this command with caution.

```
nlvm [nlvm_options] unmount <pool_name>
```

### Command Option

**pool\_name**

Mandatory. Specify the name of the NSS pool to unmount.

Use the `unmount` command to temporarily unload a pool in order to manage underlying devices. Pools are by design auto mounted. Therefore, running the `nssmu` utility, or running most `nlvm` commands without the `-m` option can cause an unmounted pool to be remounted if underlying devices and partitions still exist. To execute an `nlvm` command without mounting the unmounted pools, you must include the `-m` option. The `nlvm mount` command internally sets the `-m` flag, so only the specified pool is mounted.

### Example

```
MYPOOL1
```

### Command Example

```
nlvm unmount MYPOOL1
```

Unmount the pool MYPOOL1.

## 6.55 Unshare

**unshare** <device\_name>

Set the specified device as not shared.

```
nlvm [nlvm_options] unshare <device_name>
```

### Command Option

**device\_name**

Mandatory. Specify the device to be unshared.

You can enter multiple devices by separating the device names with a comma and no spaces.

### Examples

```
sdb  
sde,sdf,sdg
```

### Command Example

```
nlvm unshare sdb
```

Sets the /dev/sdb device as not shared.

```
nlvm unshare sde,sdf,sdg
```

Sets the /dev/sde, /dev/sdf, and /dev/sdg devices as not shared.

## 6.56 Volume Mount

**volume mount** <volume\_name> [encryption\_password]

Mount a specified NSS volume. This also activates the volume before mounting it.

```
nlvm [nlvm_options] volume mount <volume_name> [encryption_password]
```

### Command Options

**volume\_name**

Mandatory. Specify the name of the NSS volume to mount.

### Example

```
MYVOL
```

**encryption\_password**

Optional. The password is required to mount an encrypted NSS volume on the first mount after a reboot. Thereafter, the password is stored encrypted in system memory until the next server reboot.

### Example

```
novell
```

### Command Examples

```
nlvm volume mount MYVOL
```

Mount the volume MYVOL.

```
nlvm volume mount MYVOL2 novell
```

Mount the encrypted volume MYVOL2 on the first mount after a reboot. Thereafter until the next reboot, the password is not used to mount the volume. For example:

```
nlvm volume mount MYVOL2
```

## 6.57 Volume Unmount

**volume unmount** <volume\_name>

Dismount a specified NSS volume. This also deactivates the volume before dismounting it.

```
nlvm [nlvm_options] volume unmount <volume_name>
```

### Command Option

**volume\_name**

Mandatory. Specify the name of the NSS volume to dismount.

### Example

```
MYVOL
```

### Command Example

```
nlvm volume unmount MYVOL
```

Dismount the volume MYVOL.

# 7 NLVM Examples for the NSS File System

This section provides examples for using the Novell Linux Volume Manager (NLVM) to manage the OES Storage Services (NSS) file system on your Open Enterprise Server (OES) servers.

For information about using NLVM commands to create and manage Linux POSIX volumes on your OES servers, see [“Managing Linux Volumes with NLVM Commands”](#) in the *OES 2023: Linux POSIX Volume Administration Guide*.

- [Section 7.1, “Creating an NSS Pool and Volume,”](#) on page 111
- [Section 7.2, “Mirroring a Pool Partition,”](#) on page 111
- [Section 7.3, “Recovering a Mirror where All Elements Report ‘Not in Sync,’”](#) on page 112
- [Section 7.4, “Logging Out of an iSCSI Device that Contains an NSS Pool,”](#) on page 113
- [Section 7.5, “Creating a Linux Volume on a Device that Contains a Novell Partition,”](#) on page 114

## 7.1 Creating an NSS Pool and Volume

Enter commands at a terminal command prompt as the `root` user.

Create an NSS pool named `MYPOOL1` with a size of 100 GB on device `/dev/sdb`. Create a volume on the new pool named `MYVOL`.

```
nlvm create pool device=sdb size=100G name=MYPOOL1
```

```
nlvm create volume name=MYVOL pool=MYPOOL1
```

The command to create an NSS pool creates the partition, pool, Device mapper object, (such as `/dev/nss/sdb1.1`), and activates the pool.

The command to create the volume creates the volume and automatically mounts it if the pool is not shared. If the pool is shared and cluster enabled, you must configure the pool cluster resource and use the OES Cluster Services commands to bring the resource and its volume online.

## 7.2 Mirroring a Pool Partition

You can mirror an existing NSS pool partition by using the `Create RAID` command with the `part=<existing_partition_name>` option as follows:

```
nlvm [nlvm_options] create raid
  name=<raid_name>
  raid=1
  [type=nss|sdb]
  part=<pool_partition_name>
  device=<device_name>
```

This command specifies the existing pool partition as the first segment of a RAID1 mirror. You must specify the `device` option one time with the device to use as its mirror. You do not specify a size in the command. The size of the existing partition determines the amount of space that is used for the mirrored segment. The partition type created for the mirror is the same type as the original partition. After you mirror the partition, you manage the RAID1 device by using the normal NSS software RAID management tools and commands.

For example, if `POOL1` uses partition `sdcl.1`, the following command creates an NSS software RAID 1 mirrored device named `POOL1RAID1`. The pool's existing partition becomes the first segment of the RAID, and its existing data is mirrored to device `sdf`.

```
nlvm create raid name=POOL1RAID1 raid=1 part=sdcl.1 device=sdf
```

## 7.3 Recovering a Mirror where All Elements Report 'Not in Sync'

If all elements of a mirrored RAID report a status of "not in sync", use the following procedure to recover the mirror.

- 1 Determine which element you believe to be the in-sync element.
- 2 Log in to the server as the `root` user, then open a terminal console.
- 3 Using the `nlvm delete raid segment` command, remove all of the elements from the mirror except the element you want to keep.

For each element that you want to remove, enter the following command. When you are prompted to confirm, type `yes`, then press Enter. Wait for the segment to be removed before you remove the next segment.

```
nlvm --force delete raid <raid_name> segment <segment_number>
```

Use the `--force` NLVM option to force the deletion of an out-of-sync segment. When you are done, you have a RAID1 device that consists of the single element that you believed to be the in-sync element.

For example, enter

```
nlvm -f delete raid MYRAID1 segment 0
```

When prompted to confirm the deletion, type `yes`, then press Enter.

- 4 Force the single RAID element to be in sync. At the command prompt, enter

```
nlvm raid force <raid_name>
```

- 5 Add elements back into the mirror as desired by using the `nlvm raid expand` command. At the command prompt, enter

```
nlvm expand raid <raid_name> device=<device_name>
```

The `device` option can be specified multiple times to specify additional segments.



## 7.4 Logging Out of an iSCSI Device that Contains an NSS Pool

Before you log out of an iSCSI device that is used for an NSS pool, you must first unmount the volumes, deactivate the pool, and unmount the pool. Log out of iSCSI immediately after you unmount the pool.

---

**IMPORTANT:** The `nlvm unmount <poolname>` command removes the pool's Device Mapper object and allows the device to be disconnected gracefully. Otherwise, a server hang can occur.

---

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

```
nssmu
```

- 3 Dismount the volumes on the pool.
  - 3a In the NSSMU main menu, select **Volumes**, then press Enter.
  - 3b Select the volume, then press F7 to dismount it.
  - 3c If the pool contains multiple volumes, repeat [Step 3b](#) for each volume.
  - 3d Press Esc to exit the Volumes page.
- 4 Deactivate the pool.
  - 4a In the NSSMU main menu, select **Pools**, then press Enter.
  - 4b Select the pool, then press F7 to deactivate it.
  - 4c Press Esc to exit the Pools page.
- 5 Press Esc to exit NSSMU.

Ensure that you have exited NSSMU before you continue. It is essential that there be no cached states for device, partition, and pool objects within NSSMU.

- 6 Use NLVM to unmount the pool.

```
nlvm unmount <poolname>
```

An unmounted pool is a temporary state. You must log out of the iSCSI connection immediately after executing the `nlvm unmount <poolname>` command before any NLVM or NSSMU command is executed. As soon as NSSMU is run, NSSMU remounts the pool in order to manage it. In addition, almost any NLVM command that is run after the unmount also causes the pool to be remounted unless you use the `-m` option.

- 7 Log out of the iSCSI connection.
  - 7a Launch YaST to manage the iSCSI client.

```
yast2 iscsi-client
```

- 7b Select the **Connected Targets** tab, then select the iSCSI device and click **Logout**.

## 7.5 Creating a Linux Volume on a Device that Contains a Novell Partition

As a best practice, disks using Novell partitions should have only Novell partitions on the device. If you mix Novell and Linux partition types on the same device, the recommended method is to create a Linux volume first, and then create the NSS pool.

In OES 11 SP2 and later, you can use the following procedure to create a Linux partition on a device that already contains a Novell type partition, and then specify the Linux partition as the location for a non-clustered Linux volume.

To add a Linux volume to an unshared device with an existing NSS partition and pool on it:

- 1 Log in to the server as the `root` user, then open a terminal console.
- 2 Create a Linux partition on the device. Enter

```
nlvm create partition type=<83|8E> device=<node_name> size=<value[K|M|G|T]|max>
```

Specify the partition type based on the type of Linux volume you plan to create.

```
type=83                (Linux native volume)
type=8E                (Linux LVM volume)
```

For example, to create an LVM partition type on device `sdd` that is 500 GB, enter

```
nlvm create partition type=8E device=sdd size=500G
```

- 3 Unmount all NSS pools on the device. Enter

```
nlvm unmount pool <poolname>
```

For example, to dismount POOL1 and POOL2 on device `sdd`, enter

```
nlvm unmount pool POOL1
nlvm unmount pool POOL2
```

- 4 Do any one of the following to allow NLVM to recognize the new Linux storage object on the device for Device Mapper:

- ◆ Mount the pools on the device. For each pool, enter

```
nlvm mount pool <poolname>
```

- ◆ Rescan the device for storage objects and allow NLVM to automatically mount all pools on the device.

```
nlvm rescan
```

- ◆ Restart the server.

- 5 Create a non-clustered Linux volume on the new partition.

```
nlvm create linux volume type=<fs_type> part=<partition_name> mp=<mount_path>
[mkopt=<opt1[,opt2,...]>] [mntopt=<opt1[opt2]...>]
[lvm] [name=<lvm_volume_name>] [group=<lvm_vg_name>]
[ncp]
```

The volume type must match the type of partition you created in [Step 2](#).

Continuing the example, on a type 8E partition named `sdd3`, create an Ext3 file system on an ncp-enabled LVM logical volume named `MYVOL3`. Enter:

```
nlvm create linux volume type=ext3 part=sdd3 mp=/usr/novell/lvm/myvol3 mntopt=rw lvm  
name=MYVOL3 ncp
```



# 8

## NLVM Examples for Clustering with OES Cluster Services

This section provides examples for using the Novell Linux Volume Manager (NLVM) with OES Cluster Services on your Open Enterprise Server (OES) servers.

- ♦ [Section 8.1, “Creating or Mirroring an SBD Partition,” on page 117](#)
- ♦ [Section 8.2, “Unmirroring a Mirrored SBD Partition with NLVM,” on page 129](#)
- ♦ [Section 8.3, “Deleting an SBD Partition with NLVM,” on page 130](#)

### 8.1 Creating or Mirroring an SBD Partition

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

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

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

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

Typically, you use the OES Cluster Services SBD Utility (`sbdutil`) to create or delete an SBD partition for a cluster, as described in [“Creating or Deleting Cluster SBD Partitions”](#) in the *OES 2023: OES Cluster Services for Linux Administration Guide*. However, you can also use NLVM commands in OES 11 SP1 and later to create or delete SBD partitions.

Use the procedures in this section to create a non-mirrored or mirrored SBD partition:

- ♦ [Section 8.1.1, “Requirements and Guidelines for Creating an SBD Partition,” on page 118](#)
- ♦ [Section 8.1.2, “Creating a Non-Mirrored SBD Partition with NLVM,” on page 120](#)
- ♦ [Section 8.1.3, “Mirroring an Existing SBD Partition with NLVM,” on page 124](#)
- ♦ [Section 8.1.4, “Creating a Mirrored SBD Partition with NLVM,” on page 126](#)

## 8.1.1 Requirements and Guidelines for Creating an SBD Partition

Consider the requirements and guidelines in this section when you create a OES Cluster Services SBD (split-brain detector) partition.

- ♦ “Preparing OES Cluster Services” on page 118
- ♦ “Using a Shared Disk System” on page 118
- ♦ “Preparing a SAN Device” on page 118
- ♦ “Working with NLVM Commands in a Cluster” on page 119
- ♦ “Initializing and Sharing a Device for the SBD” on page 119
- ♦ “Determining the SBD Partition Size” on page 119
- ♦ “Replacing an Existing SBD Partition” on page 119

### Preparing OES Cluster Services

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

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

### Using a Shared Disk System

You must have a shared disk system (such as a Fibre Channel SAN or an iSCSI SAN) connected to your cluster nodes before you create a split-brain-detector (SBD) partition. For information, see “[Shared Disk Configuration Requirements](#)” in the *OES 2023: OES Cluster Services for Linux Administration Guide*.

### Preparing a SAN Device

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

For device fault tolerance, you can use the `nlvm create raid` command to mirror the SBD partition on another SAN device. Before you mirror the device, you must carve a second LUN of the same size, and connect the LUN device to all nodes in the cluster.

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

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

```
nlvm -s rescan
```

## Working with NLVM Commands in a Cluster

If an SBD does not exist in the cluster, NLVM cannot detect whether a node is a member of the cluster, and therefore, it cannot get exclusive locks to the physically shared storage. In this state, you must use the `-s` NLVM option to override the shared locking requirement and force NLVM to execute the commands you use to create the SBD partition. To minimize the risk of possible corruption, you are responsible for ensuring that you have exclusive access to the shared storage at this time.

## Initializing and Sharing a Device for the SBD

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

When you use NLVM to create an SBD, the `nlvm create partition` command can accept an initialized or uninitialized device when you use the `type=lad` option. NLVM checks the specified device to see if it is initialized, and takes the following actions:

- ♦ **Uninitialized device:** NLVM initializes the device, marks it as Shareable for Clustering, and creates the requested SBD partition.
- ♦ **Initialized and shared device:** NLVM creates the requested SBD partition.
- ♦ **Initialized and unshared device:** NLVM creates the requested SBD partition, but does not alter the shared state. It returns an error warning that the SBD partition is not shared.

You must manually mark the device as Shareable for Clustering after the partition is created. You can use the `nlvm share` command to share the device.

## Determining the SBD Partition Size

When you create the SBD partition by using the `nlvm create partition` command, you can specify how much free space to use for the SBD, or you can specify the `max` option to use the entire device. If you specify a device to use as a mirror, the same amount of space is used. If you specify to use the maximum size and the mirror device is bigger than the SBD device, you will not be able to use the excess free space on the mirror for other purposes.

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

## Replacing an Existing SBD Partition

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

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

## 8.1.2 Creating a Non-Mirrored SBD Partition with NLVM

Use the procedure in this section to create a new SBD partition. If an SBD partition already exists, you must first delete the SBD as described in [Section 8.3, “Deleting an SBD Partition with NLVM,”](#) on page 130.

- 1 Ensure that nobody else is changing any storage on any nodes at this time.

Until the SBD exists and the cluster is set up for shared disk access, you are responsible for ensuring that you have exclusive access to the shared storage.

- 2 Take the cluster down:

- 2a Log in to any node in the cluster as the `root` user, then open a terminal console.

- 2b At the command prompt, enter

```
cluster down
```

- 3 On each cluster node, stop OES Cluster Services:

- 3a Log in to the cluster node as the `root` user, then open a terminal console.

- 3b At the command prompt, enter

```
rcnovell-ncs stop
```

- 3c After you have stopped OES Cluster Services on all nodes, continue with the next step.

- 4 Prepare a SAN device to use for the SBD partition:

- 4a Use the SAN storage array software to carve a device to use exclusively for the SBD partition.

- 4b Attach the device to all nodes in the cluster.

- 4c On each node, log in as the `root` user and rescan for devices:

```
nlvm -s rescan
```

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

- 5 Log in to any node in the cluster as the `root` user, then open a terminal console.

- 6 View a list of the devices and identify the leaf node name (such as `sdC`) of the SAN device that you want to use for the SBD partition. At the command prompt, enter

```
nlvm -s list devices --terse
```

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

The device information shows the leaf node name, the size, the amount of free available space, the partitioning format (such as MSDOS or GPT), the shared state (whether it is marked as Shareable for Clustering), and the RAID state (whether the device is an NSS software RAID device). Do not use an NSS software RAID for the device.



For example, the uninitialized device `sdc` reports a used and free size of 0 KB, a format of `None`, and a shared state of `No`:

| Name             | Size     | Used     | Free     | Format | Shared | RAID | Enabled |
|------------------|----------|----------|----------|--------|--------|------|---------|
| <code>sda</code> | 20.00GB  | 19.99GB  | 1008KB   | MSDOS  | No     | No   |         |
| <code>sdb</code> | 1.00GB   | 400.01MB | 623.98MB | MSDOS  | No     | No   |         |
| <code>sdc</code> | 102.00MB | 0KB      | 0KB      | None   | No     | No   |         |
| <code>sdd</code> | 8.00GB   | 50.01MB  | 7.95GB   | MSDOS  | Yes    | No   |         |

**7 Initialize and share the device.** At the command prompt, enter

```
nlvm -s init <device_name> format=msdos shared
```

---

**WARNING:** Initializing a device destroys all data on the device.

---

Replace `device_name` with the leaf node name (such as `sdc`) of the SAN device you want to use as the SBD partition.

Specify a partitioning format of `msdos`.

Specify the `shared` option to mark the device as Shareable for Clustering.

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

You can list the devices to visually verify that the device is formatted and shared:

```
nlvm -s list devices
```

For example, the formatted device `sdc` reports values for used and free size, a format of `MSDOS`, and a shared state of `Yes`:

| Name             | Size     | Used     | Free     | Format | Shared | RAID | Enabled |
|------------------|----------|----------|----------|--------|--------|------|---------|
| <code>sda</code> | 20.00GB  | 19.99GB  | 1008KB   | MSDOS  | No     | No   |         |
| <code>sdb</code> | 1.00GB   | 400.01MB | 623.98MB | MSDOS  | No     | No   |         |
| <code>sdc</code> | 102.00MB | 16KB     | 101.98MB | MSDOS  | Yes    | No   |         |
| <code>sdd</code> | 8.00GB   | 50.01MB  | 7.95GB   | MSDOS  | Yes    | No   |         |

**8 Create the SBD partition.** At the command prompt, enter (all on the same line):

```
nlvm -s create partition
      type=lad
      device=<device_name>
      size=<value[K|M|G|T]|max>
      label="<cluster_name>"
```

Specify a type of `lad` to create the SBD partition type.

Replace `device_name` with the leaf node name (such as `sdc`) of the SAN device you want to use as the SBD partition.

Replace `value` with the amount of space to use for the SBD partition and select a unit of measure as its multiplier, or specify `max` to use the entire device. If you specify a value without a multiplier, gigabytes (G) is assumed.

Replace `cluster_name` with the name of the cluster, such as `cluster1`. This name must match the name of an existing cluster that has a Cluster object in eDirectory. The name is case sensitive.

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

For example, to create an SBD partition for a cluster named `cluster1` on device `sdc` that has already been initialized and shared, enter

```
nlvm -s create partition type=1ad device=sdc size=max label="cluster1"
```

A partition is created named `cluster1.sbd`. It uses all available free space on the specified device.

- 9 View a list of partitions and verify that the new partition appears in the list. At the command prompt, enter

```
nlvm -s list partitions
```

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

The partition information shows the partition name, the leaf node name of the device, the partition type (1AD), the starting location, and the partition size. Because an SBD partition must end at a cylinder boundary, the partition size might be slightly smaller than the device size, or the size you specified for the partition.

For example, for device `sdc` that is 102 MB in size, the partition created is 99.59 MB in size:

| Name         | Type          | Start   | Size     | Device |
|--------------|---------------|---------|----------|--------|
| sda1         | 83 (Linux)    | 2048    | 297.00MB | sda    |
| sda2         | 82Linux_Swap) | 610304  | 1.00GB   | sda    |
| sda3         | 83 (Linux)    | 2715648 | 7.99GB   | sda    |
| cluster1.sbd | 1AD(Cluster)  | 32      | 99.59MB  | sdc    |

If you specified the maximum size for the SBD partition, you can list devices again to see that all space on the device is used for the SBD partition:

| Name | Size     | Used     | Free     | Format | Shared | RAID | Enabled |
|------|----------|----------|----------|--------|--------|------|---------|
| sda  | 20.00GB  | 19.99GB  | 1008KB   | MSDOS  | No     | No   |         |
| sdb  | 1.00GB   | 400.01MB | 623.98MB | MSDOS  | No     | No   |         |
| sdc  | 102.00MB | 102.00MB | 0KB      | MSDOS  | Yes    | No   |         |
| sdd  | 8.00GB   | 50.01MB  | 7.95GB   | MSDOS  | Yes    | No   |         |

- 10 Modify the Cluster object in eDirectory to enable its NCS: Shared Disk Flag attribute.

This step is required only if the cluster has never had an SBD partition. However, it does no harm to verify that the **NCS: Shared Disk Flag** attribute is enabled.

- 10a In a web browser, open iManager, then log in to the eDirectory tree that contains the cluster you want to manage.

---

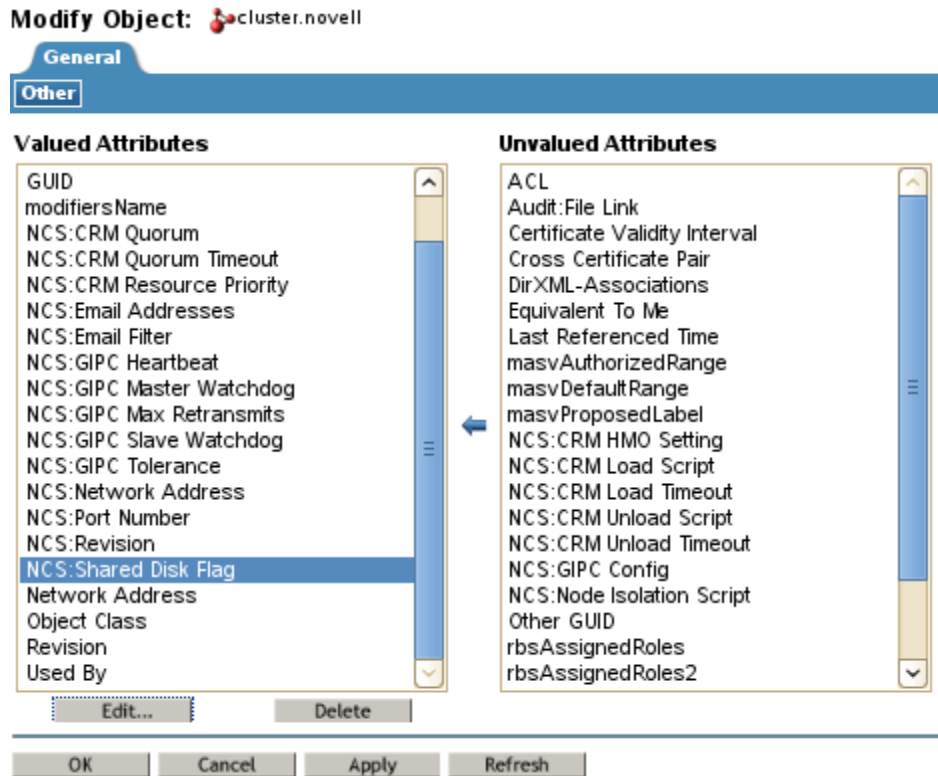
**IMPORTANT:** Log in as an administrator user who has sufficient rights in eDirectory to delete and modify eDirectory objects.

---

- 10b Select **Directory Administration**, then select **Modify Object**.

- 10c Browse to locate and select the Cluster object of the cluster you want to manage, then click **OK**.

**10d** Under **Valued Attributes**, select the **NCS: Shared Disk Flag**, then click **Edit**.



**10e** Select (enable) the **NCS: Shared Disk Flag** check box, then click **OK**.



**10f** Click **Apply** to save changes.

**11** On each cluster node, start OES Cluster Services:

**11a** Log in to the cluster node as the `root` user, then open a terminal console.

**11b** At the command prompt, enter

```
rcnovell-ncs start
```

**11c** After you have restarted OES Cluster Services on all nodes, continue with the next step.

**12** On each cluster node, join the cluster. At the command prompt, enter

```
cluster join
```

**13** (Optional) Continue with [Section 8.1.3, “Mirroring an Existing SBD Partition with NLVM,”](#) on page 124.

## 8.1.3 Mirroring an Existing SBD Partition with NLVM

You can mirror an existing OES Cluster Services SBD partition to provide device fault tolerance. It is not necessary to take the cluster down or stop the cluster software.

- 1 Prepare a SAN device to use as the mirror segment for the SBD partition:
  - 1a Use the SAN storage array software to carve a device that is at least the size of the existing SBD partition's device.
  - 1b Attach the device to all nodes in the cluster.
  - 1c On each node, log in as the `root` user and rescan for devices:

```
nlvm rescan
```

- 2 Log in to any member node of the cluster as the `root` user, then open a terminal console.
- 3 View a list of the devices and identify the leaf node name (such as `sde`) of the SAN device that you want to use as the mirror for the existing SBD partition. At the command prompt, enter

```
nlvm list devices
```

For example, the uninitialized device `sde` reports a used and free size of 0 KB, a format of `None` and a shared state of `No`:

| Name | Size     | Used     | Free     | Format | Shared | RAID | Enabled |
|------|----------|----------|----------|--------|--------|------|---------|
| sda  | 20.00GB  | 19.99GB  | 1008KB   | MSDOS  | No     | No   |         |
| sdb  | 1.00GB   | 400.01MB | 623.98MB | MSDOS  | No     | No   |         |
| sdc  | 102.00MB | 102.00MB | 0KB      | MSDOS  | Yes    | No   |         |
| sdd  | 8.00GB   | 50.01MB  | 7.95GB   | MSDOS  | Yes    | No   |         |
| sde  | 102.00MB | 0KB      | 0KB      | None   | No     | No   |         |

- 4 Initialize and share the device. At the command prompt, enter

```
nlvm init <device_name> format=msdos shared
```

---

**WARNING:** Initializing a device destroys all data on the device.

---

Replace `device_name` with the leaf node name (such as `sde`) of the SAN device you want to use as the mirror for the existing SBD partition.

Specify a partitioning format of `msdos`.

Specify the `shared` option to mark the device as Shareable for Clustering.

You can list the devices to visually verify that the device is formatted and shared:

```
nlvm list devices
```

For example, the formatted device `sde` reports a format of `MSDOS` and a shared state of `Yes`:

| Name | Size     | Used     | Free     | Format | Shared | RAID | Enabled |
|------|----------|----------|----------|--------|--------|------|---------|
| sda  | 20.00GB  | 19.99GB  | 1008KB   | MSDOS  | No     | No   |         |
| sdb  | 1.00GB   | 400.01MB | 623.98MB | MSDOS  | No     | No   |         |
| sdc  | 102.00MB | 102.00MB | 0KB      | MSDOS  | Yes    | No   |         |
| sdd  | 8.00GB   | 50.01MB  | 7.95GB   | MSDOS  | Yes    | No   |         |
| sde  | 102.00MB | 16KB     | 101.98MB | MSDOS  | Yes    | No   |         |

- 5 Mirror the SBD partition. At the command prompt, enter (all on the same line):

```
nlvm create raid
    raid=1
    name=<cluster_name>
    type=sd
    part=<partition_name>
    device=<device_name>
```

Specify a RAID type of 1 for mirroring.

Replace **cluster\_name** with the name of the SBD's cluster, such as `cluster1`. This name must match the name of an existing cluster that has a Cluster object in eDirectory. The name is case sensitive.

Specify a type of `sd` to create SBD partitions on the RAID1 device. The `type` option must precede the `part` option in the command.

Replace *partition\_name* with the partition name of the existing SBD partition.

Replace *device\_name* with the leaf node name (such as `sde`) of the SAN device you want to use as the mirror for the existing SBD partition. The device must be at least the same size as the partition you want to mirror.

You do not specify a size in the command. The size of the existing partition determines the amount of space that is used for the mirrored segment.

For example, to mirror the SBD partition `cluster1.sd` with device `sde` for a cluster named `cluster1`, enter

```
nlvm create raid raid=1 type=sd name=cluster1 part=cluster1.sd
device=sde
```

For our example, a RAID1 (mirror) device is created named `cluster1.sd` that is made up of device `sdc` and device `sde`. The existing SBD partition is renamed from `cluster1.sd` to `cluster1.msbd0`. A new partition named `cluster1.msbd1` is created on device `sde`.

- 6** View a list of devices to verify the current state of both devices and to verify that a RAID1 device named `cluster1.sd` was created. At the command prompt, enter

```
nlvm list devices
```

The entries of interest in the devices list are the devices that you use for the SBD partition (such as `sdc` and `sde`) and the newly created RAID1 device:

| Name        | Size     | Used     | Free     | Format | Shared | RAID | Enabled |
|-------------|----------|----------|----------|--------|--------|------|---------|
| sda         | 20.00GB  | 19.99GB  | 1008KB   | MSDOS  | No     | No   |         |
| sdb         | 1.00GB   | 400.01MB | 623.98MB | MSDOS  | No     | No   |         |
| sdc         | 102.00MB | 102.00MB | 0KB      | MSDOS  | Yes    | No   |         |
| sdd         | 8.00GB   | 50.01MB  | 7.95GB   | MSDOS  | Yes    | No   |         |
| sde         | 102.00MB | 102.00MB | 0KB      | MSDOS  | Yes    | No   |         |
| cluster1.sd | 99.57MB  | 97.57MB  | 0KB      | None   | Yes    | 1    | Yes     |

- 7** View a list of partitions to verify the status of mirrored SBD partitions `cluster1.msbd0` and `cluster1.msbd1`. At the command prompt, enter

```
nlvm list partitions
```

The entries of interest in the list are `cluster1.msbd0` and `cluster1.msbd1`:

| Name           | Type          | Start   | Size     | Device |
|----------------|---------------|---------|----------|--------|
| sda1           | 83(Linux)     | 2048    | 297.00MB | sda    |
| sda2           | 82Linux Swap) | 610304  | 1.00GB   | sda    |
| sda3           | 83(Linux)     | 2715648 | 7.99GB   | sda    |
| cluster1.msbd0 | 1AD(Cluster)  | 32      | 99.59MB  | sdc    |
| cluster1.msbd1 | 1AD(Cluster)  | 32      | 99.59MB  | sde    |

## 8.1.4 Creating a Mirrored SBD Partition with NLVM

You can create a mirrored OES Cluster Services SBD partition to provide device fault tolerance for the SBD. You must take the cluster down and stop the cluster software. If an SBD partition already exists, you must first delete the SBD as described in [Section 8.3, “Deleting an SBD Partition with NLVM,”](#) on page 130.

Use the procedure in this section to create a new mirrored SBD partition by using NLVM commands.

- 1 Ensure that nobody else is changing any storage on any nodes at this time.

Until the SBD exists and the cluster is set up for shared disk access, you are responsible for ensuring that you have exclusive access to the shared storage.

- 2 Take the cluster down:

- 2a Log in to any node in the cluster as the `root` user, then open a terminal console.

- 2b At the command prompt, enter

```
cluster down
```

- 3 On each cluster node, stop OES Cluster Services:

- 3a Log in to the cluster node as the `root` user, then open a terminal console.

- 3b At the command prompt, enter

```
rcnovell-ncs stop
```

- 3c After you have stopped OES Cluster Services on all nodes, continue with the next step.

- 4 Prepare two SAN devices to use for the mirrored SBD partition:

- 4a Use the SAN storage array software to carve two devices of equal size to use exclusively for the mirrored SBD partition.

- 4b Attach the devices to all nodes in the cluster.

- 4c On each node, log in as the `root` user and rescan for devices:

```
nlvm -s rescan
```

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

- 5 Log in to any node in the cluster as the `root` user, then open a terminal console.

- 6 View a list of the devices and identify the leaf node name (such as `sdC`) of the two SAN devices that you want to use for the mirrored SBD partition. At the command prompt, enter

```
nlvm -s list devices
```

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

The device information shows the leaf node name, the size, the amount of free available space, the partitioning format (such as MSDOS or GPT), the shared state (whether it is marked as Shareable for Clustering), and the RAID state (whether the device is an NSS software RAID device). Do not use an NSS software RAID for the device.

- 7 Initialize and share the two devices. At the command prompt, enter

```
nlvm -s init <device_name1>,<device_name2> format=msdos shared
```

---

**WARNING:** Initializing a device destroys all data on the device.

---

Replace *device\_name1* and *device\_name2* with the leaf node names (such as *sdc* and *sdd*) of the two SAN devices you want to use for the mirrored SBD partition.

Specify a partitioning format of *msdos*.

Specify the *shared* option to mark the devices as Shareable for Clustering.

Use the *-s* NLVM option to override the shared locking requirement and force the command to execute.

For example, to initialize devices *sdc* and *sdd*, enter

```
nlvm -s init sdc,sdd format=msdos shared
```

You can list the devices to visually verify that the device is formatted and shared:

```
nlvm -s list devices
```

- 8 Create the mirrored SBD partition. At the command prompt, enter (all on the same line):

```
nlvm -s create raid
raid=1
type=sd
name=<cluster_name>
device=<device_name1> size=max
device=<device_name2>
```

Specify a RAID type of 1 for mirroring.

Specify a type of *sd* to create SBD partitions on the RAID1 device.

Replace **cluster\_name** with the name of the cluster, such as *cluster1*. This name must match the name of an existing cluster that has a Cluster object in eDirectory. The name is case sensitive.

Replace *device\_name1* and *device\_name2* with the leaf node names (such as *sdc* and *sdd*) of the two SAN devices you want to use for the mirrored SBD partition. The *cluster1.msbd0* mirrored SBD partition is created on the first device option instance in the command. The *cluster1.msbd1* mirrored SBD partition is created on the second device option instance in the command.

Specify a size of *max* to use all of the available space. Specify the size only once. Both devices should be the same size; however, if they are not, the size of the RAID segments is determined by the size of the smaller device.

Use the *-s* NLVM option to override the shared locking requirement and force the command to execute.

For example, to create a mirrored SBD for a cluster named *cluster1* with devices *sdc* and *sdd* that have already been initialized and shared, enter

```
nlvm -s create raid raid=1 type=sd name="cluster1" device=sdc size=max
device=sdd
```

A RAID1 device is created named *cluster1.sbd*. The *cluster1.msbd0* partition is created on device *sdc*. The *cluster1.msbd1* partition is created on device *sdd*.

- 9 View a list of devices to verify the current state of both devices and to verify that a RAID1 device named *cluster1.sbd* was created. At the command prompt, enter

```
nlvm -s list devices
```

- 10 View a list of partitions and verify that the new partitions appear in the list. At the command prompt, enter

```
nlvm -s list partitions
```

The partition information shows the partition name, the leaf node name of the device, the partition type (1AD), the starting location, and the partition size. Because an SBD partition must end at a cylinder boundary, the partition size might be slightly smaller than the device size, or the size you specified for the partition.

You can list devices again to see the amount of space that is unused beyond the cylinder boundary. Our example devices show 2.39 MB of free space after the partition is created, as shown in [Step 9](#).

- 11 Modify the Cluster object in eDirectory to enable its NCS: Shared Disk Flag attribute.

This step is required only if the cluster has never had an SBD partition. However, it does no harm to verify that the NCS: Shared Disk Flag attribute is enabled.

- 11a In a web browser, open iManager, then log in to the eDirectory tree that contains the cluster you want to manage.

---

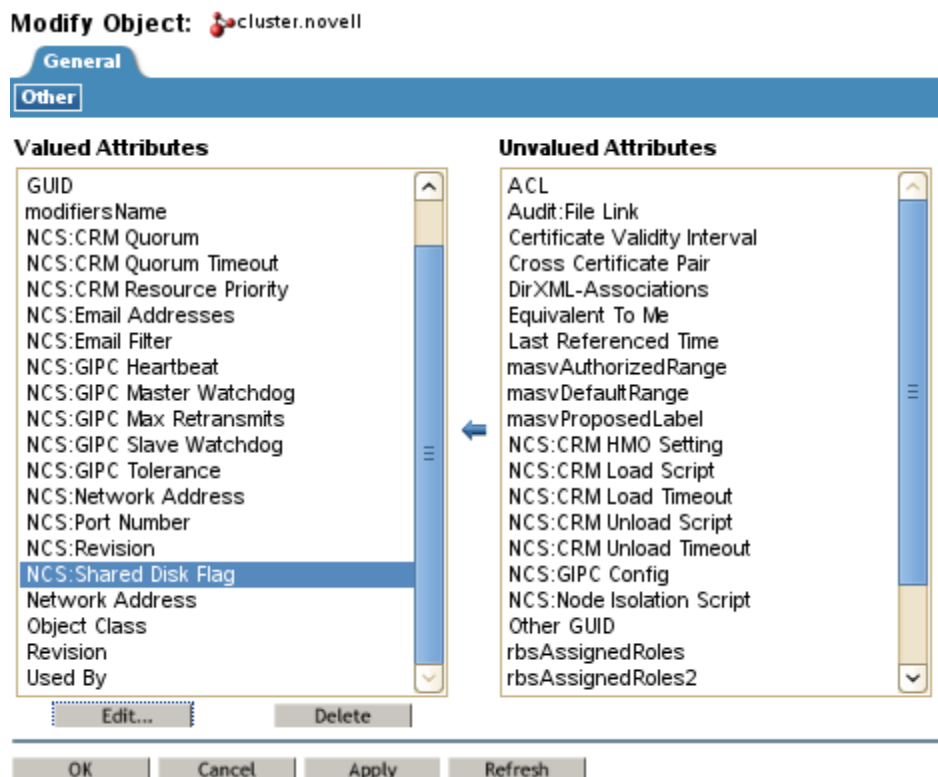
**IMPORTANT:** Log in as an administrator user who has sufficient rights in eDirectory to delete and modify eDirectory objects.

---

- 11b Select **Directory Administration**, then select **Modify Object**.

- 11c Browse to locate and select the Cluster object of the cluster you want to manage, then click **OK**.

- 11d Under **Valued Attributes**, select the **NCS: Shared Disk Flag**, then click **Edit**.





- 11e Select (enable) the **NCS: Shared Disk Flag** check box, then click **OK**.



- 11f Click **Apply** to save changes.
- 12 On each cluster node, start OES Cluster Services:
  - 12a Log in to the cluster node as the `root` user, then open a terminal console.
  - 12b At the command prompt, enter

```
rcnovell-ncs start
```
  - 12c After you have restarted OES Cluster Services on all nodes, continue with the next step.
- 13 On each cluster node, join the cluster. At the command prompt, enter

```
cluster join
```

## 8.2 Unmirroring a Mirrored SBD Partition with NLVM

Use the procedure in this section to remove the mirrored segment from a mirrored SBD partition, and then to remove the single element mirror from the SBD. This leaves a single device that contains an SBD partition.

- 1 Log in to any node as the `root` user, then launch a terminal console.
- 2 Delete the mirrored segment from the mirrored SBD partition. At the command prompt, enter

```
nlvm [--force] [--no-prompt] delete raid <raid_name> segment  
<segment_number>
```

You are automatically prompted to confirm the delete action. Respond by typing `yes` or `no`, then press Enter. You can use the `--no-prompt` NLVM option to suppress the confirmation prompt.

Replace `raid_name` with the name of the mirrored SBD RAID device that contains the segment to be deleted, such as `cluster1.sbd`. The RAID name is case sensitive.

Replace `segment_number` with the segment index (zero relative) to be removed. For a mirrored SBD RAID the possible values are 0 and 1.

Use the `--force` NLVM option to remove out-of-sync segments.

For example, to delete segment 1 of the `cluster1.sbd` RAID1 device, enter

```
nlvm delete raid cluster1.sbd segment 1
```

- 3 View a list of partitions and verify that the SBD partition named `<cluster_name>.msbd1` has been deleted. At the command prompt, enter

```
nlvm list partitions
```

- 4 View a list of RAID1 devices and verify that the SBD RAID1 device `<cluster_name>.sbd` still exists. At the command prompt, enter

```
nlvm list devices exclude=nonraid
```

- 5 Delete the single element mirror from the SBD. At the command prompt, enter

```
nlvm [--no-prompt] delete raid <raid_name>
```

You are automatically prompted to confirm the delete action. Respond by typing `yes` or `no`, then press Enter. You can use the `--no-prompt` NLVM option to suppress the confirmation prompt.

Replace *raid\_name* with the name of the mirrored SBD RAID device that contains the segment to be deleted, such as `cluster1.sbd`. The RAID name is case sensitive.

Because the RAID device is now a single element RAID1, this command removes the single element mirror from the SBD, and leaves the SBD partition on the device. The SBD partition is renamed from `<cluster_name>.msbd0` to `<cluster_name>.sbd`, and the RAID1 device `<cluster_name>.sbd` is deleted.

- 6 View a list of RAID1s and verify that the SBD RAID1 device `<cluster_name>.sbd` has been removed. At the command prompt, enter

```
nlvm list devices exclude=nonraid
```

- 7 View a list of partitions and verify that the SBD partition name has been changed from `<cluster_name>.msbd0` to `<cluster_name>.sbd`. At the command prompt, enter

```
nlvm list partitions
```

For example, the SBD partition entry is now:

```
cluster1.sbd device=sdC type=1AD(Cluster) start=32 size=99.59MB(203968)
```

## 8.3 Deleting an SBD Partition with NLVM

You might need to delete and re-create a OES Cluster Services SBD partition if the SBD becomes corrupted or its device fails. Use the procedure in this section to delete the SBD partition, and then to create a new SBD partition by using one of the methods in [Section 8.1, “Creating or Mirroring an SBD Partition,” on page 117](#).

---

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

---

- 1 Ensure that nobody else is changing any storage on any nodes at this time.

Until the SBD exists and the cluster is set up for shared disk access, you are responsible for ensuring that you have exclusive access to the shared storage.

- 2 Take the cluster down:

- 2a Log in to any node in the cluster as the `root` user, then open a terminal console.

- 2b At the command prompt, enter

```
cluster down
```

**3** On each cluster node, stop OES Cluster Services:

**3a** Log in to the cluster node as the `root` user, then open a terminal console.

**3b** At the command prompt, enter

```
rcnovell-ncs stop
```

**3c** After you have stopped OES Cluster Services on all nodes, continue with the next step.

**4** Log in to any node in the cluster as the `root` user, then launch a terminal console.

**5** If the SBD partition is mirrored, unmirror the SBD partition:

**5a** Delete the mirrored segment from the mirrored SBD partition. At the command prompt, enter

```
nlvm -s [--force] [--no-prompt] delete raid <raid_name> segment <segment_number>
```

You are automatically prompted to confirm the delete action. Respond by typing `yes` or `no`, then press Enter. You can use the `--no-prompt` NLVM option to suppress the confirmation prompt.

Replace `raid_name` with the name of the mirrored SBD RAID device that contains the segment to be deleted, such as `cluster1.sbd`. The RAID name is case sensitive.

Replace `segment_number` with the segment index (zero relative) to be removed. For a mirrored SBD RAID the possible values are 0 and 1.

Use the `--force` NLVM option to remove out-of-sync segments.

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

For example, to delete segment 1 of the `cluster1.sbd` RAID1 device, enter

```
nlvm -s --force delete raid cluster1.sbd segment 1
```

**5b** Delete the single element mirror from the SBD. At the command prompt, enter

```
nlvm -s [--no-prompt] delete raid <raid_name>
```

Because the RAID device is now a single element RAID1, this command removes the single element mirror from the SBD, and leaves the SBD partition on the device.

You are automatically prompted to confirm the delete action. Respond by typing `yes` or `no`, then press Enter. You can use the `--no-prompt` NLVM option to suppress the confirmation prompt.

Replace `raid_name` with the name of the mirrored SBD RAID device that contains the segment to be deleted, such as `cluster1.sbd`. The RAID name is case sensitive.

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

For example, to delete the single element mirror from the `cluster1.sbd` device, enter

```
nlvm -s delete raid cluster1.sbd
```

**5c** View a list of RAID devices and verify that the SBD RAID1 device `<cluster_name>.sbd` has been removed. At the command prompt, enter

```
nlvm -s list raids
```

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

- 5d** View a list of partitions and verify that the SBD partition name has been changed from `<cluster_name>.msbd0` to `<cluster_name>.sbd`. At the command prompt, enter

```
nlvm -s list partitions
```

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

For example, the SBD partition entry is now:

```
cluster1.sbd device=sdc type=1AD(Cluster) start=32
size=99.59MB(203968)
```

- 6** Delete the SBD partition. At the command prompt, enter

```
nlvm -s delete partition <partition_name>
```

Replace *partition\_name* with the name of the SBD partition, such as `cluster1.sbd`. The partition name is case sensitive.

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

For example, to delete the single element mirror from the `cluster1.sbd` device, enter

```
nlvm -s delete partition cluster1.sbd
```

- 7** If you plan to reuse the device for the SBD, initialize and share the device. At the command prompt, enter

```
nlvm -s init <device_name> format=msdos shared
```

---

**WARNING:** Initializing a device destroys all data on the device.

---

Replace *device\_name* with the leaf node name (such as `sde`) of the SAN device.

Specify a partitioning format of `msdos`.

Specify the `shared` option to mark the device as Shareable for Clustering.

Use the `-s` NLVM option to override the shared locking requirement and force the command to execute.

You can list the devices to visually verify that the device is formatted and shared:

```
nlvm -s list devices
```

- 8** To re-create the SBD partition, continue with [Section 8.1, “Creating or Mirroring an SBD Partition,” on page 117](#).

Do not restart OES Cluster services and rejoin nodes to the cluster until after you create the new SBD.

# 9 Troubleshooting NLVM

This section identifies common problems and troubleshooting tips for Novell Linux Volume Manager (NLVM) on your Open Enterprise Server (OES) server.

- ♦ Section 9.1, “Viewing Error Code Messages,” on page 133
- ♦ Section 9.2, “Failure to Create an LVM Volume Group,” on page 133
- ♦ Section 9.3, “Failure to Create a Clustered LVM Volume Group,” on page 134
- ♦ Section 9.4, “Device Is Not Available for Use in an LVM Volume Group,” on page 134
- ♦ Section 9.5, “NLVM Pool Move Fails and Deactivates the Pool,” on page 134
- ♦ Section 9.6, “Error 20897 - This node is not a cluster member,” on page 134
- ♦ Section 9.7, “NLVM Error Codes,” on page 135
- ♦ Section 9.8, “NSS Error Codes,” on page 144

For additional troubleshooting information, see the [Micro Focus Technical Support website \(http://www.novell.com/support\)](http://www.novell.com/support).

## 9.1 Viewing Error Code Messages

If an error message for a failed NLVM command line operation provides an error code without a corresponding message, you can use the `nss /err` command to view the message. At a command prompt, enter

```
nss /err=<error_code_number>
```

You can also use the following command to view the error code message in the NSS Console (`nsscon`):

```
nsscon /ErrorCode=<error_code_number>
```

Type `exit` and press Enter to close the NSS console and return to the command prompt.

## 9.2 Failure to Create an LVM Volume Group

When you create an LVM volume group or clustered LVM volume group, the command fails with the following error:

```
Error 23384: Not enough free space to handle requested size
```

This error occurs if any one of the devices you used for the volume group is not initialized. Uninitialized devices report that there is no available free space on the device. Initialize the device and try again.

## 9.3 Failure to Create a Clustered LVM Volume Group

When you create a clustered Logical Volume Manager (LVM) volume group, the command fails with the following error:

```
Error 23384: Device /dev/sde is not shared by clvmd
```

This error can occur if the installed Linux kernel does not contain the latest Clustered LVM software. Clustered LVM requires the Linux kernel 2.6.32.45-0.3 or later. You can get the latest kernel version by using the SUSE Linux Enterprise Server (SLES) update channel for OES. For information about applying patches for your server, see “[Updating \(Patching\) an OES Server](#)” in the *OES 2023: Installation Guide*.

The correct version of CLVM software is included in the SLES 11 SP2 or later releases.

## 9.4 Device Is Not Available for Use in an LVM Volume Group

A device cannot be used to create an LVM volume group if any of the following conditions exist:

- ♦ The device is not initialized.
- ♦ The device contains partitions.
- ♦ The device is marked as Shareable for Clustering, which adds a 4 KB partition on the device to store the shared state.

## 9.5 NLVM Pool Move Fails and Deactivates the Pool

If a hardware error is encountered during an `nlvm move`, the pool move fails, and the pool is automatically deactivated. Currently, no error is returned, but the pool will not activate. The pool move cannot continue because of the hardware error. You must delete the move to clear the move:

```
nlvm delete move [<poolname>|<movename>]
```

After the move is deleted, you can activate the pool.

Because of the hardware error, you cannot use the `nlvm move` command to move the pool. You can move the pool's data to another SAN device by restoring files from backup media, or by copying the files from the old pool to a new pool.

## 9.6 Error 20897 - This node is not a cluster member

If OES Cluster Services is installed on a node, but an SBD does not exist, NLVM commands return the following error:

```
Error 20897 - This node is not a cluster member.
```

In a OES Cluster Services cluster, NLVM uses the cluster's SBD to detect whether a node is a cluster member and to lock against concurrent changes to physically shared storage. Without an SBD, NLVM cannot detect whether a node is a member of the cluster and cannot acquire the locks it needs to

execute tasks. In this state, you can use the `-s` option with NLVM commands to prepare a device and create an SBD partition. To minimize the risk of corruption, you must ensure that nobody else is changing any storage on any nodes at the same time.

For information about creating an SBD partition by using NLVM commands, see [Section 8.1, “Creating or Mirroring an SBD Partition,”](#) on page 117.

## 9.7 NLVM Error Codes

Use the information in this section to manage your storage when Novell Linux Volume Manager (NLVM) error conditions exist. NLVM error codes are usually displayed in positive decimal numbers with a message and a `status=` prefix. For example:

```
Pool is not active: status=23357
```

- ◆ [Section 9.7.1, “NLVM Error List,”](#) on page 135
- ◆ [Section 9.7.2, “NLVM Error Descriptions,”](#) on page 138

### 9.7.1 NLVM Error List

NLVM error code numbers can be categorized as follows:

- ◆ NLVM General Errors (23300 to 23309)
  - ◆ [23300 zERR NLVM LOCKED](#)
  - ◆ [23301 zERR NLVM BOOT DEVICE](#)
  - ◆ [23302 zERR NLVM DEVICE HAS RAID](#)
  - ◆ [23303 zERR NLVM NO LOCK](#)
  - ◆ [23304 zERR NLVM VLDB SYMBOL ERROR](#)
  - ◆ [23305 zERR NLVM NOT PERMITTED](#)
  - ◆ [23306 zERR NLVM PARSE ERROR](#)
  - ◆ [23307 zERR NLVM INVALID PARAMETER](#)
- ◆ NLVM Device Errors (23310 to 23319)
  - ◆ [23310 zERR NLVM CSM DEVICE](#)
  - ◆ [23311 zERR NLVM DEVICE NOT FOUND](#)
  - ◆ [23312 zERR NLVM PART NOT FOUND](#)
  - ◆ [23313 zERR NLVM READ FAILURE](#)
  - ◆ [23314 zERR NLVM WRITE FAILURE](#)
  - ◆ [23315 zERR NLVM PART EXPAND FAILURE](#)
  - ◆ [23316 zERR NLVM SIZE TOO SMALL](#)
  - ◆ [23317 zERR NLVM SIZE TOO BIG](#)
  - ◆ [23318 zERR NLVM INVALID PART TYPE](#)
  - ◆ [23319 zERR NLVM DEVICE NOT INIT](#)

- ◆ NLVM General File System Errors (23320 to 23329)
  - ◆ 23320 zERR NLVM ERROR OPENING DB
  - ◆ 23321 zERR NLVM DB MATCH ERROR
  - ◆ 23322 zERR NLVM INVALID MODE
  - ◆ 23323 zERR NLVM ERROR OPENING CONFIG
  - ◆ 23324 zERR NLVM ERROR OPENING DEBUG
  - ◆ 23325 zERR NLVM ERROR OPENING DEV
  - ◆ 23326 zERR NLVM ERROR READING DEV
  - ◆ 23327 zERR NLVM INVALID VERSION
  - ◆ 23328 (reserved)
  - ◆ 23329 (reserved)
- ◆ NLVM Device Mapper Errors (23330 to 23340)
  - ◆ 23330 zERR NLVM ERROR OPENING DM
  - ◆ 23331 zERR NLVM DM IOCTL ERROR
  - ◆ 23332 zERR NLVM BAD SEGMENT COUNT
  - ◆ 23333 (reserved)
  - ◆ 23334 zERR NLVM BAD IDENTIFIER
  - ◆ 23335 zERR NLVM DM OBJECT NOT FOUND
  - ◆ 23336 zERR NLVM INVALID OBJECT
  - ◆ 23337 zERR NLVM OBJECT EXISTS
  - ◆ 23338 zERR NLVM OBJECT BUSY
  - ◆ 23339 zERR NLVM INVALID TYPE
  - ◆ 23340 zERR NLVM LOAD ERROR
- ◆ NLVM Create Snapshot Error (23341)
  - ◆ 23341 zERR NLVM SNAP NOT FOUND
- ◆ NLVM Create Partition Errors (23342 to 23345)
  - ◆ 23342 zERR NLVM LIMIT ERROR
  - ◆ 23343 zERR NLVM PART CREATE
  - ◆ 23344 zERR NLVM PART DELETE
  - ◆ 23345 zERR NLVM PART WRITE
- ◆ NLVM NSS Pool and Volume Errors (23346 to 23359)
  - ◆ 23346 zERR NLVM UNABLE TO EXPAND POOL
  - ◆ 23347 zERR NLVM UNABLE TO CREATE POOL
  - ◆ 23348 zERR NLVM SHARED MISMATCH
  - ◆ 23349 zERR NLVM TYPE MISMATCH
  - ◆ 23350 zERR NLVM HAS POOL
  - ◆ 23351 zERR NLVM DIRECTORY TOO LONG
  - ◆ 23352 zERR NLVM UNABLE TO CREATE DIR



- ◆ 23353 zERR NLVM UNABLE TO CREATE NODE
- ◆ 23354 zERR NLVM POOL UPDATE
- ◆ 23355 zERR NLVM POOL MOUNT ERROR
- ◆ 23356 zERR NLVM POOL MAX SIZE
- ◆ 23358 zERR NLVM GROUP NOT FOUND
- ◆ 23359 (reserved)
- ◆ NLVM NSS Pool Snapshot Errors (23360 to 23369)
  - ◆ 23360 zERR NLVM SNAPSHOT ERROR
  - ◆ 23361 to 23369 (reserved)
- ◆ NLVM NSS Software RAID Errors (23370 to 23379)
  - ◆ 23370 zERR NLVM DUPLICATE DEVICE
  - ◆ 23371 zERR NLVM MAX ELEMENTS
  - ◆ 23372 zERR NLVM TOO FEW ELEMENTS
  - ◆ 23373 zERR NLVM SIZE MISMATCH
  - ◆ 23374 zERR NLVM NOT A RAID
  - ◆ 23375 zERR NLVM NOT A MIRROR
  - ◆ 23376 zERR NLVM TOO MANY PARTITIONS
  - ◆ 23377 zERR NLVM RAID NOT IN SYNC
  - ◆ 23378 zERR NLVM RAID NOT ENABLED
  - ◆ 23379 zERR NLVM RAID NONE IN SYNC
- ◆ NLVM Linux POSIX Volume Errors (23380 to 23390)
  - ◆ 23380 zERR NLVM FSTAB UPDATE
  - ◆ 23381 zERR NLVM OPEN ERROR
  - ◆ 23382 zERR NLVM NO VOLUME NAME
  - ◆ 23383 zERR NLVM NO IP ADDRESS
  - ◆ 23384 zERR NLVM ERROR CREATING LVM VOL
  - ◆ 23385 zERR NLVM ERROR MAKING FS
  - ◆ 23386 zERR NLVM ERROR DELETING RES
  - ◆ 23387 zERR NLVM ERROR DELETING LVM VOL
  - ◆ 23388 zERR NLVM ERROR SENDING CMD
  - ◆ 23389 zERR NLVM NCP ERROR
  - ◆ 23390 zERR NLVM DUPLICATE MP
- ◆ NLVM eDirectory Errors (23391 to 23392)
  - ◆ 23391 zERR NLVM EDIR OBJECT NOT FOUND
  - ◆ 23392 zERR NLVM Invalid CRC
  - ◆ 23393 to 23399 (reserved)

## 9.7.2 NLVM Error Descriptions

- ♦ “NLVM General Errors (23300 to 23309)” on page 138
- ♦ “NLVM Device Errors (23310 to 23319)” on page 138
- ♦ “NLVM General File System Errors (23320 to 23329)” on page 139
- ♦ “NLVM Device Mapper Errors (23330 to 23340)” on page 140
- ♦ “NLVM Create Snapshot Error (23341)” on page 140
- ♦ “NLVM Create Partition Errors (23342 to 23345)” on page 141
- ♦ “NLVM NSS Pool and Volume Errors (23346 to 23369)” on page 141
- ♦ “NLVM NSS Pool Snapshot Errors (23360 to 23369)” on page 142
- ♦ “NLVM NSS Software RAID Errors (23370 to 23379)” on page 142
- ♦ “NLVM Linux POSIX Volume Errors (23380 to 23390)” on page 143
- ♦ “NLVM eDirectory Errors (23391 to 233 99)” on page 143

### NLVM General Errors (23300 to 23309)

#### 23300 zERR NLVM LOCKED

The NLVM lock is already locked.

#### 23301 zERR NLVM BOOT DEVICE

This device contains `/boot`, `root (/)`, or `swap` partitions.

#### 23302 zERR NLVM DEVICE HAS RAID

This device contains RAID partitions.

#### 23303 zERR NLVM NO LOCK

A function was called without the NLVM lock.

#### 23304 zERR NLVM VLDB SYMBOL ERROR

An error occurred when importing the OES Distributed File Services (DFS) VLDB (volume location database) library or functions.

#### 23305 zERR NLVM NOT PERMITTED

This request is not permitted.

#### 23306 zERR NLVM PARSE ERROR

An error occurred when parsing the data.

#### 23307 zERR NLVM INVALID PARAMETER

An invalid parameter was passed in.

### NLVM Device Errors (23310 to 23319)

#### 23310 zERR NLVM CSM DEVICE

This device contains a Cluster Segment Manager (CSM) container.

**23311 zERR NLVM DEVICE NOT FOUND**

The device was not found in NLVM.

**23312 zERR NLVM PART NOT FOUND**

The partition was not found in NLVM.

**23313 zERR NLVM READ FAILURE**

An error occurred while reading a stamp from the disk.

**23314 zERR NLVM WRITE FAILURE**

An error occurred while writing a stamp to the disk.

**23315 zERR NLVM PART EXPAND FAILURE**

An error occurred while expanding the partition.

**23316 zERR NLVM SIZE TOO SMALL**

The specified size is too small.

**23317 zERR NLVM SIZE TOO BIG**

Unable to find a space big enough for the request.

**23318 zERR NLVM INVALID PART TYPE**

The specified partition type is invalid.

**23319 zERR NLVM DEVICE NOT INIT**

The device is not initialized.

**NLVM General File System Errors (23320 to 23329)****23320 zERR NLVM ERROR OPENING DB**

An error occurred while opening the database file.

**23321 zERR NLVM DB MATCH ERROR**

The current object does not match the database object.

**23322 zERR NLVM INVALID MODE**

Invalid mode opening the database file.

**23323 zERR NLVM ERROR OPENING CONFIG**

An error occurred while opening the NLVM configuration file.

**23324 zERR NLVM ERROR OPENING DEBUG**

An error occurred while opening the NLVM debug file.

**23325 zERR NLVM ERROR OPENING DEV**

An error occurred while opening the device for I/O.

**23326 zERR NLVM ERROR READING DEV**

An error occurred while reading from the device.

**23327 zERR NLVM INVALID VERSION**

The stamps have an unsupported version.

**22328 (reserved)**

Not used.

**22329 (reserved)**

Not used.

**NLVM Device Mapper Errors (23330 to 23340)****23330 zERR NLVM ERROR OPENING DM**

An error occurred while opening the Device Mapper.

**23331 zERR NLVM DM IOCTL ERROR**

An error occurred while sending Device Mapper I/O Control (`ioctl`).

**23332 zERR NLVM BAD SEGMENT COUNT**

A segment count mismatch occurred.

**22333 (reserved)**

Not used.

**23334 zERR NLVM BAD IDENTIFIER**

The object identifier does not match a Device Mapper object ID.

**23335 zERR NLVM DM OBJECT NOT FOUND**

The Device Mapper object was not found.

**23336 zERR NLVM INVALID OBJECT**

The object is invalid.

**23337 zERR NLVM OBJECT EXISTS**

The object already exists in Device Mapper.

**23338 zERR NLVM OBJECT BUSY**

The object is busy.

**23339 zERR NLVM INVALID TYPE**

Invalid type parameter.

**23340 zERR NLVM LOAD ERROR**

An error occurred while loading a module.

**NLVM Create Snapshot Error (23341)****23341 zERR NLVM SNAP NOT FOUND**

The NSS pool snapshot was not found.

## **NLVM Create Partition Errors (23342 to 23345)**

### **23342 zERR NLVM LIMIT ERROR**

An error occurred while getting the device limits.

### **23343 zERR NLVM PART CREATE**

An error occurred while creating a partition object.

### **23344 zERR NLVM PART DELETE**

An error occurred while deleting a partition object.

### **23345 zERR NLVM PART WRITE**

An error occurred while writing to a partition object.

## **NLVM NSS Pool and Volume Errors (23346 to 23369)**

### **23346 zERR NLVM UNABLE TO EXPAND POOL**

Unable to expand the NSS pool.

### **23347 zERR NLVM UNABLE TO CREATE POOL**

Unable to create the NSS pool.

### **23348 zERR NLVM SHARED MISMATCH**

The shared states do not match.

### **23349 zERR NLVM TYPE MISMATCH**

The partition types do not match.

### **23350 zERR NLVM HAS POOL**

The partition already has an NSS pool.

### **23351 zERR NLVM DIRECTORY TOO LONG**

The specified directory is too long.

### **23352 zERR NLVM UNABLE TO CREATE DIR**

Unable to create the directory.

### **23353 zERR NLVM UNABLE TO CREATE NODE**

Unable to create the device node.

### **23354 zERR NLVM POOL UPDATE**

An error occurred while updating the NSS pool.

### **23355 zERR NLVM POOL MOUNT ERROR**

An error occurred while updating the NSS pool.

### **23356 zERR NLVM POOL MAX SIZE**

The NSS pool is already at the maximum size.

**23357 zERR NLVM POOL NOT ACTIVE**

The NSS pool is not active.

**23358 zERR NLVM GROUP NOT FOUND**

The group was not found in NLVM.

**23359 (reserved)**

Not used.

**NLVM NSS Pool Snapshot Errors (23360 to 23369)****23360 zERR NLVM SNAPSHOT ERROR**

A pool snapshot error occurred.

**23361 to 23369 (reserved)**

Not used.

**NLVM NSS Software RAID Errors (23370 to 23379)****23370 zERR NLVM DUPLICATE DEVICE**

The device is already used in this RAID.

**23371 zERR NLVM MAX ELEMENTS**

The RAID already has the maximum number of elements.

**23372 zERR NLVM TOO FEW ELEMENTS**

There are too few elements to create the RAID.

**23373 zERR NLVM SIZE MISMATCH**

The element sizes do not match.

**23374 zERR NLVM NOT A RAID**

The device is not a RAID device.

**23375 zERR NLVM NOT A MIRROR**

The device is not a RAID1 device.

**23376 zERR NLVM TOO MANY PARTITIONS**

You are trying to add too many partitions to a RAID.

**23377 zERR NLVM RAID NOT IN SYNC**

The RAID is not in sync.

**23378 zERR NLVM RAID NOT ENABLED**

The RAID is not enabled.

**23379 zERR NLVM RAID NONE IN SYNC**

No partition of the RAID device is in sync.

## **NLVM Linux POSIX Volume Errors (23380 to 23390)**

### **23380 zERR NLVM FSTAB UPDATE**

An error occurred while updating the `/etc/fstab` file.

### **23381 zERR NLVM OPEN ERROR**

An error occurred while opening the file.

### **23382 zERR NLVM NO VOLUME NAME**

No volume name was specified.

### **23383 zERR NLVM NO IP ADDRESS**

No IP address was specified.

### **23384 zERR NLVM ERROR CREATING LVM VOL**

An error occurred while creating the LVM2 volume.

### **23385 zERR NLVM ERROR MAKING FS**

An error occurred while making the file system on a volume.

### **23386 zERR NLVM ERROR DELETING RES**

An error occurred while deleting a cluster resource for a volume.

### **23387 zERR NLVM ERROR DELETING LVM VOL**

An error occurred while deleting the LVM2 volume.

### **23388 zERR NLVM ERROR SENDING CMD**

An error occurred while sending the XML command.

### **23389 zERR NLVM NCP ERROR**

An error occurred while adding a volume to NCP (NetWare Core Protocol).

### **23390 zERR NLVM DUPLICATE MP**

A duplicate mount point was specified.

## **NLVM eDirectory Errors (23391 to 233 99)**

### **23391 zERR NLVM EDIR OBJECT NOT FOUND**

The eDirectory object was not found.

### **23392 zERR NLVM Invalid CRC**

Invalid CRC (cyclic redundancy check) in GPT (GUID partition table) partitions.

### **23393 to 23399 (reserved)**

Not used.

## 9.8 NSS Error Codes

For information about OES Storage Services error codes, see the *OES Storage Services Error Codes* (<http://www.novell.com/documentation/nwec/nwec/data/al3s3ui.html>).



# 10 Security Considerations

This section describes the security considerations for the Novell Linux Volume Manager (NLVM) on a Open Enterprise Server (OES) server.

- ♦ [Section 10.1, “Root User Privileges,” on page 145](#)
- ♦ [Section 10.2, “Files,” on page 145](#)

## 10.1 Root User Privileges

The Linux system `root` user privileges are required to use NLVM commands.

## 10.2 Files

### **`/dev/nss/`**

Location where NSS software RAID and SBD partition device mapper objects are created.

### **`/dev/pool/`**

Location where NSS pool device mapper objects are created.

### **`/etc/opt/novell/nss/nlvm.conf`**

Location of the NLVM configuration file.

### **`/opt/novell/nss/mnt/.pools/`**

Location where NSS pool objects are mounted.

### **`/opt/novell/nss/nlvm/`**

Location of the NLVM storage configuration database files. The database files are named `nlvm.<number.>db`, such as `nlvm.db`, `nlvm.1.db`, and so on. The default is to keep the 10 most recent files. The number of NLVM database files to keep is set in the `/etc/opt/novell/nss/nlvm.conf` file.

### **`/opt/novell/nss/sbin/nlvm`**

Location of the NLVM utility. It also has a link in the `sbin` directory so that it is in the search path.

### **`/var/opt/novell/log/nss/debug/`**

Location of the debug log files when debug is enabled. The debug files are named `nlvm_debug.<number.>log`, such as `nlvm_debug.log`, `nlvm_debug.1.log`, and so on. The default is to keep the 10 most recent files. The number of debug log files to keep is set in the `/etc/opt/novell/nss/nlvm.conf` file.

### **`/var/run/novell-nss/nlvm.lock`**

Local lock file for NLVM.



# A

## Configuring Settings for the NLVM Library

The Novell Linux Volume Manager (NLVM) library software has some configurable settings that are exposed in the `/etc/opt/novell/nss/nlvm.conf` file. The default settings are automatically configured. To modify the default behavior, use the options described in [Table A-1](#).

**Table A-1** Default Settings for the NLVM Library

| Parameter                   | Description  |
|-----------------------------|--|
| <code>Debug on</code>       | <p>If this line is enabled, the command allows the debug feature of the NLVM utility to run every time without needing to use the <code>-d</code> option.</p> <p>The default is off (commented out). You can enable debug as needed by using the <code>-d</code> option when you start the utility.</p> <p>To enable debug to run every time, you can uncomment the <code>Debug on</code> command in the <code>nlvm.conf</code> file.</p> <p>To return to the default debug behavior with the <code>-d</code> option, you can comment out the <code>Debug on</code> command again.</p>   |
| <code>Debug files 10</code> | <p>If this line is enabled, the command specifies the number of NLVM debug log files to keep before deleting the oldest file. A log file shows actions that were performed by the NLVM library.</p> <p>The default is to keep the 10 most recent files. The minimum value is 1. The default setting applies when the command is commented out.</p> <p>To modify the number of files kept, uncomment the line and specify a new value.</p> <p>To use the default setting, comment out the command again.</p> <p>When debug runs, a debug log file is opened in the <code>/var/opt/novell/log/nss/debug</code> directory. The debug files are named <code>nlvm_debug.&lt;number.&gt;log</code>, such as <code>nlvm_debug.log</code>, <code>nlvm_debug.1.log</code>, and so on.</p> |

| Parameter                             | Description   |
|---------------------------------------|---|
| <code>Data base files 10</code>       | <p>If this line is enabled, the command specifies the number of NLVM database files to keep before deleting the oldest file. Database files are stored every time a change is made to the system with the NLVM library.</p> <p>The default is to keep the 10 most recent files. The minimum value is 1. The default setting applies when the command is commented out.</p> <p>To modify the number of files kept, uncomment the line and specify a new value.</p> <p>To use the default setting, comment out the command.</p> <p>When a change is made to the system, a database file is opened in the <code>/opt/novell/nss/nlvm/</code> directory. The database files are named <code>nlvm.&lt;number.&gt;db</code>, such as <code>nlvm.db</code>, <code>nlvm.1.db</code>, and so on.</p> |
| <code>Auto refresh off</code>         | <p>If this line is enabled, the command turns off the autorefresh. The system gets its information from the database files. This results in much faster load times for utilities, but might require a refresh within the utility. If the autorefresh is off, a refresh can be triggered by using the <code>-r</code> option when you start the NLVM utility.</p> <p><b>IMPORTANT:</b> If OES Cluster Services is on, the autorefresh is always on.</p> <p>The default is that autorefresh is enabled (the line is commented out). This allows the NLVM library to refresh the system each time it is used. The autorefresh picks up any changes to the system that happened outside the library.</p>  |
| <code>Blk_discard_on_delete on</code> | <p>If this parameter is enabled, the command discards the blocks on a thin-provisioned storage disk when you delete a pool on that disk. This allows all the allocated blocks for that pool to be made available to the storage disk.</p> <p>By default, this parameter is off (the line is commented out).</p>   |