



Red Hat Enterprise Linux 7

Logical Volume Manager Administration

LVM Administrator Guide

Red Hat Enterprise Linux 7 Logical Volume Manager Administration

LVM Administrator Guide

Steven Levine

Red Hat Customer Content Services

slevine@redhat.com

Legal Notice

Copyright © 2018 Red Hat, Inc.

This document is licensed by Red Hat under the [Creative Commons Attribution-ShareAlike 3.0 Unported License](https://creativecommons.org/licenses/by-sa/3.0/). If you distribute this document, or a modified version of it, you must provide attribution to Red Hat, Inc. and provide a link to the original. If the document is modified, all Red Hat trademarks must be removed.

Red Hat, as the licensor of this document, waives the right to enforce, and agrees not to assert, Section 4d of CC-BY-SA to the fullest extent permitted by applicable law.

Red Hat, Red Hat Enterprise Linux, the Shadowman logo, JBoss, OpenShift, Fedora, the Infinity logo, and RHCE are trademarks of Red Hat, Inc., registered in the United States and other countries.

Linux ® is the registered trademark of Linus Torvalds in the United States and other countries.

Java ® is a registered trademark of Oracle and/or its affiliates.

XFS ® is a trademark of Silicon Graphics International Corp. or its subsidiaries in the United States and/or other countries.

MySQL ® is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Node.js ® is an official trademark of Joyent. Red Hat Software Collections is not formally related to or endorsed by the official Joyent Node.js open source or commercial project.

The OpenStack ® Word Mark and OpenStack logo are either registered trademarks/service marks or trademarks/service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation, or the OpenStack community.

All other trademarks are the property of their respective owners.

Abstract

This book describes the LVM logical volume manager, including information on running LVM in a clustered environment.

Table of Contents

CHAPTER 1. THE LVM LOGICAL VOLUME MANAGER	4
1.1. NEW AND CHANGED FEATURES	4
1.2. LOGICAL VOLUMES	5
1.3. LVM ARCHITECTURE OVERVIEW	6
1.4. LVM LOGICAL VOLUMES IN A RED HAT HIGH AVAILABILITY CLUSTER	7
CHAPTER 2. LVM COMPONENTS	9
2.1. PHYSICAL VOLUMES	9
2.2. VOLUME GROUPS	10
2.3. LVM LOGICAL VOLUMES	11
CHAPTER 3. LVM ADMINISTRATION OVERVIEW	18
3.1. LOGICAL VOLUME CREATION OVERVIEW	18
3.2. GROWING A FILE SYSTEM ON A LOGICAL VOLUME	18
3.3. LOGICAL VOLUME BACKUP	19
3.4. LOGGING	19
3.5. THE METADATA DAEMON (LVMETAD)	19
3.6. DISPLAYING LVM INFORMATION WITH THE LVM COMMAND	20
CHAPTER 4. LVM ADMINISTRATION WITH CLI COMMANDS	22
4.1. USING CLI COMMANDS	22
4.2. PHYSICAL VOLUME ADMINISTRATION	23
4.3. VOLUME GROUP ADMINISTRATION	26
4.4. LOGICAL VOLUME ADMINISTRATION	34
4.5. CONTROLLING LVM DEVICE SCANS WITH FILTERS	81
4.6. ONLINE DATA RELOCATION	82
4.7. ACTIVATING LOGICAL VOLUMES ON INDIVIDUAL NODES IN A CLUSTER	83
4.8. CUSTOMIZED REPORTING FOR LVM	83
CHAPTER 5. LVM CONFIGURATION EXAMPLES	98
5.1. CREATING AN LVM LOGICAL VOLUME ON THREE DISKS	98
5.2. CREATING A STRIPED LOGICAL VOLUME	99
5.3. SPLITTING A VOLUME GROUP	100
5.4. REMOVING A DISK FROM A LOGICAL VOLUME	102
5.5. CREATING A MIRRORED LVM LOGICAL VOLUME IN A CLUSTER	104
CHAPTER 6. LVM TROUBLESHOOTING	108
6.1. TROUBLESHOOTING DIAGNOSTICS	108
6.2. DISPLAYING INFORMATION ON FAILED DEVICES	108
6.3. RECOVERING FROM LVM MIRROR FAILURE	109
6.4. RECOVERING PHYSICAL VOLUME METADATA	112
6.5. REPLACING A MISSING PHYSICAL VOLUME	114
6.6. REMOVING LOST PHYSICAL VOLUMES FROM A VOLUME GROUP	114
6.7. INSUFFICIENT FREE EXTENTS FOR A LOGICAL VOLUME	114
6.8. DUPLICATE PV WARNINGS FOR MULTIPATHED DEVICES	115
APPENDIX A. THE DEVICE MAPPER	118
A.1. DEVICE TABLE MAPPINGS	118
A.2. THE DMSETUP COMMAND	128
A.3. DEVICE MAPPER SUPPORT FOR THE UDEV DEVICE MANAGER	132
APPENDIX B. THE LVM CONFIGURATION FILES	136
B.1. THE LVM CONFIGURATION FILES	136

B.2. THE LVMCONFIG COMMAND	136
B.3. LVM PROFILES	137
B.4. SAMPLE LVM.CONF FILE	138
APPENDIX C. LVM SELECTION CRITERIA	177
C.1. SELECTION CRITERIA FIELD TYPES	177
C.2. SELECTION CRITERIA OPERATORS	178
C.3. SELECTION CRITERIA FIELDS	180
C.4. SPECIFYING TIME VALUES	192
C.5. SELECTION CRITERIA DISPLAY EXAMPLES	194
C.6. SELECTION CRITERIA PROCESSING EXAMPLES	196
APPENDIX D. LVM OBJECT TAGS	199
D.1. ADDING AND REMOVING OBJECT TAGS	199
D.2. HOST TAGS	199
D.3. CONTROLLING ACTIVATION WITH TAGS	200
APPENDIX E. LVM VOLUME GROUP METADATA	201
E.1. THE PHYSICAL VOLUME LABEL	201
E.2. METADATA CONTENTS	201
E.3. SAMPLE METADATA	202
APPENDIX F. REVISION HISTORY	205
INDEX	206

CHAPTER 1. THE LVM LOGICAL VOLUME MANAGER

This chapter provides a summary of the features of the LVM logical volume manager that are new since the initial release of Red Hat Enterprise Linux 7. This chapter also provides a high-level overview of the components of the Logical Volume Manager (LVM).

1.1. NEW AND CHANGED FEATURES

This section lists features of the LVM logical volume manager that are new since the initial release of Red Hat Enterprise Linux 7.

1.1.1. New and Changed Features for Red Hat Enterprise Linux 7.1

Red Hat Enterprise Linux 7.1 includes the following documentation and feature updates and changes.

- The documentation for thinly-provisioned volumes and thinly-provisioned snapshots has been clarified. Additional information about LVM thin provisioning is now provided in the **lvmthin(7)** man page. For general information on thinly-provisioned logical volumes, see [Section 2.3.4, “Thinly-Provisioned Logical Volumes \(Thin Volumes\)”](#). For information on thinly-provisioned snapshot volumes, see [Section 2.3.6, “Thinly-Provisioned Snapshot Volumes”](#).
- This manual now documents the **lvm dumpconfig** command in [Section B.2, “The lvmconfig Command”](#). Note that as of the Red Hat Enterprise Linux 7.2 release, this command was renamed **lvmconfig**, although the old format continues to work.
- This manual now documents LVM profiles in [Section B.3, “LVM Profiles”](#).
- This manual now documents the **lvm** command in [Section 3.6, “Displaying LVM Information with the lvm Command”](#).
- In the Red Hat Enterprise Linux 7.1 release, you can control activation of thin pool snapshots with the **-k** and **-K** options of the **lvcreate** and **lvchange** command, as documented in [Section 4.4.20, “Controlling Logical Volume Activation”](#).
- This manual documents the **--force** argument of the **vgimport** command. This allows you to import volume groups that are missing physical volumes and subsequently run the **vgreduce - -removemissing** command. For information on the **vgimport** command, refer to [Section 4.3.15, “Moving a Volume Group to Another System”](#).
- This manual documents the **--mirroronly** argument of the **vgreduce** command. This allows you remove only the logical volumes that are mirror images from a physical volume that has failed. For information on using this option, refer to [Section 4.3.15, “Moving a Volume Group to Another System”](#).

In addition, small technical corrections and clarifications have been made throughout the document.

1.1.2. New and Changed Features for Red Hat Enterprise Linux 7.2

Red Hat Enterprise Linux 7.2 includes the following documentation and feature updates and changes.

- Many LVM processing commands now accept the **-S** or **--select** option to define selection criteria for those commands. LVM selection criteria are documented in the new appendix [Appendix C, LVM Selection Criteria](#).

- This document provides basic procedures for creating cache logical volumes in [Section 4.4.8, “Creating LVM Cache Logical Volumes”](#).
- The troubleshooting chapter of this document includes a new section, [Section 6.8, “Duplicate PV Warnings for Multipathed Devices”](#).
- As of the Red Hat Enterprise Linux 7.2 release, the **lvm dumpconfig** command was renamed **lvmconfig**, although the old format continues to work. This change is reflected throughout this document.

In addition, small technical corrections and clarifications have been made throughout the document.

1.1.3. New and Changed Features for Red Hat Enterprise Linux 7.3

Red Hat Enterprise Linux 7.3 includes the following documentation and feature updates and changes.

- LVM supports RAID0 segment types. RAID0 spreads logical volume data across multiple data subvolumes in units of stripe size. For information on creating RAID0 volumes, see [Section 4.4.3.1, “Creating RAID0 Volumes \(Red Hat Enterprise Linux 7.3 and Later\)”](#).
- You can report information about physical volumes, volume groups, logical volumes, physical volume segments, and logical volume segments all at once with the **lvm fullreport** command. For information on this command and its capabilities, see the **lvm-fullreport(8)** man page.
- LVM supports log reports, which contain a log of operations, messages, and per-object status with complete object identification collected during LVM command execution. For an example of an LVM log report, see [Section 4.8.6, “Command Log Reporting \(Red Hat Enterprise Linux 7.3 and later\)”](#). For further information about the LVM log report, see the **lvmreport(7)** man page.
- You can use the **--reportformat** option of the LVM display commands to display the output in JSON format. For an example of output displayed in JSON format, see [Section 4.8.5, “JSON Format Output \(Red Hat Enterprise Linux 7.3 and later\)”](#).
- You can now configure your system to track thin snapshot and thin logical volumes that have been removed by enabling the **record_lvs_history** metadata option in the **lvm.conf** configuration file. This allows you to display a full thin snapshot dependency chain that includes logical volumes that have been removed from the original dependency chain and have become *historical* logical volumes. For information on historical logical volumes, see [Section 4.4.21, “Tracking and Displaying Historical Logical Volumes \(Red Hat Enterprise Linux 7.3 and Later\)”](#).

In addition, small technical corrections and clarifications have been made throughout the document.

1.1.4. New and Changed Features for Red Hat Enterprise Linux 7.4

Red Hat Enterprise Linux 7.4 includes the following documentation and feature updates and changes.

- Red Hat Enterprise Linux 7.4 provides support for RAID takeover and RAID reshaping. For a summary of these features, see [Section 4.4.3.12, “RAID Takeover \(Red Hat Enterprise Linux 7.4 and Later\)”](#) and [Section 4.4.3.13, “Reshaping a RAID Logical Volume \(Red Hat Enterprise Linux 7.4 and Later\)”](#).

1.2. LOGICAL VOLUMES

Volume management creates a layer of abstraction over physical storage, allowing you to create logical

storage volumes. This provides much greater flexibility in a number of ways than using physical storage directly. With a logical volume, you are not restricted to physical disk sizes. In addition, the hardware storage configuration is hidden from the software so it can be resized and moved without stopping applications or unmounting file systems. This can reduce operational costs.

Logical volumes provide the following advantages over using physical storage directly:

- Flexible capacity

When using logical volumes, file systems can extend across multiple disks, since you can aggregate disks and partitions into a single logical volume.

- Resizeable storage pools

You can extend logical volumes or reduce logical volumes in size with simple software commands, without reformatting and repartitioning the underlying disk devices.

- Online data relocation

To deploy newer, faster, or more resilient storage subsystems, you can move data while your system is active. Data can be rearranged on disks while the disks are in use. For example, you can empty a hot-swappable disk before removing it.

- Convenient device naming

Logical storage volumes can be managed in user-defined and custom named groups.

- Disk striping

You can create a logical volume that stripes data across two or more disks. This can dramatically increase throughput.

- Mirroring volumes

Logical volumes provide a convenient way to configure a mirror for your data.

- Volume Snapshots

Using logical volumes, you can take device snapshots for consistent backups or to test the effect of changes without affecting the real data.

The implementation of these features in LVM is described in the remainder of this document.

1.3. LVM ARCHITECTURE OVERVIEW



NOTE

LVM2 is backwards compatible with LVM1, with the exception of snapshot and cluster support. You can convert a volume group from LVM1 format to LVM2 format with the **vgconvert** command. For information on converting LVM metadata format, see the **vgconvert(8)** man page.

The underlying physical storage unit of an LVM logical volume is a block device such as a partition or whole disk. This device is initialized as an LVM *physical volume* (PV).

To create an LVM logical volume, the physical volumes are combined into a *volume group* (VG). This creates a pool of disk space out of which LVM logical volumes (LVs) can be allocated. This process is analogous to the way in which disks are divided into partitions. A logical volume is used by file systems and applications (such as databases).

Figure 1.1, “LVM Logical Volume Components” shows the components of a simple LVM logical volume:




Figure 1.1. LVM Logical Volume Components

For detailed information on the components of an LVM logical volume, see [Chapter 2, LVM Components](#).

1.4. LVM LOGICAL VOLUMES IN A RED HAT HIGH AVAILABILITY CLUSTER

The Red Hat High Availability Add-On provides support for LVM volumes in two distinct cluster configurations:

- High availability LVM volumes (HA-LVM) in an active/passive failover configurations in which only a single node of the cluster accesses the storage at any one time.
- LVM volumes that use the Clustered Logical Volume (CLVM) extensions in an active/active configurations in which more than one node of the cluster requires access to the storage at the same time. CLVM is part of the Resilient Storage Add-On.

1.4.1. Choosing CLVM or HA-LVM

When to use CLVM or HA-LVM should be based on the needs of the applications or services being deployed.

- If multiple nodes of the cluster require simultaneous read/write access to LVM volumes in an active/active system, then you must use CLVMD. CLVMD provides a system for coordinating activation of and changes to LVM volumes across nodes of a cluster concurrently. CLVMD's clustered-locking service provides protection to LVM metadata as various nodes of the cluster

interact with volumes and make changes to their layout. This protection is contingent upon appropriately configuring the volume groups in question, including setting **locking_type** to 3 in the **lvm.conf** file and setting the clustered flag on any volume group that will be managed by CLVMD and activated simultaneously across multiple cluster nodes.

- If the high availability cluster is configured to manage shared resources in an active/passive manner with only one single member needing access to a given LVM volume at a time, then you can use HA-LVM without the CLVMD clustered-locking service

Most applications will run better in an active/passive configuration, as they are not designed or optimized to run concurrently with other instances. Choosing to run an application that is not cluster-aware on clustered logical volumes may result in degraded performance if the logical volume is mirrored. This is because there is cluster communication overhead for the logical volumes themselves in these instances. A cluster-aware application must be able to achieve performance gains above the performance losses introduced by cluster file systems and cluster-aware logical volumes. This is achievable for some applications and workloads more easily than others. Determining what the requirements of the cluster are and whether the extra effort toward optimizing for an active/active cluster will pay dividends is the way to choose between the two LVM variants. Most users will achieve the best HA results from using HA-LVM.

HA-LVM and CLVM are similar in the fact that they prevent corruption of LVM metadata and its logical volumes, which could otherwise occur if multiple machines are allowed to make overlapping changes. HA-LVM imposes the restriction that a logical volume can only be activated exclusively; that is, active on only one machine at a time. This means that only local (non-clustered) implementations of the storage drivers are used. Avoiding the cluster coordination overhead in this way increases performance. CLVM does not impose these restrictions and a user is free to activate a logical volume on all machines in a cluster; this forces the use of cluster-aware storage drivers, which allow for cluster-aware file systems and applications to be put on top.

1.4.2. Configuring LVM volumes in a cluster

In Red Hat Enterprise Linux 7, clusters are managed through Pacemaker. Both HA-LVM and CLVM logical volumes are supported only in conjunction with Pacemaker clusters, and must be configured as cluster resources.

- For a procedure for configuring an HA-LVM volume as part of a Pacemaker cluster, see [An active/passive Apache HTTP Server in a Red Hat High Availability Cluster in High Availability Add-On Administration](#). Note that this procedure includes the following steps:
 - Configuring an LVM logical volume
 - Ensuring that only the cluster is capable of activating the volume group
 - Configuring the LVM volume as a cluster resource
- For a procedure for configuring a CLVM volume in a cluster, see [Configuring a GFS2 File System in a Cluster in Global File System 2](#).

CHAPTER 2. LVM COMPONENTS

This chapter describes the components of an LVM Logical volume.

2.1. PHYSICAL VOLUMES

The underlying physical storage unit of an LVM logical volume is a block device such as a partition or whole disk. To use the device for an LVM logical volume, the device must be initialized as a physical volume (PV). Initializing a block device as a physical volume places a label near the start of the device.

By default, the LVM label is placed in the second 512-byte sector. You can overwrite this default by placing the label on any of the first 4 sectors when you create the physical volume. This allows LVM volumes to co-exist with other users of these sectors, if necessary.

An LVM label provides correct identification and device ordering for a physical device, since devices can come up in any order when the system is booted. An LVM label remains persistent across reboots and throughout a cluster.

The LVM label identifies the device as an LVM physical volume. It contains a random unique identifier (the UUID) for the physical volume. It also stores the size of the block device in bytes, and it records where the LVM metadata will be stored on the device.

The LVM metadata contains the configuration details of the LVM volume groups on your system. By default, an identical copy of the metadata is maintained in every metadata area in every physical volume within the volume group. LVM metadata is small and stored as ASCII.

Currently LVM allows you to store 0, 1 or 2 identical copies of its metadata on each physical volume. The default is 1 copy. Once you configure the number of metadata copies on the physical volume, you cannot change that number at a later time. The first copy is stored at the start of the device, shortly after the label. If there is a second copy, it is placed at the end of the device. If you accidentally overwrite the area at the beginning of your disk by writing to a different disk than you intend, a second copy of the metadata at the end of the device will allow you to recover the metadata.

For detailed information about the LVM metadata and changing the metadata parameters, see [Appendix E, *LVM Volume Group Metadata*](#).

2.1.1. LVM Physical Volume Layout

[Figure 2.1, “Physical Volume layout”](#) shows the layout of an LVM physical volume. The LVM label is on the second sector, followed by the metadata area, followed by the usable space on the device.



NOTE

In the Linux kernel (and throughout this document), sectors are considered to be 512 bytes in size.




Figure 2.1. Physical Volume layout

2.1.2. Multiple Partitions on a Disk

LVM allows you to create physical volumes out of disk partitions. Red Hat recommends that you create a single partition that covers the whole disk to label as an LVM physical volume for the following reasons:

- Administrative convenience

It is easier to keep track of the hardware in a system if each real disk only appears once. This becomes particularly true if a disk fails. In addition, multiple physical volumes on a single disk may cause a kernel warning about unknown partition types at boot.

- Striping performance

LVM cannot tell that two physical volumes are on the same physical disk. If you create a striped logical volume when two physical volumes are on the same physical disk, the stripes could be on different partitions on the same disk. This would result in a decrease in performance rather than an increase.

Although it is not recommended, there may be specific circumstances when you will need to divide a disk into separate LVM physical volumes. For example, on a system with few disks it may be necessary to move data around partitions when you are migrating an existing system to LVM volumes. Additionally, if you have a very large disk and want to have more than one volume group for administrative purposes then it is necessary to partition the disk. If you do have a disk with more than one partition and both of those partitions are in the same volume group, take care to specify which partitions are to be included in a logical volume when creating striped volumes.

2.2. VOLUME GROUPS

Physical volumes are combined into volume groups (VGs). This creates a pool of disk space out of which logical volumes can be allocated.

Within a volume group, the disk space available for allocation is divided into units of a fixed-size called extents. An extent is the smallest unit of space that can be allocated. Within a physical volume, extents are referred to as physical extents.

A logical volume is allocated into logical extents of the same size as the physical extents. The extent size is thus the same for all logical volumes in the volume group. The volume group maps the logical extents to physical extents.

2.3. LVM LOGICAL VOLUMES

In LVM, a volume group is divided up into logical volumes. The following sections describe the different types of logical volumes.

2.3.1. Linear Volumes

A linear volume aggregates space from one or more physical volumes into one logical volume. For example, if you have two 60GB disks, you can create a 120GB logical volume. The physical storage is concatenated.

Creating a linear volume assigns a range of physical extents to an area of a logical volume in order. For example, as shown in [Figure 2.2, “Extent Mapping”](#) logical extents 1 to 99 could map to one physical volume and logical extents 100 to 198 could map to a second physical volume. From the point of view of the application, there is one device that is 198 extents in size.




Figure 2.2. Extent Mapping

The physical volumes that make up a logical volume do not have to be the same size. [Figure 2.3, “Linear Volume with Unequal Physical Volumes”](#) shows volume group **VG1** with a physical extent size of 4MB. This volume group includes 2 physical volumes named **PV1** and **PV2**. The physical volumes are divided

into 4MB units, since that is the extent size. In this example, **PV1** is 200 extents in size (800MB) and **PV2** is 100 extents in size (400MB). You can create a linear volume any size between 1 and 300 extents (4MB to 1200MB). In this example, the linear volume named **LV1** is 300 extents in size.




Figure 2.3. Linear Volume with Unequal Physical Volumes

You can configure more than one linear logical volume of whatever size you require from the pool of physical extents. [Figure 2.4, “Multiple Logical Volumes”](#) shows the same volume group as in [Figure 2.3, “Linear Volume with Unequal Physical Volumes”](#), but in this case two logical volumes have been carved out of the volume group: **LV1**, which is 250 extents in size (1000MB) and **LV2** which is 50 extents in size (200MB).




Figure 2.4. Multiple Logical Volumes

2.3.2. Striped Logical Volumes

When you write data to an LVM logical volume, the file system lays the data out across the underlying

physical volumes. You can control the way the data is written to the physical volumes by creating a striped logical volume. For large sequential reads and writes, this can improve the efficiency of the data I/O.

Striping enhances performance by writing data to a predetermined number of physical volumes in round-robin fashion. With striping, I/O can be done in parallel. In some situations, this can result in near-linear performance gain for each additional physical volume in the stripe.

The following illustration shows data being striped across three physical volumes. In this figure:

- the first stripe of data is written to the first physical volume
- the second stripe of data is written to the second physical volume
- the third stripe of data is written to the third physical volume
- the fourth stripe of data is written to the first physical volume

In a striped logical volume, the size of the stripe cannot exceed the size of an extent.




Figure 2.5. Striping Data Across Three PVs

Striped logical volumes can be extended by concatenating another set of devices onto the end of the first set. In order to extend a striped logical volume, however, there must be enough free space on the set of underlying physical volumes that make up the volume group to support the stripe. For example, if you have a two-way stripe that uses up an entire volume group, adding a single physical volume to the

volume group will not enable you to extend the stripe. Instead, you must add at least two physical volumes to the volume group. For more information on extending a striped volume, see [Section 4.4.17, “Extending a Striped Volume”](#).

2.3.3. RAID Logical Volumes

LVM supports RAID0/1/4/5/6/10. An LVM RAID volume has the following characteristics:

- RAID logical volumes created and managed by means of LVM leverage the MD kernel drivers.
- RAID1 images can be temporarily split from the array and merged back into the array later.
- LVM RAID volumes support snapshots.

For information on creating RAID logical volumes, see [Section 4.4.3, “RAID Logical Volumes”](#).



NOTE

RAID logical volumes are not cluster-aware. While RAID logical volumes can be created and activated exclusively on one machine, they cannot be activated simultaneously on more than one machine. If you require non-exclusive mirrored volumes, you must create the volumes with a **mirror** segment type, as described in [Section 4.4.4, “Creating Mirrored Volumes”](#).

2.3.4. Thinly-Provisioned Logical Volumes (Thin Volumes)

Logical volumes can be thinly provisioned. This allows you to create logical volumes that are larger than the available extents. Using thin provisioning, you can manage a storage pool of free space, known as a thin pool, which can be allocated to an arbitrary number of devices when needed by applications. You can then create devices that can be bound to the thin pool for later allocation when an application actually writes to the logical volume. The thin pool can be expanded dynamically when needed for cost-effective allocation of storage space.



NOTE

Thin volumes are not supported across the nodes in a cluster. The thin pool and all its thin volumes must be exclusively activated on only one cluster node.

By using thin provisioning, a storage administrator can overcommit the physical storage, often avoiding the need to purchase additional storage. For example, if ten users each request a 100GB file system for their application, the storage administrator can create what appears to be a 100GB file system for each user but which is backed by less actual storage that is used only when needed. When using thin provisioning, it is important that the storage administrator monitor the storage pool and add more capacity if it starts to become full.

To make sure that all available space can be used, LVM supports data discard. This allows for re-use of the space that was formerly used by a discarded file or other block range.

For information on creating thin volumes, see [Section 4.4.5, “Creating Thinly-Provisioned Logical Volumes”](#).

Thin volumes provide support for a new implementation of copy-on-write (COW) snapshot logical volumes, which allow many virtual devices to share the same data in the thin pool. For information on thin snapshot volumes, see [Section 2.3.6, “Thinly-Provisioned Snapshot Volumes”](#).

2.3.5. Snapshot Volumes

The LVM snapshot feature provides the ability to create virtual images of a device at a particular instant without causing a service interruption. When a change is made to the original device (the origin) after a snapshot is taken, the snapshot feature makes a copy of the changed data area as it was prior to the change so that it can reconstruct the state of the device.



NOTE

LVM supports thinly-provisioned snapshots. For information on thinly provisioned snapshot volumes, see [Section 2.3.6, “Thinly-Provisioned Snapshot Volumes”](#).



NOTE

LVM snapshots are not supported across the nodes in a cluster. You cannot create a snapshot volume in a clustered volume group.

Because a snapshot copies only the data areas that change after the snapshot is created, the snapshot feature requires a minimal amount of storage. For example, with a rarely updated origin, 3-5 % of the origin's capacity is sufficient to maintain the snapshot.



NOTE

Snapshot copies of a file system are virtual copies, not an actual media backup for a file system. Snapshots do not provide a substitute for a backup procedure.

The size of the snapshot governs the amount of space set aside for storing the changes to the origin volume. For example, if you made a snapshot and then completely overwrote the origin the snapshot would have to be at least as big as the origin volume to hold the changes. You need to dimension a snapshot according to the expected level of change. So for example a short-lived snapshot of a read-mostly volume, such as `/usr`, would need less space than a long-lived snapshot of a volume that sees a greater number of writes, such as `/home`.

If a snapshot runs full, the snapshot becomes invalid, since it can no longer track changes on the origin volume. You should regularly monitor the size of the snapshot. Snapshots are fully resizable, however, so if you have the storage capacity you can increase the size of the snapshot volume to prevent it from getting dropped. Conversely, if you find that the snapshot volume is larger than you need, you can reduce the size of the volume to free up space that is needed by other logical volumes.

When you create a snapshot file system, full read and write access to the origin stays possible. If a chunk on a snapshot is changed, that chunk is marked and never gets copied from the original volume.

There are several uses for the snapshot feature:

- Most typically, a snapshot is taken when you need to perform a backup on a logical volume without halting the live system that is continuously updating the data.
- You can execute the **fsck** command on a snapshot file system to check the file system integrity and determine whether the original file system requires file system repair.
- Because the snapshot is read/write, you can test applications against production data by taking a snapshot and running tests against the snapshot, leaving the real data untouched.
- You can create LVM volumes for use with Red Hat Virtualization. LVM snapshots can be used

to create snapshots of virtual guest images. These snapshots can provide a convenient way to modify existing guests or create new guests with minimal additional storage. For information on creating LVM-based storage pools with Red Hat Virtualization, see the *Virtualization Administration Guide*.

For information on creating snapshot volumes, see [Section 4.4.6, “Creating Snapshot Volumes”](#).

You can use the `-merge` option of the `lvconvert` command to merge a snapshot into its origin volume. One use for this feature is to perform system rollback if you have lost data or files or otherwise need to restore your system to a previous state. After you merge the snapshot volume, the resulting logical volume will have the origin volume's name, minor number, and UUID and the merged snapshot is removed. For information on using this option, see [Section 4.4.9, “Merging Snapshot Volumes”](#).

2.3.6. Thinly-Provisioned Snapshot Volumes

Red Hat Enterprise Linux provides support for thinly-provisioned snapshot volumes. Thin snapshot volumes allow many virtual devices to be stored on the same data volume. This simplifies administration and allows for the sharing of data between snapshot volumes.

As for all LVM snapshot volumes, as well as all thin volumes, thin snapshot volumes are not supported across the nodes in a cluster. The snapshot volume must be exclusively activated on only one cluster node.

Thin snapshot volumes provide the following benefits:

- A thin snapshot volume can reduce disk usage when there are multiple snapshots of the same origin volume.
- If there are multiple snapshots of the same origin, then a write to the origin will cause one COW operation to preserve the data. Increasing the number of snapshots of the origin should yield no major slowdown.
- Thin snapshot volumes can be used as a logical volume origin for another snapshot. This allows for an arbitrary depth of recursive snapshots (snapshots of snapshots of snapshots...).
- A snapshot of a thin logical volume also creates a thin logical volume. This consumes no data space until a COW operation is required, or until the snapshot itself is written.
- A thin snapshot volume does not need to be activated with its origin, so a user may have only the origin active while there are many inactive snapshot volumes of the origin.
- When you delete the origin of a thinly-provisioned snapshot volume, each snapshot of that origin volume becomes an independent thinly-provisioned volume. This means that instead of merging a snapshot with its origin volume, you may choose to delete the origin volume and then create a new thinly-provisioned snapshot using that independent volume as the origin volume for the new snapshot.

Although there are many advantages to using thin snapshot volumes, there are some use cases for which the older LVM snapshot volume feature may be more appropriate to your needs:

- You cannot change the chunk size of a thin pool. If the thin pool has a large chunk size (for example, 1MB) and you require a short-living snapshot for which a chunk size that large is not efficient, you may elect to use the older snapshot feature.
- You cannot limit the size of a thin snapshot volume; the snapshot will use all of the space in the thin pool, if necessary. This may not be appropriate for your needs.

In general, you should consider the specific requirements of your site when deciding which snapshot format to use.

For information on configuring thin snapshot volumes, see [Section 4.4.7, “Creating Thinly-Provisioned Snapshot Volumes”](#).

2.3.7. Cache Volumes

As of the Red Hat Enterprise Linux 7.1 release, LVM supports the use of fast block devices (such as SSD drives) as write-back or write-through caches for larger slower block devices. Users can create cache logical volumes to improve the performance of their existing logical volumes or create new cache logical volumes composed of a small and fast device coupled with a large and slow device.

For information on creating LVM cache volumes, see [Section 4.4.8, “Creating LVM Cache Logical Volumes”](#).

CHAPTER 3. LVM ADMINISTRATION OVERVIEW

This chapter provides an overview of the administrative procedures you use to configure LVM logical volumes. This chapter is intended to provide a general understanding of the steps involved. For specific step-by-step examples of common LVM configuration procedures, see [Chapter 5, LVM Configuration Examples](#).

For descriptions of the CLI commands you can use to perform LVM administration, see [Chapter 4, LVM Administration with CLI Commands](#).

3.1. LOGICAL VOLUME CREATION OVERVIEW

The following is a summary of the steps to perform to create an LVM logical volume.

1. Initialize the partitions you will use for the LVM volume as physical volumes (this labels them).
2. Create a volume group.
3. Create a logical volume.

After creating the logical volume you can create and mount the file system. The examples in this document use GFS2 file systems.

1. Create a GFS2 file system on the logical volume with the **mkfs.gfs2** command.
2. Create a new mount point with the **mkdir** command. In a clustered system, create the mount point on all nodes in the cluster.
3. Mount the file system. You may want to add a line to the **fstab** file for each node in the system.



NOTE

Although a GFS2 file system can be implemented in a standalone system or as part of a cluster configuration, for the Red Hat Enterprise Linux 7 release Red Hat does not support the use of GFS2 as a single-node file system. Red Hat will continue to support single-node GFS2 file systems for mounting snapshots of cluster file systems (for example, for backup purposes).

Creating the LVM volume is machine independent, since the storage area for LVM setup information is on the physical volumes and not the machine where the volume was created. Servers that use the storage have local copies, but can recreate that from what is on the physical volumes. You can attach physical volumes to a different server if the LVM versions are compatible.

3.2. GROWING A FILE SYSTEM ON A LOGICAL VOLUME

To grow a file system on a logical volume, perform the following steps:

1. Determine whether there is sufficient unallocated space in the existing volume group to extend the logical volume. If not, perform the following procedure:
 1. Create a new physical volume with the **pvcreate** command.
 2. Use the **vgextend** command to extend the volume group that contains the logical volume with the file system you are growing to include the new physical volume.

2. Once the volume group is large enough to include the larger file system, extend the logical volume with the **lvresize** command.
3. Resize the file system on the logical volume.

Note that you can use the **-r** option of the **lvresize** command to extend the logical volume and resize the underlying file system with a single command

3.3. LOGICAL VOLUME BACKUP

Metadata backups and archives are automatically created whenever there is a configuration change for a volume group or logical volume, unless this feature is disabled in the **lvm.conf** file. By default, the metadata backup is stored in the **/etc/lvm/backup** file and the metadata archives are stored in the **/etc/lvm/archive** file. How long the metadata archives stored in the **/etc/lvm/archive** file are kept and how many archive files are kept is determined by parameters you can set in the **lvm.conf** file. A daily system backup should include the contents of the **/etc/lvm** directory in the backup.

Note that a metadata backup does not back up the user and system data contained in the logical volumes.

You can manually back up the metadata to the **/etc/lvm/backup** file with the **vgcfgbackup** command. You can restore metadata with the **vgcfgrestore** command. The **vgcfgbackup** and **vgcfgrestore** commands are described in [Section 4.3.13, “Backing Up Volume Group Metadata”](#).

3.4. LOGGING

All message output passes through a logging module with independent choices of logging levels for:

- standard output/error
- syslog
- log file
- external log function

The logging levels are set in the **/etc/lvm/lvm.conf** file, which is described in [Appendix B, *The LVM Configuration Files*](#).

3.5. THE METADATA DAEMON (LVMETAD)

LVM can optionally use a central metadata cache, implemented through a daemon (**lvmemd**) and a **udev** rule. The metadata daemon has two main purposes: it improves performance of LVM commands and it allows **udev** to automatically activate logical volumes or entire volume groups as they become available to the system.

LVM is configured to make use of the daemon when the **global/use_lvmemd** variable is set to 1 in the **lvm.conf** configuration file. This is the default value. For information on the **lvm.conf** configuration file, see [Appendix B, *The LVM Configuration Files*](#).

**NOTE**

The **lvmetad** daemon is not currently supported across the nodes of a cluster, and requires that the locking type be local file-based locking. When you use the **lvmetad --enable-cluster/--disable-cluster** command, the **lvm.conf** file is configured appropriately, including the **use_lvmetad** setting (which should be 0 for **locking_type=3**). Note, however, that in a Pacemaker cluster, the **ocf:heartbeat:clvm** resource agent itself sets these parameters as part of the start procedure.

If you change the value of **use_lvmetad** from 1 to 0, you must reboot or stop the **lvmetad** service manually with the following command.

```
# systemctl stop lvm2-lvmetad.service
```

Normally, each LVM command issues a disk scan to find all relevant physical volumes and to read volume group metadata. However, if the metadata daemon is running and enabled, this expensive scan can be skipped. Instead, the **lvmetad** daemon scans each device only once, when it becomes available, using **udev** rules. This can save a significant amount of I/O and reduce the time required to complete LVM operations, particularly on systems with many disks.

When a new volume group is made available at runtime (for example, through hotplug or iSCSI), its logical volumes must be activated in order to be used. When the **lvmetad** daemon is enabled, the **activation/auto_activation_volume_list** option in the **lvm.conf** configuration file can be used to configure a list of volume groups or logical volumes (or both) that should be automatically activated. Without the **lvmetad** daemon, a manual activation is necessary.

**NOTE**

When the **lvmetad** daemon is running, the **filter =** setting in the **/etc/lvm/lvm.conf** file does not apply when you execute the **pvscan --cache device** command. To filter devices, you need to use the **global_filter =** setting. Devices that fail the global filter are not opened by LVM and are never scanned. You may need to use a global filter, for example, when you use LVM devices in VMs and you do not want the contents of the devices in the VMs to be scanned by the physical host.

3.6. DISPLAYING LVM INFORMATION WITH THE **lvm** COMMAND

The **lvm** command provides several built-in options that you can use to display information about LVM support and configuration.

- **lvm devtypes**

Displays the recognized build-in block device types (Red Hat Enterprise Linux release 6.6 and later).

- **lvm formats**

Displays recognized metadata formats.

- **lvm help**

Displays LVM help text.

- **lvm segtypes**

Displays recognized logical volume segment types.

- **lvm tags**

Displays any tags defined on this host. For information on LVM object tags, see [Appendix D, LVM Object Tags](#).

- **lvm version**

Displays the current version information.

CHAPTER 4. LVM ADMINISTRATION WITH CLI COMMANDS

This chapter summarizes the individual administrative tasks you can perform with the LVM Command Line Interface (CLI) commands to create and maintain logical volumes.

In addition to the LVM Command Line Interface (CLI), you can use System Storage Manager (SSM) to configure LVM logical volumes. For information on using SSM with LVM, see the *Storage Administration Guide*.

4.1. USING CLI COMMANDS

There are several general features of all LVM CLI commands.

When sizes are required in a command line argument, units can always be specified explicitly. If you do not specify a unit, then a default is assumed, usually KB or MB. LVM CLI commands do not accept fractions.

When specifying units in a command line argument, LVM is case-insensitive; specifying M or m is equivalent, for example, and powers of 2 (multiples of 1024) are used. However, when specifying the **--units** argument in a command, lower-case indicates that units are in multiples of 1024 while upper-case indicates that units are in multiples of 1000.

Where commands take volume group or logical volume names as arguments, the full path name is optional. A logical volume called **lv010** in a volume group called **vg0** can be specified as **vg0/lv010**. Where a list of volume groups is required but is left empty, a list of all volume groups will be substituted. Where a list of logical volumes is required but a volume group is given, a list of all the logical volumes in that volume group will be substituted. For example, the **lvdisplay vg0** command will display all the logical volumes in volume group **vg0**.

All LVM commands accept a **-v** argument, which can be entered multiple times to increase the output verbosity. For example, the following examples shows the default output of the **lvcreate** command.

```
# lvcreate -L 50MB new_vg
Rounding up size to full physical extent 52.00 MB
Logical volume "lv010" created
```

The following command shows the output of the **lvcreate** command with the **-v** argument.

```
# lvcreate -v -L 50MB new_vg
Finding volume group "new_vg"
Rounding up size to full physical extent 52.00 MB
Archiving volume group "new_vg" metadata (seqno 4).
Creating logical volume lv010
Creating volume group backup "/etc/lvm/backup/new_vg" (seqno 5).
Found volume group "new_vg"
Creating new_vg-lv010
Loading new_vg-lv010 table
Resuming new_vg-lv010 (253:2)
Clearing start of logical volume "lv010"
Creating volume group backup "/etc/lvm/backup/new_vg" (seqno 5).
Logical volume "lv010" created
```

You could also have used the **-vv**, **-vvv** or the **-vvvv** argument to display increasingly more details about the command execution. The **-vvvv** argument provides the maximum amount of information at

this time. The following example shows only the first few lines of output for the **lvcreate** command with the **-vvvv** argument specified.

```
# lvcreate -vvvv -L 50MB new_vg
#lvmcmdline.c:913      Processing: lvcreate -vvvv -L 50MB new_vg
#lvmcmdline.c:916      O_DIRECT will be used
#config/config.c:864   Setting global/locking_type to 1
#locking/locking.c:138 File-based locking selected.
#config/config.c:841   Setting global/locking_dir to /var/lock/lvm
#activate/activate.c:358 Getting target version for linear
#ioctl/libdm-iface.c:1569 dm version  OF  [16384]
#ioctl/libdm-iface.c:1569 dm versions  OF  [16384]
#activate/activate.c:358 Getting target version for striped
#ioctl/libdm-iface.c:1569 dm versions  OF  [16384]
#config/config.c:864   Setting activation/mirror_region_size to 512
...
```

You can display help for any of the LVM CLI commands with the **--help** argument of the command.

```
# commandname --help
```

To display the man page for a command, execute the **man** command:

```
# man commandname
```

The **man lv** command provides general online information about LVM.

All LVM objects are referenced internally by a UUID, which is assigned when you create the object. This can be useful in a situation where you remove a physical volume called **/dev/sdf** which is part of a volume group and, when you plug it back in, you find that it is now **/dev/sdk**. LVM will still find the physical volume because it identifies the physical volume by its UUID and not its device name. For information on specifying the UUID of a physical volume when creating a physical volume, see [Section 6.4, “Recovering Physical Volume Metadata”](#).

4.2. PHYSICAL VOLUME ADMINISTRATION

This section describes the commands that perform the various aspects of physical volume administration.

4.2.1. Creating Physical Volumes

The following subsections describe the commands used for creating physical volumes.

4.2.1.1. Setting the Partition Type

If you are using a whole disk device for your physical volume, the disk must have no partition table. For DOS disk partitions, the partition id should be set to 0x8e using the **fdisk** or **cfdisk** command or an equivalent. For whole disk devices only the partition table must be erased, which will effectively destroy all data on that disk. You can remove an existing partition table by zeroing the first sector with the following command:

```
# dd if=/dev/zero of=PhysicalVolume bs=512 count=1
```

4.2.1.2. Initializing Physical Volumes

Use the **pvcreate** command to initialize a block device to be used as a physical volume. Initialization is analogous to formatting a file system.

The following command initializes **/dev/sdd**, **/dev/sde**, and **/dev/sdf** as LVM physical volumes for later use as part of LVM logical volumes.

```
# pvcreate /dev/sdd /dev/sde /dev/sdf
```

To initialize partitions rather than whole disks: run the **pvcreate** command on the partition. The following example initializes the partition **/dev/hdb1** as an LVM physical volume for later use as part of an LVM logical volume.

```
# pvcreate /dev/hdb1
```

4.2.1.3. Scanning for Block Devices

You can scan for block devices that may be used as physical volumes with the **lvmdiskscan** command, as shown in the following example.

```
# lvmdiskscan
/dev/ram0          [      16.00 MB]
/dev/sda          [      17.15 GB]
/dev/root         [      13.69 GB]
/dev/ram          [      16.00 MB]
/dev/sda1         [      17.14 GB] LVM physical volume
/dev/VolGroup00/LogVol01 [    512.00 MB]
/dev/ram2         [      16.00 MB]
/dev/new_vg/lvol0 [      52.00 MB]
/dev/ram3         [      16.00 MB]
/dev/pkl_new_vg/sparkie_lv [    7.14 GB]
/dev/ram4         [      16.00 MB]
/dev/ram5         [      16.00 MB]
/dev/ram6         [      16.00 MB]
/dev/ram7         [      16.00 MB]
/dev/ram8         [      16.00 MB]
/dev/ram9         [      16.00 MB]
/dev/ram10        [      16.00 MB]
/dev/ram11        [      16.00 MB]
/dev/ram12        [      16.00 MB]
/dev/ram13        [      16.00 MB]
/dev/ram14        [      16.00 MB]
/dev/ram15        [      16.00 MB]
/dev/sdb          [      17.15 GB]
/dev/sdb1         [      17.14 GB] LVM physical volume
/dev/sdc          [      17.15 GB]
/dev/sdc1         [      17.14 GB] LVM physical volume
/dev/sdd          [      17.15 GB]
/dev/sdd1         [      17.14 GB] LVM physical volume
7 disks
17 partitions
0 LVM physical volume whole disks
4 LVM physical volumes
```

4.2.2. Displaying Physical Volumes

There are three commands you can use to display properties of LVM physical volumes: **pvs**, **pvdisplay**, and **pvscan**.

The **pvs** command provides physical volume information in a configurable form, displaying one line per physical volume. The **pvs** command provides a great deal of format control, and is useful for scripting. For information on using the **pvs** command to customize your output, see [Section 4.8, “Customized Reporting for LVM”](#).

The **pvdisplay** command provides a verbose multi-line output for each physical volume. It displays physical properties (size, extents, volume group, and so on) in a fixed format.

The following example shows the output of the **pvdisplay** command for a single physical volume.

```
# pvdisplay
--- Physical volume ---
PV Name           /dev/sdc1
VG Name           new_vg
PV Size           17.14 GB / not usable 3.40 MB
Allocatable       yes
PE Size (KByte)   4096
Total PE          4388
Free PE           4375
Allocated PE      13
PV UUID           Joqlch-yWSj - kuEn - IdwM - 01S9 - X08M - mcpsVe
```

The **pvscan** command scans all supported LVM block devices in the system for physical volumes.

The following command shows all physical devices found:

```
# pvscan
PV /dev/sdb2    VG vg0    lvm2 [964.00 MB / 0    free]
PV /dev/sdc1    VG vg0    lvm2 [964.00 MB / 428.00 MB free]
PV /dev/sdc2    VG          lvm2 [964.84 MB]
Total: 3 [2.83 GB] / in use: 2 [1.88 GB] / in no VG: 1 [964.84 MB]
```

You can define a filter in the **lvm.conf** file so that this command will avoid scanning specific physical volumes. For information on using filters to control which devices are scanned, see [Section 4.5, “Controlling LVM Device Scans with Filters”](#).

4.2.3. Preventing Allocation on a Physical Volume

You can prevent allocation of physical extents on the free space of one or more physical volumes with the **pvchange** command. This may be necessary if there are disk errors, or if you will be removing the physical volume.

The following command disallows the allocation of physical extents on **/dev/sdk1**.

```
# pvchange -x n /dev/sdk1
```

You can also use the **-xy** arguments of the **pvchange** command to allow allocation where it had previously been disallowed.

4.2.4. Resizing a Physical Volume

If you need to change the size of an underlying block device for any reason, use the **pvresize** command to update LVM with the new size. You can execute this command while LVM is using the physical volume.

4.2.5. Removing Physical Volumes

If a device is no longer required for use by LVM, you can remove the LVM label with the **pvremove** command. Executing the **pvremove** command zeroes the LVM metadata on an empty physical volume.

If the physical volume you want to remove is currently part of a volume group, you must remove it from the volume group with the **vgreduce** command, as described in [Section 4.3.7, “Removing Physical Volumes from a Volume Group”](#).

```
# pvremove /dev/ram15
Labels on physical volume "/dev/ram15" successfully wiped
```

4.3. VOLUME GROUP ADMINISTRATION

This section describes the commands that perform the various aspects of volume group administration.

4.3.1. Creating Volume Groups

To create a volume group from one or more physical volumes, use the **vgcreate** command. The **vgcreate** command creates a new volume group by name and adds at least one physical volume to it.

The following command creates a volume group named **vg1** that contains physical volumes **/dev/sdd1** and **/dev/sde1**.

```
# vgcreate vg1 /dev/sdd1 /dev/sde1
```

When physical volumes are used to create a volume group, its disk space is divided into 4MB extents, by default. This extent is the minimum amount by which the logical volume may be increased or decreased in size. Large numbers of extents will have no impact on I/O performance of the logical volume.

You can specify the extent size with the **-s** option to the **vgcreate** command if the default extent size is not suitable. You can put limits on the number of physical or logical volumes the volume group can have by using the **-p** and **-l** arguments of the **vgcreate** command.

By default, a volume group allocates physical extents according to common-sense rules such as not placing parallel stripes on the same physical volume. This is the **normal** allocation policy. You can use the **--alloc** argument of the **vgcreate** command to specify an allocation policy of **contiguous**, **anywhere**, or **cling**. In general, allocation policies other than **normal** are required only in special cases where you need to specify unusual or nonstandard extent allocation. For further information on how LVM allocates physical extents, see [Section 4.3.2, “LVM Allocation”](#).

LVM volume groups and underlying logical volumes are included in the device special file directory tree in the **/dev** directory with the following layout:

```
/dev/vg/lv/
```

For example, if you create two volume groups **myvg1** and **myvg2**, each with three logical volumes named **lv01**, **lv02**, and **lv03**, this creates six device special files:

```
/dev/myvg1/lv01
/dev/myvg1/lv02
/dev/myvg1/lv03
/dev/myvg2/lv01
/dev/myvg2/lv02
/dev/myvg2/lv03
```

The device special files are not present if the corresponding logical volume is not currently active.

The maximum device size with LVM is 8 Exabytes on 64-bit CPUs.

4.3.2. LVM Allocation

When an LVM operation needs to allocate physical extents for one or more logical volumes, the allocation proceeds as follows:

- The complete set of unallocated physical extents in the volume group is generated for consideration. If you supply any ranges of physical extents at the end of the command line, only unallocated physical extents within those ranges on the specified physical volumes are considered.
- Each allocation policy is tried in turn, starting with the strictest policy (**contiguous**) and ending with the allocation policy specified using the **--alloc** option or set as the default for the particular logical volume or volume group. For each policy, working from the lowest-numbered logical extent of the empty logical volume space that needs to be filled, as much space as possible is allocated, according to the restrictions imposed by the allocation policy. If more space is needed, LVM moves on to the next policy.

The allocation policy restrictions are as follows:

- An allocation policy of **contiguous** requires that the physical location of any logical extent that is not the first logical extent of a logical volume is adjacent to the physical location of the logical extent immediately preceding it.

When a logical volume is striped or mirrored, the **contiguous** allocation restriction is applied independently to each stripe or mirror image (leg) that needs space.

- An allocation policy of **cling** requires that the physical volume used for any logical extent be added to an existing logical volume that is already in use by at least one logical extent earlier in that logical volume. If the configuration parameter **allocation/cling_tag_list** is defined, then two physical volumes are considered to match if any of the listed tags is present on both physical volumes. This allows groups of physical volumes with similar properties (such as their physical location) to be tagged and treated as equivalent for allocation purposes. For more information on using the **cling** policy in conjunction with LVM tags to specify which additional physical volumes to use when extending an LVM volume, see [Section 4.4.19, “Extending a Logical Volume with the **cling** Allocation Policy”](#).

When a Logical Volume is striped or mirrored, the **cling** allocation restriction is applied independently to each stripe or mirror image (leg) that needs space.

- An allocation policy of **normal** will not choose a physical extent that shares the same physical volume as a logical extent already allocated to a parallel logical volume (that is, a different stripe or mirror image/leg) at the same offset within that parallel logical volume.

When allocating a mirror log at the same time as logical volumes to hold the mirror data, an allocation policy of **normal** will first try to select different physical volumes for the log and the data. If that is not possible and the **allocation/mirror_logs_require_separate_pvs** configuration parameter is set to 0, it will then allow the log to share physical volume(s) with part of the data.

Similarly, when allocating thin pool metadata, an allocation policy of **normal** will follow the same considerations as for allocation of a mirror log, based on the value of the **allocation/thin_pool_metadata_require_separate_pvs** configuration parameter.

- If there are sufficient free extents to satisfy an allocation request but a **normal** allocation policy would not use them, the **anywhere** allocation policy will, even if that reduces performance by placing two stripes on the same physical volume.

The allocation policies can be changed using the **vgchange** command.



NOTE

If you rely upon any layout behavior beyond that documented in this section according to the defined allocation policies, you should note that this might change in future versions of the code. For example, if you supply on the command line two empty physical volumes that have an identical number of free physical extents available for allocation, LVM currently considers using each of them in the order they are listed; there is no guarantee that future releases will maintain that property. If it is important to obtain a specific layout for a particular Logical Volume, then you should build it up through a sequence of **lvcreate** and **lvconvert** steps such that the allocation policies applied to each step leave LVM no discretion over the layout.

To view the way the allocation process currently works in any specific case, you can read the debug logging output, for example by adding the **-vvvv** option to a command.

4.3.3. Creating Volume Groups in a Cluster

You create CLVM volume groups in a cluster environment with the **vgcreate** command, just as you create them on a single node.



NOTE

In Red Hat Enterprise Linux 7, clusters are managed through Pacemaker. Clustered LVM logical volumes are supported only in conjunction with Pacemaker clusters, and must be configured as cluster resources. For general information on configuring LVM volumes in a cluster, see [Section 1.4, “LVM Logical Volumes in a Red Hat High Availability Cluster”](#).

Volume groups that are shared by members of the cluster should be created with the clustered attribute set with the **vgcreate -cy** or **vgchange -cy** command. The clustered attribute is set automatically if CLVMD is running. This clustered attribute signals that this volume group should be managed and protected by CLVMD. When creating any volume group that is not shared by the cluster and should only be visible to a single host, this clustered attribute should be disabled with the **vgcreate -cn** or **vgchange -cn** command.

By default, volume groups created with with the clustered attribute on shared storage are visible to all computers that have access to the shared storage. It is possible, however, to create volume groups that are local, visible only to one node in the cluster, by using the **-cn** option of the **vgcreate** command.

The following command, when executed in a cluster environment, creates a volume group that is local to the node from which the command was executed. The command creates a local volume named **vg1** that contains physical volumes **/dev/sdd1** and **/dev/sde1**.

```
# vgcreate -c n vg1 /dev/sdd1 /dev/sde1
```

You can change whether an existing volume group is local or clustered with the **-c** option of the **vgchange** command, which is described in [Section 4.3.9, “Changing the Parameters of a Volume Group”](#).

You can check whether an existing volume group is a clustered volume group with the **vgs** command, which displays the **c** attribute if the volume is clustered. The following command displays the attributes of the volume groups **VolGroup00** and **testvg1**. In this example, **VolGroup00** is not clustered, while **testvg1** is clustered, as indicated by the **c** attribute under the **Attr** heading.

```
# vgs
VG                #PV #LV #SN Attr   VSize  VFree
VolGroup00        1  2  0 wz--n- 19.88G  0
testvg1           1  1  0 wz--nc 46.00G 8.00M
```

For more information on the **vgs** command, see [Section 4.3.5, “Displaying Volume Groups”](#) [Section 4.8, “Customized Reporting for LVM”](#), and the **vgs** man page.

4.3.4. Adding Physical Volumes to a Volume Group

To add additional physical volumes to an existing volume group, use the **vgextend** command. The **vgextend** command increases a volume group's capacity by adding one or more free physical volumes.

The following command adds the physical volume **/dev/sdf1** to the volume group **vg1**.

```
# vgextend vg1 /dev/sdf1
```

4.3.5. Displaying Volume Groups

There are two commands you can use to display properties of LVM volume groups: **vgs** and **vgdisplay**.

The **vgscan** command, which scans all the disks for volume groups and rebuilds the LVM cache file, also displays the volume groups. For information on the **vgscan** command, see [Section 4.3.6, “Scanning Disks for Volume Groups to Build the Cache File”](#).

The **vgs** command provides volume group information in a configurable form, displaying one line per volume group. The **vgs** command provides a great deal of format control, and is useful for scripting. For information on using the **vgs** command to customize your output, see [Section 4.8, “Customized Reporting for LVM”](#).

The **vgdisplay** command displays volume group properties (such as size, extents, number of physical volumes, and so on) in a fixed form. The following example shows the output of the **vgdisplay** command for the volume group **new_vg**. If you do not specify a volume group, all existing volume groups

are displayed.

```
# vgdisplay new_vg
--- Volume group ---
VG Name                new_vg
System ID
Format                 lvm2
Metadata Areas         3
Metadata Sequence No   11
VG Access               read/write
VG Status               resizable
MAX LV                 0
Cur LV                 1
Open LV                 0
Max PV                 0
Cur PV                 3
Act PV                 3
VG Size                 51.42 GB
PE Size                 4.00 MB
Total PE                13164
Alloc PE / Size         13 / 52.00 MB
Free PE / Size           13151 / 51.37 GB
VG UUID                 jxQJ0a-ZKk0-OpM0-0118-nlw0-wwqd-fD5D32
```

4.3.6. Scanning Disks for Volume Groups to Build the Cache File

The **vgscan** command scans all supported disk devices in the system looking for LVM physical volumes and volume groups. This builds the LVM cache file in the `/etc/lvm/cache/.cache` file, which maintains a listing of current LVM devices.

LVM runs the **vgscan** command automatically at system startup and at other times during LVM operation, such as when you execute the **vgcreate** command or when LVM detects an inconsistency.



NOTE

You may need to run the **vgscan** command manually when you change your hardware configuration and add or delete a device from a node, causing new devices to be visible to the system that were not present at system bootup. This may be necessary, for example, when you add new disks to the system on a SAN or hotplug a new disk that has been labeled as a physical volume.

You can define a filter in the `/etc/lvm/lvm.conf` file to restrict the scan to avoid specific devices. For information on using filters to control which devices are scanned, see [Section 4.5, “Controlling LVM Device Scans with Filters”](#).

The following example shows the output of the **vgscan** command.

```
# vgscan
Reading all physical volumes. This may take a while...
Found volume group "new_vg" using metadata type lvm2
Found volume group "officevg" using metadata type lvm2
```

4.3.7. Removing Physical Volumes from a Volume Group

To remove unused physical volumes from a volume group, use the **vgreduce** command. The **vgreduce** command shrinks a volume group's capacity by removing one or more empty physical volumes. This frees those physical volumes to be used in different volume groups or to be removed from the system.

Before removing a physical volume from a volume group, you can make sure that the physical volume is not used by any logical volumes by using the **pvdisplay** command.

```
# pvdisplay /dev/hda1

-- Physical volume ---
PV Name           /dev/hda1
VG Name           myvg
PV Size           1.95 GB / NOT usable 4 MB [LVM: 122 KB]
PV#               1
PV Status          available
Allocatable        yes (but full)
Cur LV           1
PE Size (KByte)    4096
Total PE          499
Free PE           0
Allocated PE       499
PV UUID           Sd44tK-9IRw-SrMC-M0kn-76iP-iftz-OVSen7
```

If the physical volume is still being used you will have to migrate the data to another physical volume using the **pvmove** command. Then use the **vgreduce** command to remove the physical volume.

The following command removes the physical volume **/dev/hda1** from the volume group **my_volume_group**.

```
# vgreduce my_volume_group /dev/hda1
```

If a logical volume contains a physical volume that fails, you cannot use that logical volume. To remove missing physical volumes from a volume group, you can use the **--removemissing** parameter of the **vgreduce** command, if there are no logical volumes that are allocated on the missing physical volumes.

If the physical volume that fails contains a mirror image of a logical volume of a **mirror** segment type, you can remove that image from the mirror with the **vgreduce --removemissing --mirroronly --force** command. This removes only the logical volumes that are mirror images from the physical volume.

For information on recovering from LVM mirror failure, see [Section 6.3, “Recovering from LVM Mirror Failure”](#). For information on removing lost physical volumes from a volume group, see [Section 6.6, “Removing Lost Physical Volumes from a Volume Group”](#)

4.3.8. Activating and Deactivating Volume Groups

When you create a volume group it is, by default, activated. This means that the logical volumes in that group are accessible and subject to change.

There are various circumstances for which you need to make a volume group inactive and thus unknown to the kernel. To deactivate or activate a volume group, use the **-a (--available)** argument of the **vgchange** command.

The following example deactivates the volume group **my_volume_group**.

```
# vgchange -a n my_volume_group
```

If clustered locking is enabled, add 'e' to activate or deactivate a volume group exclusively on one node or 'l' to activate or deactivate a volume group only on the local node. Logical volumes with single-host snapshots are always activated exclusively because they can only be used on one node at once.

You can deactivate individual logical volumes with the **lvchange** command, as described in [Section 4.4.11, “Changing the Parameters of a Logical Volume Group”](#). For information on activating logical volumes on individual nodes in a cluster, see [Section 4.7, “Activating Logical Volumes on Individual Nodes in a Cluster”](#).

4.3.9. Changing the Parameters of a Volume Group

The **vgchange** command is used to deactivate and activate volume groups, as described in [Section 4.3.8, “Activating and Deactivating Volume Groups”](#). You can also use this command to change several volume group parameters for an existing volume group.

The following command changes the maximum number of logical volumes of volume group **vg00** to 128.

```
# vgchange -l 128 /dev/vg00
```

For a description of the volume group parameters you can change with the **vgchange** command, see the **vgchange(8)** man page.

4.3.10. Removing Volume Groups

To remove a volume group that contains no logical volumes, use the **vgremove** command.

```
# vgremove officevg
Volume group "officevg" successfully removed
```

4.3.11. Splitting a Volume Group

To split the physical volumes of a volume group and create a new volume group, use the **vgsplit** command.

Logical volumes cannot be split between volume groups. Each existing logical volume must be entirely on the physical volumes forming either the old or the new volume group. If necessary, however, you can use the **pvmove** command to force the split.

The following example splits the new volume group **smallvg** from the original volume group **bigvg**.

```
# vgsplit bigvg smallvg /dev/ram15
Volume group "smallvg" successfully split from "bigvg"
```

4.3.12. Combining Volume Groups

To combine two volume groups into a single volume group, use the **vgmerge** command. You can merge an inactive "source" volume with an active or an inactive "destination" volume if the physical extent sizes of the volume are equal and the physical and logical volume summaries of both volume groups fit into the

destination volume groups limits.

The following command merges the inactive volume group **my_vg** into the active or inactive volume group **databases** giving verbose runtime information.

```
# vgmerge -v databases my_vg
```

4.3.13. Backing Up Volume Group Metadata

Metadata backups and archives are automatically created on every configuration change to a volume group or logical volume unless disabled in the **lvm.conf** file. By default, the metadata backup is stored in the **/etc/lvm/backup** file and the metadata archives are stored in the **/etc/lvm/archives** file. You can manually back up the metadata to the **/etc/lvm/backup** file with the **vgcfgbackup** command.

The **vgcfgrestore** command restores the metadata of a volume group from the archive to all the physical volumes in the volume groups.

For an example of using the **vgcfgrestore** command to recover physical volume metadata, see [Section 6.4, “Recovering Physical Volume Metadata”](#).

4.3.14. Renaming a Volume Group

Use the **vgrename** command to rename an existing volume group.

Either of the following commands renames the existing volume group **vg02** to **my_volume_group**

```
# vgrename /dev/vg02 /dev/my_volume_group
```

```
# vgrename vg02 my_volume_group
```

4.3.15. Moving a Volume Group to Another System

You can move an entire LVM volume group to another system. It is recommended that you use the **vgexport** and **vgimport** commands when you do this.



NOTE

You can use the **--force** argument of the **vgimport** command. This allows you to import volume groups that are missing physical volumes and subsequently run the **vgreduce --removemissing** command.

The **vgexport** command makes an inactive volume group inaccessible to the system, which allows you to detach its physical volumes. The **vgimport** command makes a volume group accessible to a machine again after the **vgexport** command has made it inactive.

To move a volume group from one system to another, perform the following steps:

1. Make sure that no users are accessing files on the active volumes in the volume group, then unmount the logical volumes.

2. Use the **-a n** argument of the **vgchange** command to mark the volume group as inactive, which prevents any further activity on the volume group.
3. Use the **vgexport** command to export the volume group. This prevents it from being accessed by the system from which you are removing it.

After you export the volume group, the physical volume will show up as being in an exported volume group when you execute the **pvscan** command, as in the following example.

```
# pvscan
PV /dev/sda1    is in exported VG myvg [17.15 GB / 7.15 GB free]
PV /dev/sdc1    is in exported VG myvg [17.15 GB / 15.15 GB free]
PV /dev/sdd1    is in exported VG myvg [17.15 GB / 15.15 GB free]
...
```

When the system is next shut down, you can unplug the disks that constitute the volume group and connect them to the new system.

4. When the disks are plugged into the new system, use the **vgimport** command to import the volume group, making it accessible to the new system.
5. Activate the volume group with the **-a y** argument of the **vgchange** command.
6. Mount the file system to make it available for use.

4.3.16. Recreating a Volume Group Directory

To recreate a volume group directory and logical volume special files, use the **vgmknodes** command. This command checks the LVM2 special files in the **/dev** directory that are needed for active logical volumes. It creates any special files that are missing and removes unused ones.

You can incorporate the **vgmknodes** command into the **vgscan** command by specifying the **mknodes** argument to the **vgscan** command.

4.4. LOGICAL VOLUME ADMINISTRATION

This section describes the commands that perform the various aspects of logical volume administration.

4.4.1. Creating Linear Logical Volumes

To create a logical volume, use the **lvcreate** command. If you do not specify a name for the logical volume, the default name **lvol#** is used where **#** is the internal number of the logical volume.

When you create a logical volume, the logical volume is carved from a volume group using the free extents on the physical volumes that make up the volume group. Normally logical volumes use up any space available on the underlying physical volumes on a next-free basis. Modifying the logical volume frees and reallocates space in the physical volumes.

The following command creates a logical volume 10 gigabytes in size in the volume group **vg1**.

```
# lvcreate -L 10G vg1
```

The default unit for logical volume size is megabytes. The following command creates a 1500 megabyte linear logical volume named **testlv** in the volume group **testvg**, creating the block device **/dev/testvg/testlv**.

```
# lvcreate -L 1500 -n testlv testvg
```

The following command creates a 50 gigabyte logical volume named **gfslv** from the free extents in volume group **vg0**.

```
# lvcreate -L 50G -n gfslv vg0
```

You can use the **-l** argument of the **lvcreate** command to specify the size of the logical volume in extents. You can also use this argument to specify the percentage of the size of a related volume group, logical volume, or set of physical volumes. The suffix **%VG** denotes the total size of the volume group, the suffix **%FREE** the remaining free space in the volume group, and the suffix **%PVS** the free space in the specified physical volumes. For a snapshot, the size can be expressed as a percentage of the total size of the origin logical volume with the suffix **%ORIGIN** (100%ORIGIN provides space for the whole origin). When expressed as a percentage, the size defines an upper limit for the number of logical extents in the new logical volume. The precise number of logical extents in the new LV is not determined until the command has completed.

The following command creates a logical volume called **mylv** that uses 60% of the total space in volume group **testvg**.

```
# lvcreate -l 60%VG -n mylv testvg
```

The following command creates a logical volume called **yourlv** that uses all of the unallocated space in the volume group **testvg**.

```
# lvcreate -l 100%FREE -n yourlv testvg
```

You can use **-l** argument of the **lvcreate** command to create a logical volume that uses the entire volume group. Another way to create a logical volume that uses the entire volume group is to use the **vgdisplay** command to find the "Total PE" size and to use those results as input to the **lvcreate** command.

The following commands create a logical volume called **mylv** that fills the volume group named **testvg**.

```
# vgdisplay testvg | grep "Total PE"
Total PE          10230
# lvcreate -l 10230 -n mylv testvg
```

The underlying physical volumes used to create a logical volume can be important if the physical volume needs to be removed, so you may need to consider this possibility when you create the logical volume. For information on removing a physical volume from a volume group, see [Section 4.3.7, "Removing Physical Volumes from a Volume Group"](#).

To create a logical volume to be allocated from a specific physical volume in the volume group, specify the physical volume or volumes at the end at the **lvcreate** command line. The following command creates a logical volume named **testlv** in volume group **testvg** allocated from the physical volume **/dev/sdg1**,

```
# lvcreate -L 1500 -n testlv testvg /dev/sdg1
```

You can specify which extents of a physical volume are to be used for a logical volume. The following example creates a linear logical volume out of extents 0 through 24 of physical volume **/dev/sda1** and extents 50 through 124 of physical volume **/dev/sdb1** in volume group **testvg**.

```
# lvcreate -l 100 -n testlv testvg /dev/sda1:0-24 /dev/sdb1:50-124
```

The following example creates a linear logical volume out of extents 0 through 25 of physical volume **/dev/sda1** and then continues laying out the logical volume at extent 100.

```
# lvcreate -l 100 -n testlv testvg /dev/sda1:0-25:100-
```

The default policy for how the extents of a logical volume are allocated is **inherit**, which applies the same policy as for the volume group. These policies can be changed using the **lvchange** command. For information on allocation policies, see [Section 4.3.1, “Creating Volume Groups”](#).

4.4.2. Creating Striped Volumes

For large sequential reads and writes, creating a striped logical volume can improve the efficiency of the data I/O. For general information about striped volumes, see [Section 2.3.2, “Striped Logical Volumes”](#).

When you create a striped logical volume, you specify the number of stripes with the **-i** argument of the **lvcreate** command. This determines over how many physical volumes the logical volume will be striped. The number of stripes cannot be greater than the number of physical volumes in the volume group (unless the **--alloc anywhere** argument is used).

If the underlying physical devices that make up a striped logical volume are different sizes, the maximum size of the striped volume is determined by the smallest underlying device. For example, in a two-legged stripe, the maximum size is twice the size of the smaller device. In a three-legged stripe, the maximum size is three times the size of the smallest device.

The following command creates a striped logical volume across 2 physical volumes with a stripe of 64 kilobytes. The logical volume is 50 gigabytes in size, is named **gfslv**, and is carved out of volume group **vg0**.

```
# lvcreate -L 50G -i 2 -I 64 -n gfslv vg0
```

As with linear volumes, you can specify the extents of the physical volume that you are using for the stripe. The following command creates a striped volume 100 extents in size that stripes across two physical volumes, is named **stripelv** and is in volume group **testvg**. The stripe will use sectors 0-49 of **/dev/sda1** and sectors 50-99 of **/dev/sdb1**.

```
# lvcreate -l 100 -i 2 -n stripelv testvg /dev/sda1:0-49 /dev/sdb1:50-99
Using default stripesize 64.00 KB
Logical volume "stripelv" created
```

4.4.3. RAID Logical Volumes

LVM supports RAID0/1/4/5/6/10.



NOTE

RAID logical volumes are not cluster-aware. While RAID logical volumes can be created and activated exclusively on one machine, they cannot be activated simultaneously on more than one machine. If you require non-exclusive mirrored volumes, you must create the volumes with a **mirror** segment type, as described in [Section 4.4.4, “Creating Mirrored Volumes”](#).

To create a RAID logical volume, you specify a raid type as the **--type** argument of the **lvcreate** command. [Table 4.1, “RAID Segment Types”](#) describes the possible RAID segment types.

Table 4.1. RAID Segment Types

Segment type	Description
raid1	RAID1 mirroring. This is the default value for the --type argument of the lvcreate command when you specify the -m but you do not specify striping.
raid4	RAID4 dedicated parity disk
raid5	Same as raid5_ls
raid5_la	RAID5 left asymmetric. Rotating parity 0 with data continuation
raid5_ra	RAID5 right asymmetric. Rotating parity N with data continuation
raid5_ls	RAID5 left symmetric. Rotating parity 0 with data restart
raid5_rs	RAID5 right symmetric. Rotating parity N with data restart
raid6	Same as raid6_zr
raid6_zr	RAID6 zero restart Rotating parity zero (left-to-right) with data restart

Segment type	Description
raid6_nr	RAID6 N restart
	Rotating parity N (left-to-right) with data restart
raid6_nc	RAID6 N continue
	Rotating parity N (left-to-right) with data continuation
raid10	Striped mirrors. This is the default value for the --type argument of the lvcreate command if you specify the -m and you specify a number of stripes that is greater than 1.
	Striping of mirror sets
raid0/raid0_meta (Red Hat Enterprise Linux 7.3 and later)	Striping. RAID0 spreads logical volume data across multiple data subvolumes in units of stripe size. This is used to increase performance. Logical volume data will be lost if any of the data subvolumes fail. For information on creating RAID0 volumes, see Section 4.4.3.1, “Creating RAID0 Volumes (Red Hat Enterprise Linux 7.3 and Later)” .

For most users, specifying one of the five available primary types (**raid1**, **raid4**, **raid5**, **raid6**, **raid10**) should be sufficient.

When you create a RAID logical volume, LVM creates a metadata subvolume that is one extent in size for every data or parity subvolume in the array. For example, creating a 2-way RAID1 array results in two metadata subvolumes (**lv_rmeta_0** and **lv_rmeta_1**) and two data subvolumes (**lv_rimage_0** and **lv_rimage_1**). Similarly, creating a 3-way stripe (plus 1 implicit parity device) RAID4 results in 4 metadata subvolumes (**lv_rmeta_0**, **lv_rmeta_1**, **lv_rmeta_2**, and **lv_rmeta_3**) and 4 data subvolumes (**lv_rimage_0**, **lv_rimage_1**, **lv_rimage_2**, and **lv_rimage_3**).

The following command creates a 2-way RAID1 array named **my_lv** in the volume group **my_vg** that is one gigabyte in size.

```
# lvcreate --type raid1 -m 1 -L 1G -n my_lv my_vg
```

You can create RAID1 arrays with different numbers of copies according to the value you specify for the **-m** argument. Similarly, you specify the number of stripes for a RAID 4/5/6 logical volume with the **-i** argument. You can also specify the stripe size with the **-I** argument.

The following command creates a RAID5 array (3 stripes + 1 implicit parity drive) named **my_lv** in the volume group **my_vg** that is one gigabyte in size. Note that you specify the number of stripes just as you do for an LVM striped volume; the correct number of parity drives is added automatically.

```
# lvcreate --type raid5 -i 3 -L 1G -n my_lv my_vg
```

The following command creates a RAID6 array (3 stripes + 2 implicit parity drives) named **my_lv** in the volume group **my_vg** that is one gigabyte in size.

```
# lvcreate --type raid6 -i 3 -L 1G -n my_lv my_vg
```

After you have created a RAID logical volume with LVM, you can activate, change, remove, display, and use the volume just as you would any other LVM logical volume.

When you create RAID10 logical volumes, the background I/O required to initialize the logical volumes with a **sync** operation can crowd out other I/O operations to LVM devices, such as updates to volume group metadata, particularly when you are creating many RAID logical volumes. This can cause the other LVM operations to slow down.

You can control the rate at which a RAID logical volume is initialized by implementing recovery throttling. You control the rate at which **sync** operations are performed by setting the minimum and maximum I/O rate for those operations with the **--minrecoveryrate** and **--maxrecoveryrate** options of the **lvcreate** command. You specify these options as follows.

- **--maxrecoveryrate *Rate*[bBsSkKmMgG]**

Sets the maximum recovery rate for a RAID logical volume so that it will not crowd out nominal I/O operations. The *Rate* is specified as an amount per second for each device in the array. If no suffix is given, then kiB/sec/device is assumed. Setting the recovery rate to 0 means it will be unbounded.

- **--minrecoveryrate *Rate*[bBsSkKmMgG]**

Sets the minimum recovery rate for a RAID logical volume to ensure that I/O for **sync** operations achieves a minimum throughput, even when heavy nominal I/O is present. The *Rate* is specified as an amount per second for each device in the array. If no suffix is given, then kiB/sec/device is assumed.

The following command creates a 2-way RAID10 array with 3 stripes that is 10 gigabytes in size with a maximum recovery rate of 128 kiB/sec/device. The array is named **my_lv** and is in the volume group **my_vg**.

```
# lvcreate --type raid10 -i 2 -m 1 -L 10G --maxrecoveryrate 128 -n my_lv my_vg
```

You can also specify minimum and maximum recovery rates for a RAID scrubbing operation. For information on RAID scrubbing, see [Section 4.4.3.11, “Scrubbing a RAID Logical Volume”](#).



NOTE

You can generate commands to create logical volumes on RAID storage with the LVM RAID Calculator application. This application uses the information you input about your current or planned storage to generate these commands. The LVM RAID Calculator application can be found at <https://access.redhat.com/labs/lvmraidcalculator/>.

The following sections describes the administrative tasks you can perform on LVM RAID devices:

- [Section 4.4.3.1, “Creating RAID0 Volumes \(Red Hat Enterprise Linux 7.3 and Later\)”](#)
- [Section 4.4.3.2, “Converting a Linear Device to a RAID Device”](#)

- [Section 4.4.3.3, “Converting an LVM RAID1 Logical Volume to an LVM Linear Logical Volume”](#)
- [Section 4.4.3.4, “Converting a Mirrored LVM Device to a RAID1 Device”](#)
- [Section 4.4.3.5, “Resizing a RAID Logical Volume”](#)
- [Section 4.4.3.6, “Changing the Number of Images in an Existing RAID1 Device”](#)
- [Section 4.4.3.7, “Splitting off a RAID Image as a Separate Logical Volume”](#)
- [Section 4.4.3.8, “Splitting and Merging a RAID Image”](#)
- [Section 4.4.3.9, “Setting a RAID fault policy”](#)
- [Section 4.4.3.10, “Replacing a RAID device”](#)
- [Section 4.4.3.11, “Scrubbing a RAID Logical Volume”](#)
- [Section 4.4.3.12, “RAID Takeover \(Red Hat Enterprise Linux 7.4 and Later\)”](#)
- [Section 4.4.3.13, “Reshaping a RAID Logical Volume \(Red Hat Enterprise Linux 7.4 and Later\)”](#)
- [Section 4.4.3.14, “Controlling I/O Operations on a RAID1 Logical Volume”](#)
- [Section 4.4.3.15, “Changing the region size on a RAID Logical Volume \(Red Hat Enterprise Linux 7.4 and later\)”](#)

4.4.3.1. Creating RAID0 Volumes (Red Hat Enterprise Linux 7.3 and Later)

The format for the command to create a RAID0 volume is as follows.

```
lvcreate --type raid0[_meta] --stripes Stripes --stripesize StripeSize
VolumeGroup [PhysicalVolumePath ...]
```

Table 4.2. RAID0 Command Creation parameters

Parameter	Description
--type raid0[_meta]	Specifying raid0 creates a RAID0 volume without metadata volumes. Specifying raid0_meta creates a RAID0 volume with metadata volumes. Because RAID0 is non-resilient, it does not have to store any mirrored data blocks as RAID1/10 or calculate and store any parity blocks as RAID4/5/6 do. Hence, it does not need metadata volumes to keep state about resynchronization progress of mirrored or parity blocks. Metadata volumes become mandatory on a conversion from RAID0 to RAID4/5/6/10, however, and specifying raid0_meta preallocates those metadata volumes to prevent a respective allocation failure.
--stripes <i>Stripes</i>	Specifies the number of devices to spread the logical volume across.
--stripesize <i>StripeSize</i>	Specifies the size of each stripe in kilobytes. This is the amount of data that is written to one device before moving to the next device.
<i>VolumeGroup</i>	Specifies the volume group to use.

Parameter	Description
<i>PhysicalVolumePath</i> ...	Specifies the devices to use. If this is not specified, LVM will choose the number of devices specified by the <i>Stripes</i> option, one for each stripe.

4.4.3.2. Converting a Linear Device to a RAID Device

You can convert an existing linear logical volume to a RAID device by using the **--type** argument of the **lvconvert** command.

The following command converts the linear logical volume **my_lv** in volume group **my_vg** to a 2-way RAID1 array.

```
# lvconvert --type raid1 -m 1 my_vg/my_lv
```

Since RAID logical volumes are composed of metadata and data subvolume pairs, when you convert a linear device to a RAID1 array, a new metadata subvolume is created and associated with the original logical volume on (one of) the same physical volumes that the linear volume is on. The additional images are added in metadata/data subvolume pairs. For example, if the original device is as follows:

```
# lvs -a -o name,copy_percent,devices my_vg
LV      Copy%   Devices
my_lv           /dev/sde1(0)
```

After conversion to a 2-way RAID1 array the device contains the following data and metadata subvolume pairs:

```
# lvconvert --type raid1 -m 1 my_vg/my_lv
# lvs -a -o name,copy_percent,devices my_vg
LV      Copy%   Devices
my_lv    6.25   my_lv_rimage_0(0),my_lv_rimage_1(0)
[my_lv_rimage_0]    /dev/sde1(0)
[my_lv_rimage_1]    /dev/sdf1(1)
[my_lv_rmeta_0]     /dev/sde1(256)
[my_lv_rmeta_1]     /dev/sdf1(0)
```

If the metadata image that pairs with the original logical volume cannot be placed on the same physical volume, the **lvconvert** will fail.

4.4.3.3. Converting an LVM RAID1 Logical Volume to an LVM Linear Logical Volume

You can convert an existing RAID1 LVM logical volume to an LVM linear logical volume with the **lvconvert** command by specifying the **-m0** argument. This removes all the RAID data subvolumes and all the RAID metadata subvolumes that make up the RAID array, leaving the top-level RAID1 image as the linear logical volume.

The following example displays an existing LVM RAID1 logical volume.

```
# lvs -a -o name,copy_percent,devices my_vg
LV      Copy%   Devices
my_lv    100.00 my_lv_rimage_0(0),my_lv_rimage_1(0)
[my_lv_rimage_0]    /dev/sde1(1)
```

```
[my_lv_rimage_1]      /dev/sdf1(1)
[my_lv_rmeta_0]       /dev/sde1(0)
[my_lv_rmeta_1]       /dev/sdf1(0)
```

The following command converts the LVM RAID1 logical volume **my_vg/my_lv** to an LVM linear device.

```
# lvconvert -m0 my_vg/my_lv
# lvs -a -o name,copy_percent,devices my_vg
LV      Copy%  Devices
my_lv    100.00 /dev/sde1(1)
```

When you convert an LVM RAID1 logical volume to an LVM linear volume, you can specify which physical volumes to remove. The following example shows the layout of an LVM RAID1 logical volume made up of two images: **/dev/sda1** and **/dev/sdb1**. In this example, the **lvconvert** command specifies that you want to remove **/dev/sda1**, leaving **/dev/sdb1** as the physical volume that makes up the linear device.

```
# lvs -a -o name,copy_percent,devices my_vg
LV      Copy%  Devices
my_lv    100.00 my_lv_rimage_0(0),my_lv_rimage_1(0)
[my_lv_rimage_0]      /dev/sda1(1)
[my_lv_rimage_1]      /dev/sdb1(1)
[my_lv_rmeta_0]       /dev/sda1(0)
[my_lv_rmeta_1]       /dev/sdb1(0)
# lvconvert -m0 my_vg/my_lv /dev/sda1
# lvs -a -o name,copy_percent,devices my_vg
LV      Copy%  Devices
my_lv    100.00 /dev/sdb1(1)
```

4.4.3.4. Converting a Mirrored LVM Device to a RAID1 Device

You can convert an existing mirrored LVM device with a segment type of **mirror** to a RAID1 LVM device with the **lvconvert** command by specifying the **--type raid1** argument. This renames the mirror subvolumes (***_mimage_***) to RAID subvolumes (***_rimage_***). In addition, the mirror log is removed and metadata subvolumes (***_rmeta_***) are created for the data subvolumes on the same physical volumes as the corresponding data subvolumes.

The following example shows the layout of a mirrored logical volume **my_vg/my_lv**.

```
# lvs -a -o name,copy_percent,devices my_vg
LV      Copy%  Devices
my_lv    15.20 my_lv_mimage_0(0),my_lv_mimage_1(0)
[my_lv_mimage_0]      /dev/sde1(0)
[my_lv_mimage_1]      /dev/sdf1(0)
[my_lv_mlog]          /dev/sdd1(0)
```

The following command converts the mirrored logical volume **my_vg/my_lv** to a RAID1 logical volume.

```
# lvconvert --type raid1 my_vg/my_lv
# lvs -a -o name,copy_percent,devices my_vg
LV      Copy%  Devices
my_lv    100.00 my_lv_rimage_0(0),my_lv_rimage_1(0)
[my_lv_rimage_0]      /dev/sde1(0)
[my_lv_rimage_1]      /dev/sdf1(0)
```

```
[my_lv_rmeta_0]      /dev/sde1(125)
[my_lv_rmeta_1]      /dev/sdf1(125)
```

4.4.3.5. Resizing a RAID Logical Volume

You can resize a RAID logical volume in the following ways;

- You can increase the size of a RAID logical volume of any type with the **lvresize** or **lvextend** command. This does not change the number of RAID images. For striped RAID logical volumes the same stripe rounding constraints apply as when you create a striped RAID logical volume. For more information on extending a RAID volume, see [Section 4.4.18, “Extending a RAID Volume”](#).
- You can reduce the size of a RAID logical volume of any type with the **lvresize** or **lvreduce** command. This does not change the number of RAID images. As with the **lvextend** command, the same stripe rounding constraints apply as when you create a striped RAID logical volume. For an example of a command to reduce the size of a logical volume, see [Section 4.4.16, “Shrinking Logical Volumes”](#).
- As of Red Hat Enterprise Linux 7.4, you can change the number of stripes on a striped RAID logical volume (**raid4/5/6/10**) with the **--stripes N** parameter of the **lvconvert** command. This increases or reduces the size of the RAID logical volume by the capacity of the stripes added or removed. Note that **raid10** volumes are capable only of adding stripes. This capability is part of the RAID *reshaping* feature that allows you to change attributes of a RAID logical volume while keeping the same RAID level. For information on RAID reshaping and examples of using the **lvconvert** command to reshape a RAID logical volume, see the **lvraid(7)** man page.

4.4.3.6. Changing the Number of Images in an Existing RAID1 Device

You can change the number of images in an existing RAID1 array just as you can change the number of images in the earlier implementation of LVM mirroring. Use the **lvconvert** command to specify the number of additional metadata/data subvolume pairs to add or remove. For information on changing the volume configuration in the earlier implementation of LVM mirroring, see [Section 4.4.4.4, “Changing Mirrored Volume Configuration”](#).

When you add images to a RAID1 device with the **lvconvert** command, you can specify the total number of images for the resulting device, or you can specify how many images to add to the device. You can also optionally specify on which physical volumes the new metadata/data image pairs will reside.

Metadata subvolumes (named ***_rmeta_***) always exist on the same physical devices as their data subvolume counterparts ***_rimage_***. The metadata/data subvolume pairs will not be created on the same physical volumes as those from another metadata/data subvolume pair in the RAID array (unless you specify **--alloc anywhere**).

The format for the command to add images to a RAID1 volume is as follows:

```
lvconvert -m new_absolute_count vg/lv [removable_PVs]
lvconvert -m +num_additional_images vg/lv [removable_PVs]
```

For example, the following command displays the LVM device **my_vg/my_lv**, which is a 2-way RAID1 array:

```
# lvs -a -o name,copy_percent,devices my_vg
LV                Copy%   Devices
my_lv             6.25     my_lv_rimage_0(0),my_lv_rimage_1(0)
[my_lv_rimage_0]   /dev/sde1(0)
[my_lv_rimage_1]   /dev/sdf1(1)
[my_lv_rmeta_0]    /dev/sde1(256)
[my_lv_rmeta_1]    /dev/sdf1(0)
```

The following command converts the 2-way RAID1 device **my_vg/my_lv** to a 3-way RAID1 device:

```
# lvconvert -m 2 my_vg/my_lv
# lvs -a -o name,copy_percent,devices my_vg
LV                Copy%   Devices
my_lv             6.25
my_lv_rimage_0(0),my_lv_rimage_1(0),my_lv_rimage_2(0)
[my_lv_rimage_0]   /dev/sde1(0)
[my_lv_rimage_1]   /dev/sdf1(1)
[my_lv_rimage_2]   /dev/sdg1(1)
[my_lv_rmeta_0]    /dev/sde1(256)
[my_lv_rmeta_1]    /dev/sdf1(0)
[my_lv_rmeta_2]    /dev/sdg1(0)
```

When you add an image to a RAID1 array, you can specify which physical volumes to use for the image. The following command converts the 2-way RAID1 device **my_vg/my_lv** to a 3-way RAID1 device, specifying that the physical volume **/dev/sdd1** be used for the array:

```
# lvs -a -o name,copy_percent,devices my_vg
LV                Copy%   Devices
my_lv             56.00  my_lv_rimage_0(0),my_lv_rimage_1(0)
[my_lv_rimage_0]   /dev/sda1(1)
[my_lv_rimage_1]   /dev/sdb1(1)
[my_lv_rmeta_0]    /dev/sda1(0)
[my_lv_rmeta_1]    /dev/sdb1(0)
# lvconvert -m 2 my_vg/my_lv /dev/sdd1
# lvs -a -o name,copy_percent,devices my_vg
LV                Copy%   Devices
my_lv             28.00
my_lv_rimage_0(0),my_lv_rimage_1(0),my_lv_rimage_2(0)
[my_lv_rimage_0]   /dev/sda1(1)
[my_lv_rimage_1]   /dev/sdb1(1)
[my_lv_rimage_2]   /dev/sdd1(1)
[my_lv_rmeta_0]    /dev/sda1(0)
[my_lv_rmeta_1]    /dev/sdb1(0)
[my_lv_rmeta_2]    /dev/sdd1(0)
```

To remove images from a RAID1 array, use the following command. When you remove images from a RAID1 device with the **lvconvert** command, you can specify the total number of images for the resulting device, or you can specify how many images to remove from the device. You can also optionally specify the physical volumes from which to remove the device.

```
lvconvert -m new_absolute_count vg/lv [removable_PVs]
lvconvert -m -num_fewer_images vg/lv [removable_PVs]
```

Additionally, when an image and its associated metadata subvolume volume are removed, any higher-

numbered images will be shifted down to fill the slot. If you remove **lv_rimage_1** from a 3-way RAID1 array that consists of **lv_rimage_0**, **lv_rimage_1**, and **lv_rimage_2**, this results in a RAID1 array that consists of **lv_rimage_0** and **lv_rimage_1**. The subvolume **lv_rimage_2** will be renamed and take over the empty slot, becoming **lv_rimage_1**.

The following example shows the layout of a 3-way RAID1 logical volume **my_vg/my_lv**.

```
# lvs -a -o name,copy_percent,devices my_vg
LV              Copy%  Devices
my_lv           100.00
my_lv_rimage_0(0),my_lv_rimage_1(0),my_lv_rimage_2(0)
[my_lv_rimage_0]      /dev/sde1(1)
[my_lv_rimage_1]      /dev/sdf1(1)
[my_lv_rimage_2]      /dev/sdg1(1)
[my_lv_rmeta_0]       /dev/sde1(0)
[my_lv_rmeta_1]       /dev/sdf1(0)
[my_lv_rmeta_2]       /dev/sdg1(0)
```

The following command converts the 3-way RAID1 logical volume into a 2-way RAID1 logical volume.

```
# lvconvert -m1 my_vg/my_lv
# lvs -a -o name,copy_percent,devices my_vg
LV              Copy%  Devices
my_lv           100.00 my_lv_rimage_0(0),my_lv_rimage_1(0)
[my_lv_rimage_0]      /dev/sde1(1)
[my_lv_rimage_1]      /dev/sdf1(1)
[my_lv_rmeta_0]       /dev/sde1(0)
[my_lv_rmeta_1]       /dev/sdf1(0)
```

The following command converts the 3-way RAID1 logical volume into a 2-way RAID1 logical volume, specifying the physical volume that contains the image to remove as **/dev/sde1**.

```
# lvconvert -m1 my_vg/my_lv /dev/sde1
# lvs -a -o name,copy_percent,devices my_vg
LV              Copy%  Devices
my_lv           100.00 my_lv_rimage_0(0),my_lv_rimage_1(0)
[my_lv_rimage_0]      /dev/sdf1(1)
[my_lv_rimage_1]      /dev/sdg1(1)
[my_lv_rmeta_0]       /dev/sdf1(0)
[my_lv_rmeta_1]       /dev/sdg1(0)
```

4.4.3.7. Splitting off a RAID Image as a Separate Logical Volume

You can split off an image of a RAID logical volume to form a new logical volume. The procedure for splitting off a RAID image is the same as the procedure for splitting off a redundant image of a mirrored logical volume, as described in [Section 4.4.4.2, “Splitting Off a Redundant Image of a Mirrored Logical Volume”](#).

The format of the command to split off a RAID image is as follows:

```
lvconvert --splitmirrors count -n splitname vg/lv [removable_PVs]
```

Just as when you are removing a RAID image from an existing RAID1 logical volume (as described in [Section 4.4.3.6, “Changing the Number of Images in an Existing RAID1 Device”](#)), when you remove a

RAID data subvolume (and its associated metadata subvolume) from the middle of the device any higher numbered images will be shifted down to fill the slot. The index numbers on the logical volumes that make up a RAID array will thus be an unbroken sequence of integers.



NOTE

You cannot split off a RAID image if the RAID1 array is not yet in sync.

The following example splits a 2-way RAID1 logical volume, **my_lv**, into two linear logical volumes, **my_lv** and **new**.

```
# lvs -a -o name,copy_percent,devices my_vg
LV          Copy%   Devices
my_lv       12.00   my_lv_rimage_0(0),my_lv_rimage_1(0)
[my_lv_rimage_0]      /dev/sde1(1)
[my_lv_rimage_1]      /dev/sdf1(1)
[my_lv_rmeta_0]       /dev/sde1(0)
[my_lv_rmeta_1]       /dev/sdf1(0)
# lvconvert --splitmirror 1 -n new my_vg/my_lv
# lvs -a -o name,copy_percent,devices my_vg
LV          Copy%   Devices
my_lv       /dev/sde1(1)
new         /dev/sdf1(1)
```

The following example splits a 3-way RAID1 logical volume, **my_lv**, into a 2-way RAID1 logical volume, **my_lv**, and a linear logical volume, **new**

```
# lvs -a -o name,copy_percent,devices my_vg
LV          Copy%   Devices
my_lv       100.00
my_lv_rimage_0(0),my_lv_rimage_1(0),my_lv_rimage_2(0)
[my_lv_rimage_0]      /dev/sde1(1)
[my_lv_rimage_1]      /dev/sdf1(1)
[my_lv_rimage_2]      /dev/sdg1(1)
[my_lv_rmeta_0]       /dev/sde1(0)
[my_lv_rmeta_1]       /dev/sdf1(0)
[my_lv_rmeta_2]       /dev/sdg1(0)
# lvconvert --splitmirror 1 -n new my_vg/my_lv
# lvs -a -o name,copy_percent,devices my_vg
LV          Copy%   Devices
my_lv       100.00   my_lv_rimage_0(0),my_lv_rimage_1(0)
[my_lv_rimage_0]      /dev/sde1(1)
[my_lv_rimage_1]      /dev/sdf1(1)
[my_lv_rmeta_0]       /dev/sde1(0)
[my_lv_rmeta_1]       /dev/sdf1(0)
new         /dev/sdg1(1)
```

4.4.3.8. Splitting and Merging a RAID Image

You can temporarily split off an image of a RAID1 array for read-only use while keeping track of any changes by using the **--trackchanges** argument in conjunction with the **--splitmirrors** argument of the **lvconvert** command. This allows you to merge the image back into the array at a later time while resyncing only those portions of the array that have changed since the image was split.

The format for the **lvconvert** command to split off a RAID image is as follows.

```
lvconvert --splitmirrors count --trackchanges vg/lv [removable_PVs]
```

When you split off a RAID image with the **--trackchanges** argument, you can specify which image to split but you cannot change the name of the volume being split. In addition, the resulting volumes have the following constraints.

- The new volume you create is read-only.
- You cannot resize the new volume.
- You cannot rename the remaining array.
- You cannot resize the remaining array.
- You can activate the new volume and the remaining array independently.

You can merge an image that was split off with the **--trackchanges** argument specified by executing a subsequent **lvconvert** command with the **--merge** argument. When you merge the image, only the portions of the array that have changed since the image was split are resynced.

The format for the **lvconvert** command to merge a RAID image is as follows.

```
lvconvert --merge raid_image
```

The following example creates a RAID1 logical volume and then splits off an image from that volume while tracking changes to the remaining array.

```
# lvcreate --type raid1 -m 2 -L 1G -n my_lv .vg
Logical volume "my_lv" created
# lvs -a -o name,copy_percent,devices my_vg
LV          Copy%  Devices
my_lv       100.00
my_lv_rimage_0(0),my_lv_rimage_1(0),my_lv_rimage_2(0)
[my_lv_rimage_0]      /dev/sdb1(1)
[my_lv_rimage_1]      /dev/sdc1(1)
[my_lv_rimage_2]      /dev/sdd1(1)
[my_lv_rmeta_0]       /dev/sdb1(0)
[my_lv_rmeta_1]       /dev/sdc1(0)
[my_lv_rmeta_2]       /dev/sdd1(0)
# lvconvert --splitmirrors 1 --trackchanges my_vg/my_lv
my_lv_rimage_2 split from my_lv for read-only purposes.
Use 'lvconvert --merge my_vg/my_lv_rimage_2' to merge back into my_lv
# lvs -a -o name,copy_percent,devices my_vg
LV          Copy%  Devices
my_lv       100.00
my_lv_rimage_0(0),my_lv_rimage_1(0),my_lv_rimage_2(0)
[my_lv_rimage_0]      /dev/sdb1(1)
[my_lv_rimage_1]      /dev/sdc1(1)
my_lv_rimage_2        /dev/sdd1(1)
[my_lv_rmeta_0]       /dev/sdb1(0)
[my_lv_rmeta_1]       /dev/sdc1(0)
[my_lv_rmeta_2]       /dev/sdd1(0)
```

The following example splits off an image from a RAID1 volume while tracking changes to the remaining array, then merges the volume back into the array.

```
# lvconvert --splitmirrors 1 --trackchanges my_vg/my_lv
  lv_rimage_1 split from my_lv for read-only purposes.
  Use 'lvconvert --merge my_vg/my_lv_rimage_1' to merge back into my_lv
# lvs -a -o name,copy_percent,devices my_vg
LV          Copy%  Devices
my_lv       100.00  my_lv_rimage_0(0),my_lv_rimage_1(0)
[my_lv_rimage_0]    /dev/sdc1(1)
my_lv_rimage_1      /dev/sdd1(1)
[my_lv_rmeta_0]     /dev/sdc1(0)
[my_lv_rmeta_1]     /dev/sdd1(0)
# lvconvert --merge my_vg/my_lv_rimage_1
my_vg/my_lv_rimage_1 successfully merged back into my_vg/my_lv
# lvs -a -o name,copy_percent,devices my_vg
LV          Copy%  Devices
my_lv       100.00  my_lv_rimage_0(0),my_lv_rimage_1(0)
[my_lv_rimage_0]    /dev/sdc1(1)
[my_lv_rimage_1]    /dev/sdd1(1)
[my_lv_rmeta_0]     /dev/sdc1(0)
[my_lv_rmeta_1]     /dev/sdd1(0)
```

Once you have split off an image from a RAID1 volume, you can make the split permanent by issuing a second **lvconvert --splitmirrors** command, repeating the initial **lvconvert** command that split the image without specifying the **--trackchanges** argument. This breaks the link that the **--trackchanges** argument created.

After you have split an image with the **--trackchanges** argument, you cannot issue a subsequent **lvconvert --splitmirrors** command on that array unless your intent is to permanently split the image being tracked.

The following sequence of commands splits an image and tracks the image and then permanently splits off the image being tracked.

```
# lvconvert --splitmirrors 1 --trackchanges my_vg/my_lv
  my_lv_rimage_1 split from my_lv for read-only purposes.
  Use 'lvconvert --merge my_vg/my_lv_rimage_1' to merge back into my_lv
# lvconvert --splitmirrors 1 -n new my_vg/my_lv
# lvs -a -o name,copy_percent,devices my_vg
LV  Copy%  Devices
my_lv      /dev/sdc1(1)
new        /dev/sdd1(1)
```

Note, however, that the following sequence of commands will fail.

```
# lvconvert --splitmirrors 1 --trackchanges my_vg/my_lv
  my_lv_rimage_1 split from my_lv for read-only purposes.
  Use 'lvconvert --merge my_vg/my_lv_rimage_1' to merge back into my_lv
# lvconvert --splitmirrors 1 --trackchanges my_vg/my_lv
  Cannot track more than one split image at a time
```

Similarly, the following sequence of commands will fail as well, since the split image is not the image being tracked.

■

```
# lvconvert --splitmirrors 1 --trackchanges my_vg/my_lv
my_lv_rimage_1 split from my_lv for read-only purposes.
Use 'lvconvert --merge my_vg/my_lv_rimage_1' to merge back into my_lv
# lvs -a -o name,copy_percent,devices my_vg
LV                               Copy%   Devices
my_lv                            100.00  my_lv_rimage_0(0),my_lv_rimage_1(0)
[my_lv_rimage_0]                  /dev/sdc1(1)
my_lv_rimage_1                   /dev/sdd1(1)
[my_lv_rmeta_0]                   /dev/sdc1(0)
[my_lv_rmeta_1]                   /dev/sdd1(0)
# lvconvert --splitmirrors 1 -n new my_vg/my_lv /dev/sdc1
Unable to split additional image from my_lv while tracking changes for
my_lv_rimage_1
```

4.4.3.9. Setting a RAID fault policy

LVM RAID handles device failures in an automatic fashion based on the preferences defined by the **raid_fault_policy** field in the **lvm.conf** file.

- If the **raid_fault_policy** field is set to **allocate**, the system will attempt to replace the failed device with a spare device from the volume group. If there is no available spare device, this will be reported to the system log.
- If the **raid_fault_policy** field is set to **warn**, the system will produce a warning and the log will indicate that a device has failed. This allows the user to determine the course of action to take.

As long as there are enough devices remaining to support usability, the RAID logical volume will continue to operate.

4.4.3.9.1. The allocate RAID Fault Policy

In the following example, the **raid_fault_policy** field has been set to **allocate** in the **lvm.conf** file. The RAID logical volume is laid out as follows.

```
# lvs -a -o name,copy_percent,devices my_vg
LV                               Copy%   Devices
my_lv                            100.00
my_lv_rimage_0(0),my_lv_rimage_1(0),my_lv_rimage_2(0)
[my_lv_rimage_0]                  /dev/sde1(1)
[my_lv_rimage_1]                  /dev/sdf1(1)
[my_lv_rimage_2]                  /dev/sdg1(1)
[my_lv_rmeta_0]                   /dev/sde1(0)
[my_lv_rmeta_1]                   /dev/sdf1(0)
[my_lv_rmeta_2]                   /dev/sdg1(0)
```

If the **/dev/sde** device fails, the system log will display error messages.

```
# grep lvm /var/log/messages
Jan 17 15:57:18 bp-01 lvm[8599]: Device #0 of raid1 array, my_vg-my_lv,
has failed.
Jan 17 15:57:18 bp-01 lvm[8599]: /dev/sde1: read failed after 0 of 2048 at
250994294784: Input/output error
Jan 17 15:57:18 bp-01 lvm[8599]: /dev/sde1: read failed after 0 of 2048 at
250994376704: Input/output error
```

```

Jan 17 15:57:18 bp-01 lvm[8599]: /dev/sde1: read failed after 0 of 2048 at
0:
Input/output error
Jan 17 15:57:18 bp-01 lvm[8599]: /dev/sde1: read failed after 0 of 2048 at
4096: Input/output error
Jan 17 15:57:19 bp-01 lvm[8599]: Couldn't find device with uuid
3lugiV-3eSP-AFAR-sdrP-H200-wM2M-qdMANY.
Jan 17 15:57:27 bp-01 lvm[8599]: raid1 array, my_vg-my_lv, is not in-sync.
Jan 17 15:57:36 bp-01 lvm[8599]: raid1 array, my_vg-my_lv, is now in-sync.

```

Since the **raid_fault_policy** field has been set to **allocate**, the failed device is replaced with a new device from the volume group.

```

# lvs -a -o name,copy_percent,devices vg
Couldn't find device with uuid 3lugiV-3eSP-AFAR-sdrP-H200-wM2M-qdMANY.
LV          Copy%   Devices
lv          100.00  lv_rimage_0(0),lv_rimage_1(0),lv_rimage_2(0)
[lv_rimage_0]          /dev/sdh1(1)
[lv_rimage_1]          /dev/sdf1(1)
[lv_rimage_2]          /dev/sdg1(1)
[lv_rmeta_0]           /dev/sdh1(0)
[lv_rmeta_1]           /dev/sdf1(0)
[lv_rmeta_2]           /dev/sdg1(0)

```

Note that even though the failed device has been replaced, the display still indicates that LVM could not find the failed device. This is because, although the failed device has been removed from the RAID logical volume, the failed device has not yet been removed from the volume group. To remove the failed device from the volume group, you can execute **vgreduce --removemissing VG**.

If the **raid_fault_policy** has been set to **allocate** but there are no spare devices, the allocation will fail, leaving the logical volume as it is. If the allocation fails, you have the option of fixing the drive, then deactivating and activating the logical volume; this is described in [Section 4.4.3.9.2, “The warn RAID Fault Policy”](#). Alternately, you can replace the failed device, as described in [Section 4.4.3.10, “Replacing a RAID device”](#).

4.4.3.9.2. The warn RAID Fault Policy

In the following example, the **raid_fault_policy** field has been set to **warn** in the **lvm.conf** file. The RAID logical volume is laid out as follows.

```

# lvs -a -o name,copy_percent,devices my_vg
LV          Copy%   Devices
my_lv       100.00
my_lv_rimage_0(0),my_lv_rimage_1(0),my_lv_rimage_2(0)
[my_lv_rimage_0]          /dev/sdh1(1)
[my_lv_rimage_1]          /dev/sdf1(1)
[my_lv_rimage_2]          /dev/sdg1(1)
[my_lv_rmeta_0]           /dev/sdh1(0)
[my_lv_rmeta_1]           /dev/sdf1(0)
[my_lv_rmeta_2]           /dev/sdg1(0)

```

If the **/dev/sdh** device fails, the system log will display error messages. In this case, however, LVM will not automatically attempt to repair the RAID device by replacing one of the images. Instead, if the device has failed you can replace the device with the **--repair** argument of the **lvconvert** command, as shown below.

```
# lvconvert --repair my_vg/my_lv
/dev/sdh1: read failed after 0 of 2048 at 250994294784: Input/output
error
/dev/sdh1: read failed after 0 of 2048 at 250994376704: Input/output
error
/dev/sdh1: read failed after 0 of 2048 at 0: Input/output error
/dev/sdh1: read failed after 0 of 2048 at 4096: Input/output error
Couldn't find device with uuid fbI0Y0-GX7x-firU-Vy5o-vzwx-vAKZ-feRxFF.
Attempt to replace failed RAID images (requires full device resync)?
[y/n]: y

# lvs -a -o name,copy_percent,devices my_vg
Couldn't find device with uuid fbI0Y0-GX7x-firU-Vy5o-vzwx-vAKZ-feRxFF.
LV              Copy%  Devices
my_lv           64.00
my_lv_rimage_0(0),my_lv_rimage_1(0),my_lv_rimage_2(0)
[my_lv_rimage_0]      /dev/sde1(1)
[my_lv_rimage_1]      /dev/sdf1(1)
[my_lv_rimage_2]      /dev/sdg1(1)
[my_lv_rmeta_0]       /dev/sde1(0)
[my_lv_rmeta_1]       /dev/sdf1(0)
[my_lv_rmeta_2]       /dev/sdg1(0)
```

Note that even though the failed device has been replaced, the display still indicates that LVM could not find the failed device. This is because, although the failed device has been removed from the RAID logical volume, the failed device has not yet been removed from the volume group. To remove the failed device from the volume group, you can execute **vgreduce --removemissing VG**.

If the device failure is a transient failure or you are able to repair the device that failed, you can initiate recovery of the failed device with the **--refresh** option of the **lvchange** command. Previously it was necessary to deactivate and then activate the logical volume.

The following command refreshes a logical volume.

```
# lvchange --refresh my_vg/my_lv
```

4.4.3.10. Replacing a RAID device

RAID is not like traditional LVM mirroring. LVM mirroring required failed devices to be removed or the mirrored logical volume would hang. RAID arrays can keep on running with failed devices. In fact, for RAID types other than RAID1, removing a device would mean converting to a lower level RAID (for example, from RAID6 to RAID5, or from RAID4 or RAID5 to RAID0). Therefore, rather than removing a failed device unconditionally and potentially allocating a replacement, LVM allows you to replace a device in a RAID volume in a one-step solution by using the **--replace** argument of the **lvconvert** command.

The format for the **lvconvert --replace** is as follows.

```
lvconvert --replace dev_to_remove vg/lv [possible_replacements]
```

The following example creates a RAID1 logical volume and then replaces a device in that volume.

```
# lvcreate --type raid1 -m 2 -L 1G -n my_lv my_vg
Logical volume "my_lv" created
```

```
# lvs -a -o name,copy_percent,devices my_vg
LV                Copy%   Devices
my_lv             100.00
my_lv_rimage_0(0),my_lv_rimage_1(0),my_lv_rimage_2(0)
[my_lv_rimage_0]   /dev/sdb1(1)
[my_lv_rimage_1]   /dev/sdb2(1)
[my_lv_rimage_2]   /dev/sdc1(1)
[my_lv_rmeta_0]    /dev/sdb1(0)
[my_lv_rmeta_1]    /dev/sdb2(0)
[my_lv_rmeta_2]    /dev/sdc1(0)
# lvconvert --replace /dev/sdb2 my_vg/my_lv
# lvs -a -o name,copy_percent,devices my_vg
LV                Copy%   Devices
my_lv             37.50
my_lv_rimage_0(0),my_lv_rimage_1(0),my_lv_rimage_2(0)
[my_lv_rimage_0]   /dev/sdb1(1)
[my_lv_rimage_1]   /dev/sdc2(1)
[my_lv_rimage_2]   /dev/sdc1(1)
[my_lv_rmeta_0]    /dev/sdb1(0)
[my_lv_rmeta_1]    /dev/sdc2(0)
[my_lv_rmeta_2]    /dev/sdc1(0)
```

The following example creates a RAID1 logical volume and then replaces a device in that volume, specifying which physical volume to use for the replacement.

```
# lvcreate --type raid1 -m 1 -L 100 -n my_lv my_vg
Logical volume "my_lv" created
# lvs -a -o name,copy_percent,devices my_vg
LV                Copy%   Devices
my_lv             100.00 my_lv_rimage_0(0),my_lv_rimage_1(0)
[my_lv_rimage_0]   /dev/sda1(1)
[my_lv_rimage_1]   /dev/sdb1(1)
[my_lv_rmeta_0]    /dev/sda1(0)
[my_lv_rmeta_1]    /dev/sdb1(0)
# pvs
PV                VG          Fmt  Attr PSize    PFree
/dev/sda1         my_vg      lvm2 a--  1020.00m  916.00m
/dev/sdb1         my_vg      lvm2 a--  1020.00m  916.00m
/dev/sdc1         my_vg      lvm2 a--  1020.00m 1020.00m
/dev/sdd1         my_vg      lvm2 a--  1020.00m 1020.00m
# lvconvert --replace /dev/sdb1 my_vg/my_lv /dev/sdd1
# lvs -a -o name,copy_percent,devices my_vg
LV                Copy%   Devices
my_lv             28.00 my_lv_rimage_0(0),my_lv_rimage_1(0)
[my_lv_rimage_0]   /dev/sda1(1)
[my_lv_rimage_1]   /dev/sdd1(1)
[my_lv_rmeta_0]    /dev/sda1(0)
[my_lv_rmeta_1]    /dev/sdd1(0)
```

You can replace more than one RAID device at a time by specifying multiple **replace** arguments, as in the following example.

```
# lvcreate --type raid1 -m 2 -L 100 -n my_lv my_vg
Logical volume "my_lv" created
# lvs -a -o name,copy_percent,devices my_vg
```

```

LV              Copy%  Devices
my_lv           100.00
my_lv_rimage_0(0),my_lv_rimage_1(0),my_lv_rimage_2(0)
[my_lv_rimage_0]      /dev/sda1(1)
[my_lv_rimage_1]      /dev/sdb1(1)
[my_lv_rimage_2]      /dev/sdc1(1)
[my_lv_rmeta_0]       /dev/sda1(0)
[my_lv_rmeta_1]       /dev/sdb1(0)
[my_lv_rmeta_2]       /dev/sdc1(0)
# lvconvert --replace /dev/sdb1 --replace /dev/sdc1 my_vg/my_lv
# lvs -a -o name,copy_percent,devices my_vg
LV              Copy%  Devices
my_lv           60.00
my_lv_rimage_0(0),my_lv_rimage_1(0),my_lv_rimage_2(0)
[my_lv_rimage_0]      /dev/sda1(1)
[my_lv_rimage_1]      /dev/sdd1(1)
[my_lv_rimage_2]      /dev/sde1(1)
[my_lv_rmeta_0]       /dev/sda1(0)
[my_lv_rmeta_1]       /dev/sdd1(0)
[my_lv_rmeta_2]       /dev/sde1(0)

```



NOTE

When you specify a replacement drive using the **lvconvert --replace** command, the replacement drives should never be allocated from extra space on drives already used in the array. For example, **lv_rimage_0** and **lv_rimage_1** should not be located on the same physical volume.

4.4.3.11. Scrubbing a RAID Logical Volume

LVM provides scrubbing support for RAID logical volumes. RAID scrubbing is the process of reading all the data and parity blocks in an array and checking to see whether they are coherent.

You initiate a RAID scrubbing operation with the **--syncaction** option of the **lvchange** command. You specify either a **check** or **repair** operation. A **check** operation goes over the array and records the number of discrepancies in the array but does not repair them. A **repair** operation corrects the discrepancies as it finds them.

The format of the command to scrub a RAID logical volume is as follows:

```
lvchange --syncaction {check|repair} vg/raid_lv
```



NOTE

The **lvchange --syncaction repair vg/raid_lv** operation does not perform the same function as the **lvconvert --repair vg/raid_lv** operation. The **lvchange --syncaction repair** operation initiates a background synchronization operation on the array, while the **lvconvert --repair** operation is designed to repair/replace failed devices in a mirror or RAID logical volume.

In support of the new RAID scrubbing operation, the **lvs** command now supports two new printable fields: **raid_sync_action** and **raid_mismatch_count**. These fields are not printed by default. To display these fields you specify them with the **-o** parameter of the **lvs**, as follows.

```
lvs -o +raid_sync_action,raid_mismatch_count vg/lv
```

The **raid_sync_action** field displays the current synchronization operation that the raid volume is performing. It can be one of the following values:

- **idle**: All sync operations complete (doing nothing)
- **resync**: Initializing an array or recovering after a machine failure
- **recover**: Replacing a device in the array
- **check**: Looking for array inconsistencies
- **repair**: Looking for and repairing inconsistencies

The **raid_mismatch_count** field displays the number of discrepancies found during a **check** operation.

The **Cpy%Sync** field of the **lvs** command now prints the progress of any of the **raid_sync_action** operations, including **check** and **repair**.

The **lv_attr** field of the **lvs** command output now provides additional indicators in support of the RAID scrubbing operation. Bit 9 of this field displays the health of the logical volume, and it now supports the following indicators.

- (*m*)ismatches indicates that there are discrepancies in a RAID logical volume. This character is shown after a scrubbing operation has detected that portions of the RAID are not coherent.
- (*r*)efresh indicates that a device in a RAID array has suffered a failure and the kernel regards it as failed, even though LVM can read the device label and considers the device to be operational. The logical volume should be (*r*)efreshed to notify the kernel that the device is now available, or the device should be (*r*)eplaced if it is suspected of having failed.

For information on the **lvs** command, see [Section 4.8.2, “Object Display Fields”](#).

When you perform a RAID scrubbing operation, the background I/O required by the **sync** operations can crowd out other I/O operations to LVM devices, such as updates to volume group metadata. This can cause the other LVM operations to slow down. You can control the rate at which the RAID logical volume is scrubbed by implementing recovery throttling.

You control the rate at which **sync** operations are performed by setting the minimum and maximum I/O rate for those operations with the **--minrecoveryrate** and **--maxrecoveryrate** options of the **lvchange** command. You specify these options as follows.

- **--maxrecoveryrate** *Rate*[bBsSkKmMgG]

Sets the maximum recovery rate for a RAID logical volume so that it will not crowd out nominal I/O operations. The *Rate* is specified as an amount per second for each device in the array. If no suffix is given, then kiB/sec/device is assumed. Setting the recovery rate to 0 means it will be unbounded.

- **--minrecoveryrate** *Rate*[bBsSkKmMgG]

Sets the minimum recovery rate for a RAID logical volume to ensure that I/O for **sync** operations achieves a minimum throughput, even when heavy nominal I/O is present. The *Rate* is specified as an amount per second for each device in the array. If no suffix is given, then kiB/sec/device is

assumed.

4.4.3.12. RAID Takeover (Red Hat Enterprise Linux 7.4 and Later)

LVM supports *Raid takeover*, which means converting a RAID logical volume from one RAID level to another (such as from RAID 5 to RAID 6). Changing the RAID level is usually done to increase or decrease resilience to device failures or to restripe logical volumes. You use the **lvconvert** for RAID takeover. For information on RAID takeover and for examples of using the **lvconvert** to convert a RAID logical volume, see the **lvraid(7)** man page.

4.4.3.13. Reshaping a RAID Logical Volume (Red Hat Enterprise Linux 7.4 and Later)

RAID *reshaping* means changing attributes of a RAID logical volume while keeping the same RAID level. Some attributes you can change include RAID layout, stripe size, and number of stripes. For information on RAID reshaping and examples of using the **lvconvert** command to reshape a RAID logical volume, see the **lvraid(7)** man page.

4.4.3.14. Controlling I/O Operations on a RAID1 Logical Volume

You can control the I/O operations for a device in a RAID1 logical volume by using the **--writemostly** and **--writebehind** parameters of the **lvchange** command. The format for using these parameters is as follows.

- **--[raid]writemostly *PhysicalVolume*[:{t|y|n}]**

Marks a device in a RAID1 logical volume as **write-mostly**. All reads to these drives will be avoided unless necessary. Setting this parameter keeps the number of I/O operations to the drive to a minimum. By default, the **write-mostly** attribute is set to yes for the specified physical volume in the logical volume. It is possible to remove the **write-mostly** flag by appending **:n** to the physical volume or to toggle the value by specifying **:t**. The **--writemostly** argument can be specified more than one time in a single command, making it possible to toggle the write-mostly attributes for all the physical volumes in a logical volume at once.

- **--[raid]writebehind *IOCount***

Specifies the maximum number of outstanding writes that are allowed to devices in a RAID1 logical volume that are marked as **write-mostly**. Once this value is exceeded, writes become synchronous, causing all writes to the constituent devices to complete before the array signals the write has completed. Setting the value to zero clears the preference and allows the system to choose the value arbitrarily.

4.4.3.15. Changing the region size on a RAID Logical Volume (Red Hat Enterprise Linux 7.4 and later)

When you create a RAID logical volume, the region size for the logical volume will be the value of the **raid_region_size** parameter in the **/etc/lvm/lvm.conf** file. You can override this default value with the **-R** option of the **lvcreate** command.

After you have created a RAID logical volume, you can change the region size of the volume with the **-R** option of the **lvconvert** command. The following example changes the region size of logical volume **vg/raidlv** to 4096K. The RAID volume must be synced in order to change the region size.

```
# lvconvert -R 4096K vg/raid1
```

```
Do you really want to change the region_size 512.00 KiB of LV vg/raid1 to
4.00 MiB? [y/n]: y
Changed region size on RAID LV vg/raid1 to 4.00 MiB.
```

4.4.4. Creating Mirrored Volumes

For the Red Hat Enterprise Linux 7.0 release, LVM supports RAID 1/4/5/6/10, as described in [Section 4.4.3, “RAID Logical Volumes”](#). RAID logical volumes are not cluster-aware. While RAID logical volumes can be created and activated exclusively on one machine, they cannot be activated simultaneously on more than one machine. If you require non-exclusive mirrored volumes, you must create the volumes with a **mirror** segment type, as described in this section.



NOTE

For information on converting an existing LVM device with a segment type of **mirror** to a RAID1 LVM device, see [Section 4.4.3.4, “Converting a Mirrored LVM Device to a RAID1 Device”](#).



NOTE

Creating a mirrored LVM logical volume in a cluster requires the same commands and procedures as creating a mirrored LVM logical volume with a segment type of **mirror** on a single node. However, in order to create a mirrored LVM volume in a cluster, the cluster and cluster mirror infrastructure must be running, the cluster must be quorate, and the locking type in the **lvm.conf** file must be set correctly to enable cluster locking. For an example of creating a mirrored volume in a cluster, see [Section 5.5, “Creating a Mirrored LVM Logical Volume in a Cluster”](#).

Attempting to run multiple LVM mirror creation and conversion commands in quick succession from multiple nodes in a cluster might cause a backlog of these commands. This might cause some of the requested operations to time out and, subsequently, fail. To avoid this issue, it is recommended that cluster mirror creation commands be executed from one node of the cluster.

When you create a mirrored volume, you specify the number of copies of the data to make with the **-m** argument of the **lvcreate** command. Specifying **-m1** creates one mirror, which yields two copies of the file system: a linear logical volume plus one copy. Similarly, specifying **-m2** creates two mirrors, yielding three copies of the file system.

The following command creates a mirrored logical volume with a single mirror. The volume is 50 gigabytes in size, is named **mirrorlv**, and is carved out of volume group **vg0**:

```
# lvcreate --type mirror -L 50G -m 1 -n mirrorlv vg0
```

An LVM mirror divides the device being copied into regions that, by default, are 512KB in size. You can use the **-R** argument of the **lvcreate** command to specify the region size in megabytes. You can also change the default region size by editing the **mirror_region_size** setting in the **lvm.conf** file.

NOTE

Due to limitations in the cluster infrastructure, cluster mirrors greater than 1.5TB cannot be created with the default region size of 512KB. Users that require larger mirrors should increase the region size from its default to something larger. Failure to increase the region size will cause LVM creation to hang and may hang other LVM commands as well.

As a general guideline for specifying the region size for mirrors that are larger than 1.5TB, you could take your mirror size in terabytes and round up that number to the next power of 2, using that number as the **-R** argument to the **lvcreate** command. For example, if your mirror size is 1.5TB, you could specify **-R 2**. If your mirror size is 3TB, you could specify **-R 4**. For a mirror size of 5TB, you could specify **-R 8**.

The following command creates a mirrored logical volume with a region size of 2MB:

```
# lvcreate --type mirror -m 1 -L 2T -R 2 -n mirror vol_group
```

When a mirror is created, the mirror regions are synchronized. For large mirror components, the sync process may take a long time. When you are creating a new mirror that does not need to be revived, you can specify the **--nosync** argument to indicate that an initial synchronization from the first device is not required.

LVM maintains a small log which it uses to keep track of which regions are in sync with the mirror or mirrors. By default, this log is kept on disk, which keeps it persistent across reboots and ensures that the mirror does not need to be re-synced every time a machine reboots or crashes. You can specify instead that this log be kept in memory with the **--mirrorlog core** argument; this eliminates the need for an extra log device, but it requires that the entire mirror be resynchronized at every reboot.

The following command creates a mirrored logical volume from the volume group **bigvg**. The logical volume is named **ondiskmirvol** and has a single mirror. The volume is 12MB in size and keeps the mirror log in memory.

```
# lvcreate --type mirror -L 12MB -m 1 --mirrorlog core -n ondiskmirvol
bigvg
Logical volume "ondiskmirvol" created
```

The mirror log is created on a separate device from the devices on which any of the mirror legs are created. It is possible, however, to create the mirror log on the same device as one of the mirror legs by using the **--alloc anywhere** argument of the **vgcreate** command. This may degrade performance, but it allows you to create a mirror even if you have only two underlying devices.

The following command creates a mirrored logical volume with a single mirror for which the mirror log is on the same device as one of the mirror legs. In this example, the volume group **vg0** consists of only two devices. This command creates a 500 MB volume named **mirrorlv** in the **vg0** volume group.

```
# lvcreate --type mirror -L 500M -m 1 -n mirrorlv -alloc anywhere vg0
```



NOTE

With clustered mirrors, the mirror log management is completely the responsibility of the cluster node with the currently lowest cluster ID. Therefore, when the device holding the cluster mirror log becomes unavailable on a subset of the cluster, the clustered mirror can continue operating without any impact, as long as the cluster node with lowest ID retains access to the mirror log. Since the mirror is undisturbed, no automatic corrective action (repair) is issued, either. When the lowest-ID cluster node loses access to the mirror log, however, automatic action will kick in (regardless of accessibility of the log from other nodes).

To create a mirror log that is itself mirrored, you can specify the **--mirrorlog mirrored** argument. The following command creates a mirrored logical volume from the volume group **bigvg**. The logical volume is named **twologvol** and has a single mirror. The volume is 12MB in size and the mirror log is mirrored, with each log kept on a separate device.

```
# lvcreate --type mirror -L 12MB -m 1 --mirrorlog mirrored -n twologvol
bigvg
Logical volume "twologvol" created
```

Just as with a standard mirror log, it is possible to create the redundant mirror logs on the same device as the mirror legs by using the **--alloc anywhere** argument of the **vgcreate** command. This may degrade performance, but it allows you to create a redundant mirror log even if you do not have sufficient underlying devices for each log to be kept on a separate device than the mirror legs.

When a mirror is created, the mirror regions are synchronized. For large mirror components, the sync process may take a long time. When you are creating a new mirror that does not need to be revived, you can specify the **--nosync** argument to indicate that an initial synchronization from the first device is not required.

You can specify which devices to use for the mirror legs and log, and which extents of the devices to use. To force the log onto a particular disk, specify exactly one extent on the disk on which it will be placed. LVM does not necessarily respect the order in which devices are listed in the command line. If any physical volumes are listed that is the only space on which allocation will take place. Any physical extents included in the list that are already allocated will get ignored.

The following command creates a mirrored logical volume with a single mirror and a single log that is not mirrored. The volume is 500 MB in size, it is named **mirrorlv**, and it is carved out of volume group **vg0**. The first leg of the mirror is on device **/dev/sda1**, the second leg of the mirror is on device **/dev/sdb1**, and the mirror log is on **/dev/sdc1**.

```
# lvcreate --type mirror -L 500M -m 1 -n mirrorlv vg0 /dev/sda1 /dev/sdb1
/dev/sdc1
```

The following command creates a mirrored logical volume with a single mirror. The volume is 500 MB in size, it is named **mirrorlv**, and it is carved out of volume group **vg0**. The first leg of the mirror is on extents 0 through 499 of device **/dev/sda1**, the second leg of the mirror is on extents 0 through 499 of device **/dev/sdb1**, and the mirror log starts on extent 0 of device **/dev/sdc1**. These are 1MB extents. If any of the specified extents have already been allocated, they will be ignored.

```
# lvcreate --type mirror -L 500M -m 1 -n mirrorlv vg0 /dev/sda1:0-499
/dev/sdb1:0-499 /dev/sdc1:0
```

**NOTE**

You can combine striping and mirroring in a single logical volume. Creating a logical volume while simultaneously specifying the number of mirrors (**--mirrors X**) and the number of stripes (**--stripes Y**) results in a mirror device whose constituent devices are striped.

4.4.4.1. Mirrored Logical Volume Failure Policy

You can define how a mirrored logical volume behaves in the event of a device failure with the **mirror_image_fault_policy** and **mirror_log_fault_policy** parameters in the **activation** section of the **lvm.conf** file. When these parameters are set to **remove**, the system attempts to remove the faulty device and run without it. When these parameters are set to **allocate**, the system attempts to remove the faulty device and tries to allocate space on a new device to be a replacement for the failed device. This policy acts like the **remove** policy if no suitable device and space can be allocated for the replacement.

By default, the **mirror_log_fault_policy** parameter is set to **allocate**. Using this policy for the log is fast and maintains the ability to remember the sync state through crashes and reboots. If you set this policy to **remove**, when a log device fails the mirror converts to using an in-memory log; in this instance, the mirror will not remember its sync status across crashes and reboots and the entire mirror will be re-synced.

By default, the **mirror_image_fault_policy** parameter is set to **remove**. With this policy, if a mirror image fails the mirror will convert to a non-mirrored device if there is only one remaining good copy. Setting this policy to **allocate** for a mirror device requires the mirror to resynchronize the devices; this is a slow process, but it preserves the mirror characteristic of the device.

**NOTE**

When an LVM mirror suffers a device failure, a two-stage recovery takes place. The first stage involves removing the failed devices. This can result in the mirror being reduced to a linear device. The second stage, if the **mirror_log_fault_policy** parameter is set to **allocate**, is to attempt to replace any of the failed devices. Note, however, that there is no guarantee that the second stage will choose devices previously in-use by the mirror that had not been part of the failure if others are available.

For information on manually recovering from an LVM mirror failure, see [Section 6.3, “Recovering from LVM Mirror Failure”](#).

4.4.4.2. Splitting Off a Redundant Image of a Mirrored Logical Volume

You can split off a redundant image of a mirrored logical volume to form a new logical volume. To split off an image, use the **--splitmirrors** argument of the **lvconvert** command, specifying the number of redundant images to split off. You must use the **--name** argument of the command to specify a name for the newly-split-off logical volume.

The following command splits off a new logical volume named **copy** from the mirrored logical volume **vg/lv**. The new logical volume contains two mirror legs. In this example, LVM selects which devices to split off.

```
# lvconvert --splitmirrors 2 --name copy vg/lv
```

You can specify which devices to split off. The following command splits off a new logical volume named **copy** from the mirrored logical volume **vg/lv**. The new logical volume contains two mirror legs consisting of devices **/dev/sdc1** and **/dev/sde1**.

```
# lvconvert --splitmirrors 2 --name copy vg/lv /dev/sd[ce]1
```

4.4.4.3. Repairing a Mirrored Logical Device

You can use the **lvconvert --repair** command to repair a mirror after a disk failure. This brings the mirror back into a consistent state. The **lvconvert --repair** command is an interactive command that prompts you to indicate whether you want the system to attempt to replace any failed devices.

- To skip the prompts and replace all of the failed devices, specify the **-y** option on the command line.
- To skip the prompts and replace none of the failed devices, specify the **-f** option on the command line.
- To skip the prompts and still indicate different replacement policies for the mirror image and the mirror log, you can specify the **--use-policies** argument to use the device replacement policies specified by the **mirror_log_fault_policy** and **mirror_device_fault_policy** parameters in the **lvm.conf** file.

4.4.4.4. Changing Mirrored Volume Configuration

You can increase or decrease the number of mirrors that a logical volume contains by using the **lvconvert** command. This allows you to convert a logical volume from a mirrored volume to a linear volume or from a linear volume to a mirrored volume. You can also use this command to reconfigure other mirror parameters of an existing logical volume, such as **corelog**.

When you convert a linear volume to a mirrored volume, you are creating mirror legs for an existing volume. This means that your volume group must contain the devices and space for the mirror legs and for the mirror log.

If you lose a leg of a mirror, LVM converts the volume to a linear volume so that you still have access to the volume, without the mirror redundancy. After you replace the leg, use the **lvconvert** command to restore the mirror. This procedure is provided in [Section 6.3, “Recovering from LVM Mirror Failure”](#).

The following command converts the linear logical volume **vg00/lvol1** to a mirrored logical volume.

```
# lvconvert -m1 vg00/lvol1
```

The following command converts the mirrored logical volume **vg00/lvol1** to a linear logical volume, removing the mirror leg.

```
# lvconvert -m0 vg00/lvol1
```

The following example adds an additional mirror leg to the existing logical volume **vg00/lvol1**. This example shows the configuration of the volume before and after the **lvconvert** command changed the volume to a volume with two mirror legs.

```
# lvs -a -o name,copy_percent,devices vg00
LV                               Copy%   Devices
```

```

lvvol1          100.00 lvvol1_mimage_0(0),lvvol1_mimage_1(0)
[lvvol1_mimage_0]      /dev/sda1(0)
[lvvol1_mimage_1]      /dev/sdb1(0)
[lvvol1_mlog]          /dev/sdd1(0)
# lvconvert -m 2 vg00/lvvol1
vg00/lvvol1: Converted: 13.0%
vg00/lvvol1: Converted: 100.0%
Logical volume lvvol1 converted.
# lvs -a -o name,copy_percent,devices vg00
LV              Copy%   Devices
lvvol1          100.00
lvvol1_mimage_0(0),lvvol1_mimage_1(0),lvvol1_mimage_2(0)
[lvvol1_mimage_0]      /dev/sda1(0)
[lvvol1_mimage_1]      /dev/sdb1(0)
[lvvol1_mimage_2]      /dev/sdc1(0)
[lvvol1_mlog]          /dev/sdd1(0)

```

4.4.5. Creating Thinly-Provisioned Logical Volumes

Logical volumes can be thinly provisioned. This allows you to create logical volumes that are larger than the available extents. Using thin provisioning, you can manage a storage pool of free space, known as a thin pool, which can be allocated to an arbitrary number of devices when needed by applications. You can then create devices that can be bound to the thin pool for later allocation when an application actually writes to the logical volume. The thin pool can be expanded dynamically when needed for cost-effective allocation of storage space.



NOTE

This section provides an overview of the basic commands you use to create and grow thinly-provisioned logical volumes. For detailed information on LVM thin provisioning as well as information on using the LVM commands and utilities with thinly-provisioned logical volumes, see the **lvmtthin(7)** man page.



NOTE

Thin volumes are not supported across the nodes in a cluster. The thin pool and all its thin volumes must be exclusively activated on only one cluster node.

To create a thin volume, perform the following tasks:

1. Create a volume group with the **vgcreate** command.
2. Create a thin pool with the **lvcreate** command.
3. Create a thin volume in the thin pool with the **lvcreate** command.

You can use the **-T** (or **--thin**) option of the **lvcreate** command to create either a thin pool or a thin volume. You can also use **-T** option of the **lvcreate** command to create both a thin pool and a thin volume in that pool at the same time with a single command.

The following command uses the **-T** option of the **lvcreate** command to create a thin pool named **mythinpool** in the volume group **vg001** and that is 100M in size. Note that since you are creating a pool of physical space, you must specify the size of the pool. The **-T** option of the **lvcreate** command

does not take an argument; it deduces what type of device is to be created from the other options the command specifies.

```
# lvcreate -L 100M -T vg001/mythinpool
Rounding up size to full physical extent 4.00 MiB
Logical volume "mythinpool" created
# lvs
LV          VG      Attr      LSize   Pool Origin Data%   Move Log Copy%
Convert
my mythinpool vg001   twi-a-tz 100.00m                0.00
```

The following command uses the **-T** option of the **lvcreate** command to create a thin volume named **thinvolume** in the thin pool **vg001/mythinpool**. Note that in this case you are specifying virtual size, and that you are specifying a virtual size for the volume that is greater than the pool that contains it.

```
# lvcreate -V 1G -T vg001/mythinpool -n thinvolume
Logical volume "thinvolume" created
# lvs
LV          VG      Attr      LSize   Pool          Origin Data%   Move
Log Copy%  Convert
mythinpool  vg001   twi-a-tz 100.00m                0.00
thinvolume  vg001   Vwi-a-tz 1.00g mythinpool          0.00
```

The following command uses the **-T** option of the **lvcreate** command to create a thin pool and a thin volume in that pool by specifying both a size and a virtual size argument for the **lvcreate** command. This command creates a thin pool named **mythinpool** in the volume group **vg001** and it also creates a thin volume named **thinvolume** in that pool.

```
# lvcreate -L 100M -T vg001/mythinpool -V 1G -n thinvolume
Rounding up size to full physical extent 4.00 MiB
Logical volume "thinvolume" created
# lvs
LV          VG      Attr      LSize   Pool          Origin Data%   Move Log
Copy%  Convert
mythinpool  vg001   twi-a-tz 100.00m                0.00
thinvolume  vg001   Vwi-a-tz 1.00g mythinpool          0.00
```

You can also create a thin pool by specifying the **--thinpool** parameter of the **lvcreate** command. Unlike the **-T** option, the **--thinpool** parameter requires an argument, which is the name of the thin pool logical volume that you are creating. The following example specifies the **--thinpool** parameter of the **lvcreate** command to create a thin pool named **mythinpool** in the volume group **vg001** and that is 100M in size:

```
# lvcreate -L 100M --thinpool mythinpool vg001
Rounding up size to full physical extent 4.00 MiB
Logical volume "mythinpool" created
# lvs
LV          VG      Attr      LSize   Pool Origin Data%   Move Log Copy%
Convert
mythinpool  vg001   twi-a-tz 100.00m                0.00
```

Striping is supported for pool creation. The following command creates a 100M thin pool named **pool1** in volume group **vg001** with two 64 kB stripes and a chunk size of 256 kB. It also creates a 1T thin volume, **vg00/thin_lv**.

```
# lvcreate -i 2 -I 64 -c 256 -L 100M -T vg00/pool -V 1T --name thin_lv
```

You can extend the size of a thin volume with the **lvextend** command. You cannot, however, reduce the size of a thin pool.

The following command resizes an existing thin pool that is 100M in size by extending it another 100M.

```
# lvextend -L+100M vg001/mythinpool
Extending logical volume mythinpool to 200.00 MiB
Logical volume mythinpool successfully resized
# lvs
LV          VG      Attr      LSize   Pool       Origin Data%   Move Log
Copy% Convert
mythinpool  vg001    twi-a-tz  200.00m                0.00
thinvolume  vg001    Vwi-a-tz   1.00g mythinpool          0.00
```

As with other types of logical volumes, you can rename the volume with the **lvrename**, you can remove the volume with the **lvremove**, and you can display information about the volume with the **lvs** and **lvdisplay** commands.

By default, the **lvcreate** command sets the size of the thin pool's metadata logical volume according to the formula $(\text{Pool_LV_size} / \text{Pool_LV_chunk_size} * 64)$. If you will have large numbers of snapshots or if you have small chunk sizes for your thin pool and thus expect significant growth of the size of the thin pool at a later time, you may need to increase the default value of the thin pool's metadata volume with the **--poolmetadatasize** parameter of the **lvcreate** command. The supported value for the thin pool's metadata logical volume is in the range between 2MiB and 16GiB.

You can use the **--thinpool** parameter of the **lvconvert** command to convert an existing logical volume to a thin pool volume. When you convert an existing logical volume to a thin pool volume, you must use the **--poolmetadata** parameter in conjunction with the **--thinpool** parameter of the **lvconvert** to convert an existing logical volume to the thin pool volume's metadata volume.



NOTE

Converting a logical volume to a thin pool volume or a thin pool metadata volume destroys the content of the logical volume, since in this case the **lvconvert** does not preserve the content of the devices but instead overwrites the content.

The following example converts the existing logical volume **lv1** in volume group **vg001** to a thin pool volume and converts the existing logical volume **lv2** in volume group **vg001** to the metadata volume for that thin pool volume.

```
# lvconvert --thinpool vg001/lv1 --poolmetadata vg001/lv2
Converted vg001/lv1 to thin pool.
```

4.4.6. Creating Snapshot Volumes



NOTE

LVM supports thinly-provisioned snapshots. For information on creating thinly-provisioned snapshot volumes, see [Section 4.4.7, “Creating Thinly-Provisioned Snapshot Volumes”](#).

Use the **-s** argument of the **lvcreate** command to create a snapshot volume. A snapshot volume is writable.



NOTE

LVM snapshots are not supported across the nodes in a cluster. You cannot create a snapshot volume in a clustered volume group. However, if you need to create a consistent backup of data on a clustered logical volume you can activate the volume exclusively and then create the snapshot. For information on activating logical volumes exclusively on one node, see [Section 4.7, “Activating Logical Volumes on Individual Nodes in a Cluster”](#).



NOTE

LVM snapshots are supported for mirrored logical volumes.

Snapshots are supported for RAID logical volumes. For information on creating RAID logical volumes, see [Section 4.4.3, “RAID Logical Volumes”](#).

LVM does not allow you to create a snapshot volume that is larger than the size of the origin volume plus needed metadata for the volume. If you specify a snapshot volume that is larger than this, the system will create a snapshot volume that is only as large as will be needed for the size of the origin.

By default, a snapshot volume is skipped during normal activation commands. For information on controlling the activation of a snapshot volume, see [Section 4.4.20, “Controlling Logical Volume Activation”](#).

The following command creates a snapshot logical volume that is 100 MB in size named **/dev/vg00/snap**. This creates a snapshot of the origin logical volume named **/dev/vg00/lvol11**. If the original logical volume contains a file system, you can mount the snapshot logical volume on an arbitrary directory in order to access the contents of the file system to run a backup while the original file system continues to get updated.

```
# lvcreate --size 100M --snapshot --name snap /dev/vg00/lvol11
```

After you create a snapshot logical volume, specifying the origin volume on the **lvdisplay** command yields output that includes a list of all snapshot logical volumes and their status (active or inactive).

The following example shows the status of the logical volume **/dev/new_vg/lvol0**, for which a snapshot volume **/dev/new_vg/newvgsnap** has been created.

```
# lvdisplay /dev/new_vg/lvol0
--- Logical volume ---
LV Name                /dev/new_vg/lvol0
VG Name                new_vg
LV UUID                LBy1Tz-sr23-0jsI-LT03-nHLC-y8XW-EhC178
LV Write Access        read/write
LV snapshot status     source of
                       /dev/new_vg/newvgsnap1 [active]
LV Status              available
# open                 0
LV Size                52.00 MB
Current LE             13
Segments               1
```

```
Allocation                inherit
Read ahead sectors        0
Block device               253:2
```

The **lvs** command, by default, displays the origin volume and the current percentage of the snapshot volume being used. The following example shows the default output for the **lvs** command for a system that includes the logical volume **/dev/new_vg/lvol0**, for which a snapshot volume **/dev/new_vg/newvgsnap** has been created.

```
# lvs
LV          VG      Attr   LSize   Origin Snap%   Move Log Copy%
lvol0       new_vg  owi-a- 52.00M
newvgsnap1  new_vg  swi-a-  8.00M  lvol0    0.20
```



WARNING

Because the snapshot increases in size as the origin volume changes, it is important to monitor the percentage of the snapshot volume regularly with the **lvs** command to be sure it does not fill. A snapshot that is 100% full is lost completely, as a write to unchanged parts of the origin would be unable to succeed without corrupting the snapshot.

In addition to the snapshot itself being invalidated when full, any mounted file systems on that snapshot device are forcibly unmounted, avoiding the inevitable file system errors upon access to the mount point. In addition, you can specify the **snapshot_autoextend_threshold** option in the **lvm.conf** file. This option allows automatic extension of a snapshot whenever the remaining snapshot space drops below the threshold you set. This feature requires that there be unallocated space in the volume group.

LVM does not allow you to create a snapshot volume that is larger than the size of the origin volume plus needed metadata for the volume. Similarly, automatic extension of a snapshot will not increase the size of a snapshot volume beyond the maximum calculated size that is necessary for the snapshot. Once a snapshot has grown large enough to cover the origin, it is no longer monitored for automatic extension.

Information on setting **snapshot_autoextend_threshold** and **snapshot_autoextend_percent** is provided in the **lvm.conf** file itself. For information about the **lvm.conf** file, see [Appendix B, The LVM Configuration Files](#).


4.4.7. Creating Thinly-Provisioned Snapshot Volumes

Red Hat Enterprise Linux provides support for thinly-provisioned snapshot volumes. For information on the benefits and limitations of thin snapshot volumes, see [Section 2.3.6, “Thinly-Provisioned Snapshot Volumes”](#).



NOTE

This section provides an overview of the basic commands you use to create and grow thinly-provisioned snapshot volumes. For detailed information on LVM thin provisioning as well as information on using the LVM commands and utilities with thinly-provisioned logical volumes, see the **lvmthin(7)** man page.



IMPORTANT

When creating a thin snapshot volume, you do not specify the size of the volume. If you specify a size parameter, the snapshot that will be created will not be a thin snapshot volume and will not use the thin pool for storing data. For example, the command **lvcreate -s vg/thinvolume -L10M** will not create a thin snapshot, even though the origin volume is a thin volume.

Thin snapshots can be created for thinly-provisioned origin volumes, or for origin volumes that are not thinly-provisioned.

You can specify a name for the snapshot volume with the **--name** option of the **lvcreate** command. The following command creates a thinly-provisioned snapshot volume of the thinly-provisioned logical volume **vg001/thinvolume** that is named **mynsnapshot1**.

```
# lvcreate -s --name mynsnapshot1 vg001/thinvolume
Logical volume "mynsnapshot1" created
# lvs
LV          VG      Attr      LSize   Pool        Origin    Data%
Move Log Copy% Convert
mynsnapshot1 vg001   Vwi-a-tz   1.00g   mythinpool  thinvolume 0.00
mythinpool   vg001   twi-a-tz  100.00m                0.00
thinvolume   vg001   Vwi-a-tz   1.00g   mythinpool                0.00
```

A thin snapshot volume has the same characteristics as any other thin volume. You can independently activate the volume, extend the volume, rename the volume, remove the volume, and even snapshot the volume.

By default, a snapshot volume is skipped during normal activation commands. For information on controlling the activation of a snapshot volume, see [Section 4.4.20, “Controlling Logical Volume Activation”](#).

You can also create a thinly-provisioned snapshot of a non-thinly-provisioned logical volume. Since the non-thinly-provisioned logical volume is not contained within a thin pool, it is referred to as an *external origin*. External origin volumes can be used and shared by many thinly-provisioned snapshot volumes, even from different thin pools. The external origin must be inactive and read-only at the time the thinly-provisioned snapshot is created.

To create a thinly-provisioned snapshot of an external origin, you must specify the **--thinpool** option. The following command creates a thin snapshot volume of the read-only inactive volume **origin_volume**. The thin snapshot volume is named **mythinsnap**. The logical volume **origin_volume** then becomes the thin external origin for the thin snapshot volume **mythinsnap** in volume group **vg001** that will use the existing thin pool **vg001/pool**. Because the origin volume must be in the same volume group as the snapshot volume, you do not need to specify the volume group when specifying the origin logical volume.

```
# lvcreate -s --thinpool vg001/pool origin_volume --name mythinsnap
```

You can create a second thinly-provisioned snapshot volume of the first snapshot volume, as in the following command.

```
# lvcreate -s vg001/mythinsnap --name my2ndthinsnap
```

As of Red Hat Enterprise Linux 7.2, you can display a list of all ancestors and descendants of a thin snapshot logical volume by specifying the **lv_ancestors** and **lv_descendants** reporting fields of the **lvs** command.

In the following example:

- **stack1** is an origin volume in volume group **vg001**.
- **stack2** is a snapshot of **stack1**
- **stack3** is a snapshot of **stack2**
- **stack4** is a snapshot of **stack3**

Additionally:

- **stack5** is also a snapshot of **stack2**
- **stack6** is a snapshot of **stack5**

```
$ lvs -o name,lv_ancestors,lv_descendants vg001
LV      Ancestors                Descendants
stack1
stack2  stack1                    stack2, stack3, stack4, stack5, stack6
stack3  stack2, stack1            stack3, stack4, stack5, stack6
stack4  stack3, stack2, stack1     stack4
stack5  stack2, stack1            stack6
stack6  stack5, stack2, stack1
pool
```

NOTE

The **lv_ancestors** and **lv_descendants** fields display existing dependencies but do not track removed entries which can break a dependency chain if the entry was removed from the middle of the chain. For example, if you remove the logical volume **stack3** from this sample configuration, the display is as follows.

```
$ lvs -o name,lv_ancestors,lv_descendants vg001
LV      Ancestors                Descendants
stack1
stack2  stack1                    stack2, stack5, stack6
stack4
stack5  stack2, stack1            stack6
stack6  stack5, stack2, stack1
pool
```

As of Red Hat Enterprise Linux 7.3, however, you can configure your system to track and display logical volumes that have been removed, and you can display the full dependency chain that includes those volumes by specifying the **lv_ancestors_full** and **lv_descendants_full** fields. For information on tracking, displaying, and removing historical logical volumes, see [Section 4.4.21, “Tracking and Displaying Historical Logical Volumes \(Red Hat Enterprise Linux 7.3 and Later\)”](#).

4.4.8. Creating LVM Cache Logical Volumes

As of the Red Hat Enterprise Linux 7.1 release, LVM provides full support for LVM cache logical volumes. A cache logical volume uses a small logical volume consisting of fast block devices (such as SSD drives) to improve the performance of a larger and slower logical volume by storing the frequently used blocks on the smaller, faster logical volume.

LVM caching uses the following LVM logical volume types. All of these associated logical volumes must be in the same volume group.

- Origin logical volume — the large, slow logical volume
- Cache pool logical volume — the small, fast logical volume, which is composed of two devices: the cache data logical volume, and the cache metadata logical volume
- Cache data logical volume — the logical volume containing the data blocks for the cache pool logical volume
- Cache metadata logical volume — the logical volume containing the metadata for the cache pool logical volume, which holds the accounting information that specifies where data blocks are stored (for example, on the origin logical volume or the cache data logical volume).
- Cache logical volume — the logical volume containing the origin logical volume and the cache pool logical volume. This is the resultant usable device which encapsulates the various cache volume components.

The following procedure creates an LVM cache logical volume.

1. Create a volume group that contains a slow physical volume and a fast physical volume. In this example, **/dev/sde1** is a slow device and **/dev/sdf1** is a fast device and both devices are contained in volume group **VG**.

```
# pvcreate /dev/sde1
# pvcreate /dev/sdf1
# vgcreate VG /dev/sde1 /dev/sdf1
```

2. Create the origin volume. This example creates an origin volume named **lv** that is ten gigabytes in size and that consists of **/dev/sde1**, the slow physical volume.

```
# lvcreate -L 10G -n lv VG /dev/sde1
```

3. Create the cache pool logical volume. This example creates the cache pool logical volume named **cpool** on the fast device **/dev/sdf1**, which is part of the volume group **VG**. The cache pool logical volume this command creates consists of the hidden cache data logical volume **cpool_cdata** and the hidden cache metadata logical volume **cpool_cmeta**.

```
# lvcreate --type cache-pool -L 5G -n cpool VG /dev/sdf1
Using default stripesize 64.00 KiB.
Logical volume "cpool" created.
# lvs -a -o name,size,attr,devices VG
LV          LSize Attr      Devices
[cpool]      5.00g Cwi---C--- cpool_cdata(0)
[cpool_cdata] 5.00g Cwi-ao---- /dev/sdf1(4)
[cpool_cmeta] 8.00m ewi-ao---- /dev/sdf1(2)
```

For more complicated configurations you may need to create the cache data and the cache metadata logical volumes individually and then combine the volumes into a cache pool logical volume. For information on this procedure, see the **lvmdcache(7)** man page.

4. Create the cache logical volume by linking the cache pool logical volume to the origin logical volume. The resulting user-accessible cache logical volume takes the name of the origin logical volume. The origin logical volume becomes a hidden logical volume with **_corig** appended to the original name.

```
# lvconvert --type cache --cachepool VG/lv cpool
Logical volume cpool is now cached.
# lvs -a -o name,size,attr,devices vg
  LV              LSize Attr      Devices
  [cpool]          5.00g Cwi---C--- cpool_cdata(0)
  [cpool_cdata]    5.00g Cwi-ao---- /dev/sdf1(4)
  [cpool_cmeta]    8.00m ewi-ao---- /dev/sdf1(2)
  lv               10.00g Cwi-a-C--- lv_corig(0)
  [lv_corig]       10.00g owi-aoC--- /dev/sde1(0)
  [lvol0_pmspare]  8.00m ewi----- /dev/sdf1(0)
```

5. Optionally, as of Red Hat Enterprise Linux release 7.2, you can convert the cached logical volume to a thin pool logical volume. Note that any thin logical volumes created from the pool will share the cache.

The following command uses the fast device, **/dev/sdf1**, for allocating the thin pool metadata (**lv_tmeta**). This is the same device that is used by the cache pool volume, which means that the thin pool metadata volume shares that device with both the cache data logical volume **cpool_cdata** and the cache metadata logical volume **cpool_cmeta**.

```
# lvconvert --type thin-pool VG/lv /dev/sdf1
WARNING: Converting logical volume VG/lv to thin pool's data
volume with metadata wiping.
THIS WILL DESTROY CONTENT OF LOGICAL VOLUME (filesystem etc.)
Do you really want to convert VG/lv? [y/n]: y
Converted VG/lv to thin pool.
# lvs -a -o name,size,attr,devices vg
  LV              LSize Attr      Devices
  [cpool]          5.00g Cwi---C--- cpool_cdata(0)
  [cpool_cdata]    5.00g Cwi-ao---- /dev/sdf1(4)
  [cpool_cmeta]    8.00m ewi-ao---- /dev/sdf1(2)
  lv               10.00g twi-a-tz-- lv_tdata(0)
  [lv_tdata]       10.00g Cwi-aoC--- lv_tdata_corig(0)
  [lv_tdata_corig] 10.00g owi-aoC--- /dev/sde1(0)
  [lv_tmeta]       12.00m ewi-ao---- /dev/sdf1(1284)
  [lvol0_pmspare]  12.00m ewi----- /dev/sdf1(0)
  [lvol0_pmspare]  12.00m ewi----- /dev/sdf1(1287)
```

For further information on LVM cache volumes, including additional administrative examples, see the **lvmdcache(7)** man page.

For information on creating thinly-provisioned logical volumes, see [Section 4.4.5, “Creating Thinly-Provisioned Logical Volumes”](#).

4.4.9. Merging Snapshot Volumes

You can use the **--merge** option of the **lvconvert** command to merge a snapshot into its origin volume. If both the origin and snapshot volume are not open, the merge will start immediately. Otherwise, the merge will start the first time either the origin or snapshot are activated and both are closed. Merging a snapshot into an origin that cannot be closed, for example a root file system, is deferred until the next time the origin volume is activated. When merging starts, the resulting logical volume will have the origin's name, minor number and UUID. While the merge is in progress, reads or writes to the origin appear as they were directed to the snapshot being merged. When the merge finishes, the merged snapshot is removed.

The following command merges snapshot volume **vg00/lvol1_snap** into its origin.

```
# lvconvert --merge vg00/lvol1_snap
```

You can specify multiple snapshots on the command line, or you can use LVM object tags to specify that multiple snapshots be merged to their respective origins. In the following example, logical volumes **vg00/lvol1**, **vg00/lvol2**, and **vg00/lvol3** are all tagged with the tag **@some_tag**. The following command merges the snapshot logical volumes for all three volumes serially: **vg00/lvol1**, then **vg00/lvol2**, then **vg00/lvol3**. If the **--background** option were used, all snapshot logical volume merges would start in parallel.

```
# lvconvert --merge @some_tag
```

For information on tagging LVM objects, see [Appendix D, LVM Object Tags](#). For further information on the **lvconvert --merge** command, see the **lvconvert(8)** man page.

4.4.10. Persistent Device Numbers

Major and minor device numbers are allocated dynamically at module load. Some applications work best if the block device is always activated with the same device (major and minor) number. You can specify these with the **lvcreate** and the **lvchange** commands by using the following arguments:

```
--persistent y --major major --minor minor
```

Use a large minor number to be sure that it has not already been allocated to another device dynamically.

If you are exporting a file system using NFS, specifying the **fsid** parameter in the exports file may avoid the need to set a persistent device number within LVM.

4.4.11. Changing the Parameters of a Logical Volume Group

To change the parameters of a logical volume, use the **lvchange** command. For a listing of the parameters you can change, see the **lvchange(8)** man page.

You can use the **lvchange** command to activate and deactivate logical volumes. To activate and deactivate all the logical volumes in a volume group at the same time, use the **vgchange** command, as described in [Section 4.3.9, “Changing the Parameters of a Volume Group”](#).

The following command changes the permission on volume **lvol1** in volume group **vg00** to be read-only.

```
# lvchange -pr vg00/lvol1
```

4.4.12. Renaming Logical Volumes

To rename an existing logical volume, use the **lvrename** command.

Either of the following commands renames logical volume **lvold** in volume group **vg02** to **lvnew**.

```
# lvrename /dev/vg02/lvold /dev/vg02/lvnew
```

```
# lvrename vg02 lvold lvnew
```

For more information on activating logical volumes on individual nodes in a cluster, see [Section 4.7](#), “[Activating Logical Volumes on Individual Nodes in a Cluster](#)”.

4.4.13. Removing Logical Volumes

To remove an inactive logical volume, use the **lvremove** command. If the logical volume is currently mounted, unmount the volume before removing it. In addition, in a clustered environment you must deactivate a logical volume before it can be removed.

The following command removes the logical volume **/dev/testvg/testlv** from the volume group **testvg**. Note that in this case the logical volume has not been deactivated.

```
# lvremove /dev/testvg/testlv
Do you really want to remove active logical volume "testlv"? [y/n]: y
Logical volume "testlv" successfully removed
```

You could explicitly deactivate the logical volume before removing it with the **lvchange -an** command, in which case you would not see the prompt verifying whether you want to remove an active logical volume.

4.4.14. Displaying Logical Volumes

There are three commands you can use to display properties of LVM logical volumes: **lvs**, **lvdisplay**, and **lvscan**.

The **lvs** command provides logical volume information in a configurable form, displaying one line per logical volume. The **lvs** command provides a great deal of format control, and is useful for scripting. For information on using the **lvs** command to customize your output, see [Section 4.8](#), “[Customized Reporting for LVM](#)”.

The **lvdisplay** command displays logical volume properties (such as size, layout, and mapping) in a fixed format.

The following command shows the attributes of **lv012** in **vg00**. If snapshot logical volumes have been created for this original logical volume, this command shows a list of all snapshot logical volumes and their status (active or inactive) as well.

```
# lvdisplay -v /dev/vg00/lv012
```

The **lvscan** command scans for all logical volumes in the system and lists them, as in the following example.

```
# lvscan
ACTIVE                               '/dev/vg0/gfslv' [1.46 GB] inherit
```

4.4.15. Growing Logical Volumes

To increase the size of a logical volume, use the **lvextend** command.

When you extend the logical volume, you can indicate how much you want to extend the volume, or how large you want it to be after you extend it.

The following command extends the logical volume **/dev/myvg/homevol** to 12 gigabytes.

```
# lvextend -L12G /dev/myvg/homevol
lvextend -- extending logical volume "/dev/myvg/homevol" to 12 GB
lvextend -- doing automatic backup of volume group "myvg"
lvextend -- logical volume "/dev/myvg/homevol" successfully extended
```

The following command adds another gigabyte to the logical volume **/dev/myvg/homevol**.

```
# lvextend -L+1G /dev/myvg/homevol
lvextend -- extending logical volume "/dev/myvg/homevol" to 13 GB
lvextend -- doing automatic backup of volume group "myvg"
lvextend -- logical volume "/dev/myvg/homevol" successfully extended
```

As with the **lvcreate** command, you can use the **-l** argument of the **lvextend** command to specify the number of extents by which to increase the size of the logical volume. You can also use this argument to specify a percentage of the volume group, or a percentage of the remaining free space in the volume group. The following command extends the logical volume called **testlv** to fill all of the unallocated space in the volume group **myvg**.

```
# lvextend -l +100%FREE /dev/myvg/testlv
Extending logical volume testlv to 68.59 GB
Logical volume testlv successfully resized
```

After you have extended the logical volume it is necessary to increase the file system size to match.

By default, most file system resizing tools will increase the size of the file system to be the size of the underlying logical volume so you do not need to worry about specifying the same size for each of the two commands.

4.4.16. Shrinking Logical Volumes

You can reduce the size of a logical volume with the **lvreduce** command.



NOTE

Shrinking is not supported on a GFS2 or XFS file system, so you cannot reduce the size of a logical volume that contains a GFS2 or XFS file system.

If the logical volume you are reducing contains a file system, to prevent data loss you must ensure that the file system is not using the space in the logical volume that is being reduced. For this reason, it is recommended that you use the **--resizefs** option of the **lvreduce** command when the logical

volume contains a file system. When you use this option, the **lvreduce** command attempts to reduce the file system before shrinking the logical volume. If shrinking the file system fails, as can occur if the file system is full or the file system does not support shrinking, then the **lvreduce** command will fail and not attempt to shrink the logical volume.



WARNING

In most cases, the **lvreduce** command warns about possible data loss and asks for a confirmation. However, you should not rely on these confirmation prompts to prevent data loss because in some cases you will not see these prompts, such as when the logical volume is inactive or the **--resizefs** option is not used.

Note that using the **--test** option of the **lvreduce** command does not indicate where the operation is safe, as this option does not check the file system or test the file system resize.

The following command shrinks the logical volume **lv01** in volume group **vg00** to be 64 megabytes. In this example, **lv01** contains a file system, which this command resizes together with the logical volume. This example shows the output to the command.

```
# lvreduce --resizefs -L 64M vg00/lv01
fsck from util-linux 2.23.2
/dev/mapper/vg00-lv01: clean, 11/25688 files, 8896/102400 blocks
resize2fs 1.42.9 (28-Dec-2013)
Resizing the filesystem on /dev/mapper/vg00-lv01 to 65536 (1k) blocks.
The filesystem on /dev/mapper/vg00-lv01 is now 65536 blocks long.

Size of logical volume vg00/lv01 changed from 100.00 MiB (25 extents)
to 64.00 MiB (16 extents).
Logical volume vg00/lv01 successfully resized.
```

Specifying the **-** sign before the resize value indicates that the value will be subtracted from the logical volume's actual size. The following example shows the command you would use if, instead of shrinking a logical volume to an absolute size of 64 megabytes, you wanted to shrink the volume by a value 64 megabytes.

```
# lvreduce --resizefs -L -64M vg00/lv01
```

4.4.17. Extending a Striped Volume

In order to increase the size of a striped logical volume, there must be enough free space on the underlying physical volumes that make up the volume group to support the stripe. For example, if you have a two-way stripe that that uses up an entire volume group, adding a single physical volume to the volume group will not enable you to extend the stripe. Instead, you must add at least two physical volumes to the volume group.

For example, consider a volume group **vg** that consists of two underlying physical volumes, as displayed with the following **vgs** command.

■

```
# vgs
VG    #PV #LV #SN Attr   VSize   VFree
vg      2  0   0 wz--n- 271.31G 271.31G
```

You can create a stripe using the entire amount of space in the volume group.

```
# lvcreate -n stripe1 -L 271.31G -i 2 vg
Using default stripesize 64.00 KB
Rounding up size to full physical extent 271.31 GB
Logical volume "stripe1" created
# lvs -a -o +devices
LV      VG    Attr   LSize   Origin Snap%   Move Log Copy%  Devices
stripe1 vg    -wi-a- 271.31G
/dev/sda1(0),/dev/sdb1(0)
```

Note that the volume group now has no more free space.

```
# vgs
VG    #PV #LV #SN Attr   VSize   VFree
vg      2  1   0 wz--n- 271.31G    0
```

The following command adds another physical volume to the volume group, which then has 135 gigabytes of additional space.

```
# vgextend vg /dev/sdc1
Volume group "vg" successfully extended
# vgs
VG    #PV #LV #SN Attr   VSize   VFree
vg      3  1   0 wz--n- 406.97G 135.66G
```

At this point you cannot extend the striped logical volume to the full size of the volume group, because two underlying devices are needed in order to stripe the data.

```
# lvextend vg/stripe1 -L 406G
Using stripesize of last segment 64.00 KB
Extending logical volume stripe1 to 406.00 GB
Insufficient suitable allocatable extents for logical volume stripe1:
34480
more required
```

To extend the striped logical volume, add another physical volume and then extend the logical volume. In this example, having added two physical volumes to the volume group we can extend the logical volume to the full size of the volume group.

```
# vgextend vg /dev/sdd1
Volume group "vg" successfully extended
# vgs
VG    #PV #LV #SN Attr   VSize   VFree
vg      4  1   0 wz--n- 542.62G 271.31G
# lvextend vg/stripe1 -L 542G
Using stripesize of last segment 64.00 KB
Extending logical volume stripe1 to 542.00 GB
Logical volume stripe1 successfully resized
```

If you do not have enough underlying physical devices to extend the striped logical volume, it is possible to extend the volume anyway if it does not matter that the extension is not striped, which may result in uneven performance. When adding space to the logical volume, the default operation is to use the same striping parameters of the last segment of the existing logical volume, but you can override those parameters. The following example extends the existing striped logical volume to use the remaining free space after the initial **lvextend** command fails.

```
# lvextend vg/stripe1 -L 406G
Using stripesize of last segment 64.00 KB
Extending logical volume stripe1 to 406.00 GB
Insufficient suitable allocatable extents for logical volume stripe1:
34480
more required
# lvextend -i1 -l+100%FREE vg/stripe1
```

4.4.18. Extending a RAID Volume

You can grow RAID logical volumes with the **lvextend** command without performing a synchronization of the new RAID regions.

If you specify the **--nosync** option when you create a RAID logical volume with the **lvcreate** command, the RAID regions are not synchronized when the logical volume is created. If you later extend a RAID logical volume that you have created with the **--nosync** option, the RAID extensions are not synchronized at that time, either.

You can determine whether an existing logical volume was created with the **--nosync** option by using the **lvs** command to display the volume's attributes. A logical volume will show "R" as the first character in the attribute field if it is a RAID volume that was created without an initial synchronization, and it will show "r" if it was created with initial synchronization.

The following command displays the attributes of a RAID logical volume named **lv** that was created without initial synchronization, showing "R" as the first character in the attribute field. The seventh character in the attribute field is "r", indicating a target type of RAID. For information on the meaning of the attribute field, see [Table 4.5, "lvs Display Fields"](#).

```
# lvs vg
LV VG Attr LSize Pool Origin Snap% Move Log Cpy%Sync Convert
lv vg Rwi-a-r- 5.00g 100.00
```

If you grow this logical volume with the **lvextend** command, the RAID extension will not be resynchronized.

If you created a RAID logical volume without specifying the **--nosync** option of the **lvcreate** command, you can grow the logical volume without resynchronizing the mirror by specifying the **--nosync** option of the **lvextend** command.

The following example extends a RAID logical volume that was created without the **--nosync** option, indicated that the RAID volume was synchronized when it was created. This example, however, specifies that the volume not be synchronized when the volume is extended. Note that the volume has an attribute of "r", but after executing the **lvextend** command with the **--nosync** option the volume has an attribute of "R".

```
# lvs vg
LV VG Attr LSize Pool Origin Snap% Move Log Cpy%Sync Convert
```

```
lv    vg    rwi-a-r- 20.00m                                100.00
# lvextend -L +5G vg/lv --nosync
Extending 2 mirror images.
Extending logical volume lv to 5.02 GiB
Logical volume lv successfully resized
# lvs vg
LV    VG    Attr      LSize Pool Origin Snap%  Move Log       Cpy%Sync
Convert
lv    vg    Rwi-a-r- 5.02g                                100.00
```

If a RAID volume is inactive, it will not automatically skip synchronization when you extend the volume, even if you create the volume with the **--nosync** option specified. Instead, you will be prompted whether to do a full resync of the extended portion of the logical volume.



NOTE

If a RAID volume is performing recovery, you cannot extend the logical volume if you created or extended the volume with the **--nosync** option specified. If you did not specify the **--nosync** option, however, you can extend the RAID volume while it is recovering.

4.4.19. Extending a Logical Volume with the **cling** Allocation Policy

When extending an LVM volume, you can use the **--alloc cling** option of the **lvextend** command to specify the **cling** allocation policy. This policy will choose space on the same physical volumes as the last segment of the existing logical volume. If there is insufficient space on the physical volumes and a list of tags is defined in the **lvm.conf** file, LVM will check whether any of the tags are attached to the physical volumes and seek to match those physical volume tags between existing extents and new extents.

For example, if you have logical volumes that are mirrored between two sites within a single volume group, you can tag the physical volumes according to where they are situated by tagging the physical volumes with **@site1** and **@site2** tags. You can then specify the following line in the **lvm.conf** file:

```
cling_tag_list = [ "@site1", "@site2" ]
```

For information on tagging physical volumes, see [Appendix D, LVM Object Tags](#).

In the following example, the **lvm.conf** file has been modified to contain the following line:

```
cling_tag_list = [ "@A", "@B" ]
```

Also in this example, a volume group **taft** has been created that consists of the physical volumes **/dev/sdb1**, **/dev/sdc1**, **/dev/sdd1**, **/dev/sde1**, **/dev/sdf1**, **/dev/sdg1**, and **/dev/sdh1**. These physical volumes have been tagged with tags **A**, **B**, and **C**. The example does not use the **C** tag, but this will show that LVM uses the tags to select which physical volumes to use for the mirror legs.

```
# pvs -a -o +pv_tags /dev/sd[bcdefgh]
PV          VG    Fmt  Attr PSize  PFree  PV Tags
/dev/sdb1   taft lvm2 a--  15.00g 15.00g A
/dev/sdc1   taft lvm2 a--  15.00g 15.00g B
/dev/sdd1   taft lvm2 a--  15.00g 15.00g B
/dev/sde1   taft lvm2 a--  15.00g 15.00g C
```

```
/dev/sdf1 taft lvm2 a-- 15.00g 15.00g C
/dev/sdg1 taft lvm2 a-- 15.00g 15.00g A
/dev/sdh1 taft lvm2 a-- 15.00g 15.00g A
```

The following command creates a 10 gigabyte mirrored volume from the volume group **taft**.

```
# lvcreate --type raid1 -m 1 -n mirror --nosync -L 10G taft
WARNING: New raid1 won't be synchronised. Don't read what you didn't
write!
Logical volume "mirror" created
```

The following command shows which devices are used for the mirror legs and RAID metadata subvolumes.

```
# lvs -a -o +devices
LV          VG   Attr      LSize   Log Cpy%Sync Devices
mirror      taft Rwi-a-r--- 10.00g   100.00
mirror_rimage_0(0),mirror_rimage_1(0)
[mirror_rimage_0] taft iwi-aor--- 10.00g   /dev/sdb1(1)
[mirror_rimage_1] taft iwi-aor--- 10.00g   /dev/sdc1(1)
[mirror_rmeta_0] taft ewi-aor--- 4.00m    /dev/sdb1(0)
[mirror_rmeta_1] taft ewi-aor--- 4.00m    /dev/sdc1(0)
```

The following command extends the size of the mirrored volume, using the **cling** allocation policy to indicate that the mirror legs should be extended using physical volumes with the same tag.

```
# lvextend --alloc cling -L +10G taft/mirror
Extending 2 mirror images.
Extending logical volume mirror to 20.00 GiB
Logical volume mirror successfully resized
```

The following display command shows that the mirror legs have been extended using physical volumes with the same tag as the leg. Note that the physical volumes with a tag of **C** were ignored.

```
# lvs -a -o +devices
LV          VG   Attr      LSize   Log Cpy%Sync Devices
mirror      taft Rwi-a-r--- 20.00g   100.00
mirror_rimage_0(0),mirror_rimage_1(0)
[mirror_rimage_0] taft iwi-aor--- 20.00g   /dev/sdb1(1)
[mirror_rimage_0] taft iwi-aor--- 20.00g   /dev/sdg1(0)
[mirror_rimage_1] taft iwi-aor--- 20.00g   /dev/sdc1(1)
[mirror_rimage_1] taft iwi-aor--- 20.00g   /dev/sdd1(0)
[mirror_rmeta_0] taft ewi-aor--- 4.00m    /dev/sdb1(0)
[mirror_rmeta_1] taft ewi-aor--- 4.00m    /dev/sdc1(0)
```

4.4.20. Controlling Logical Volume Activation

You can flag a logical volume to be skipped during normal activation commands with the **-k** or **--setactivationskip {y|n}** option of the **lvcreate** or **lvchange** command. This flag is not applied during deactivation.

You can determine whether this flag is set for a logical volume with the **lvs** command, which displays the **k** attribute as in the following example.

```
# lvs vg/thin1s1
LV          VG Attr          LSize Pool  Origin
thin1s1     vg  Vwi---tz-k 1.00t pool0 thin1
```

By default, thin snapshot volumes are flagged for activation skip. You can activate a logical volume with the **k** attribute set by using the **-K** or **--ignoreactivationskip** option in addition to the standard **-ay** or **--activate y** option.

The following command activates a thin snapshot logical volume.

```
# lvchange -ay -K VG/SnapLV
```

The persistent "activation skip" flag can be turned off when the logical volume is created by specifying the **-kn** or **--setactivationskip n** option of the **lvcreate** command. You can turn the flag off for an existing logical volume by specifying the **-kn** or **--setactivationskip n** option of the **lvchange** command. You can turn the flag on again with the **-ky** or **--setactivationskip y** option.

The following command creates a snapshot logical volume without the activation skip flag

```
# lvcreate --type thin -n SnapLV -kn -s ThinLV --thinpool VG/ThinPoolLV
```

The following command removes the activation skip flag from a snapshot logical volume.

```
# lvchange -kn VG/SnapLV
```

You can control the default activation skip setting with the **auto_set_activation_skip** setting in the **/etc/lvm/lvm.conf** file.

4.4.21. Tracking and Displaying Historical Logical Volumes (Red Hat Enterprise Linux 7.3 and Later)

As of Red Hat Enterprise Linux 7.3, you can configure your system to track thin snapshot and thin logical volumes that have been removed by enabling the **record_lvs_history** metadata option in the **lvm.conf** configuration file. This allows you to display a full thin snapshot dependency chain that includes logical volumes that have been removed from the original dependency chain and have become *historical* logical volumes.

You can configure your system to retain historical volumes for a defined period of time by specifying the retention time, in seconds, with the **lvs_history_retention_time** metadata option in the **lvm.conf** configuration file.

A historical logical volume retains a simplified representation of the logical volume that has been removed, including the following reporting fields for the volume:

- **lv_time_removed**: the removal time of the logical volume
- **lv_time**: the creation time of the logical volume
- **lv_name**: the name of the logical volume
- **lv_uuid**: the UUID of the logical volume
- **vg_name**: the volume group that contains the logical volume.

When a volume is removed, the historical logical volume name acquires a hyphen as a prefix. For example, when you remove the logical volume **lvvol1**, the name of the historical volume is **-lvvol1**. A historical logical volume cannot be reactivated.

Even when the **record_lvs_history** metadata option enabled, you can prevent the retention of historical logical volumes on an individual basis when you remove a logical volume by specifying the **--nohistory** option of the **lvremove** command.

To include historical logical volumes in volume display, you specify the **-H|--history** option of an LVM display command. You can display a full thin snapshot dependency chain that includes historical volumes by specifying the **lv_full_ancestors** and **lv_full_descendants** reporting fields along with the **-H** option.

The following series of commands provides examples of how you can display and manage historical logical volumes.

1. Ensure that historical logical volumes are retained by setting **record_lvs_history=1** in the **lvm.conf** file. This metadata option is not enabled by default.
2. Enter the following command to display a thin provisioned snapshot chain.

In this example:

- **lvvol1** is an origin volume, the first volume in the chain.
- **lvvol2** is a snapshot of **lvvol1**.
- **lvvol3** is a snapshot of **lvvol2**.
- **lvvol4** is a snapshot of **lvvol3**.
- **lvvol5** is also a snapshot of **lvvol3**.

Note that even though the example **lvs** display command includes the **-H** option, no thin snapshot volume has yet been removed and there are no historical logical volumes to display.

```
# lvs -H -o name,full_ancestors,full_descendants
LV      FAncestors          FDescendants
lvvol1                                     lvvol2,lvvol3,lvvol4,lvvol5
lvvol2 lvvol1               lvvol3,lvvol4,lvvol5
lvvol3 lvvol2,lvvol1        lvvol4,lvvol5
lvvol4 lvvol3,lvvol2,lvvol1
lvvol5 lvvol3,lvvol2,lvvol1
pool
```

3. Remove logical volume **lvvol3** from the snapshot chain, then run the following **lvs** command again to see how historical logical volumes are displayed, along with their ancestors and descendants.

```
# lvremove -f vg/lvol3
Logical volume "lvvol3" successfully removed
# lvs -H -o name,full_ancestors,full_descendants
LV      FAncestors          FDescendants
lvvol1                                     lvvol2,-lvvol3,lvvol4,lvvol5
lvvol2 lvvol1               -lvvol3,lvvol4,lvvol5
```

```

-lvol3 lvol2,lvol1          lvol4,lvol5
lvol4 -lvol3,lvol2,lvol1
lvol5 -lvol3,lvol2,lvol1
pool

```

4. You can use the **lv_time_removed** reporting field to display the time a historical volume was removed.

```

# lvs -H -o name,full_ancestors,full_descendants,time_removed
LV      FAncestors          FDescendants          RTime
lvol1                  lvol2, -lvol3,lvol4,lvol5
lvol2 lvol1                  -lvol3,lvol4,lvol5
-lvol3 lvol2,lvol1          lvol4,lvol5          2016-03-14
14:14:32 +0100
lvol4 -lvol3,lvol2,lvol1
lvol5 -lvol3,lvol2,lvol1
pool

```

5. You can reference historical logical volumes individually in a display command by specifying the *vgname/lvname* format, as in the following example. Note that the fifth bit in the **lv_attr** field is set to **h** to indicate the volume is a historical volume.

```

# lvs -H vg/-lvol3
LV      VG      Attr          LSize
-lvol3  vg      ---h-----      0

```

6. LVM does not keep historical logical volumes if the volume has no live descendant. This means that if you remove a logical volume at the end of a snapshot chain, the logical volume is not retained as a historical logical volume.

```

# lvremove -f vg/lvol5
Automatically removing historical logical volume vg/-lvol5.
Logical volume "lvol5" successfully removed
# lvs -H -o name,full_ancestors,full_descendants
LV      FAncestors          FDescendants
lvol1                  lvol2, -lvol3,lvol4
lvol2 lvol1                  -lvol3,lvol4
-lvol3 lvol2,lvol1          lvol4
lvol4 -lvol3,lvol2,lvol1
pool

```

7. Run the following commands to remove the volume **lvol1** and **lvol2** and to see how the **lvs** command displays the volumes once they have been removed.

```

# lvremove -f vg/lvol1 vg/lvol2
Logical volume "lvol1" successfully removed
Logical volume "lvol2" successfully removed
# lvs -H -o name,full_ancestors,full_descendants
LV      FAncestors          FDescendants
-lvol1                  -lvol2, -lvol3,lvol4
-lvol2 -lvol1                  -lvol3,lvol4
-lvol3 -lvol2, -lvol1          lvol4
lvol4 -lvol3, -lvol2, -lvol1
pool

```

8. To remove a historical logical volume completely, you can run the **lvremove** command again, specifying the name of the historical volume that now includes the hyphen, as in the following example.

```
# lvremove -f vg/-lvol3
Historical logical volume "lvol3" successfully removed
# lvs -H -o name,full_ancestors,full_descendants
LV      FAncestors      FDescendants
-lvol1              -lvol2,lvol4
-lvol2 -lvol1        lvol4
lvol4   -lvol2,-lvol1
pool
```

9. A historical logical volumes is retained as long as there is a chain that includes live volumes in its descendants. This means that removing a historical logical volume also removes all of the logical volumes in the chain if no existing descendant is linked to them, as shown in the following example.

```
# lvremove -f vg/lvol4
Automatically removing historical logical volume vg/-lvol1.
Automatically removing historical logical volume vg/-lvol2.
Automatically removing historical logical volume vg/-lvol4.
Logical volume "lvol4" successfully removed
```

4.5. CONTROLLING LVM DEVICE SCANS WITH FILTERS

At startup, the **vgscan** command is run to scan the block devices on the system looking for LVM labels, to determine which of them are physical volumes and to read the metadata and build up a list of volume groups. The names of the physical volumes are stored in the LVM cache file of each node in the system, **/etc/lvm/cache/.cache**. Subsequent commands may read that file to avoiding rescanning.

You can control which devices LVM scans by setting up filters in the **lvm.conf** configuration file. The filters in the **lvm.conf** file consist of a series of simple regular expressions that get applied to the device names in the **/dev** directory to decide whether to accept or reject each block device found.

The following examples show the use of filters to control which devices LVM scans. Note that some of these examples do not necessarily represent recommended practice, as the regular expressions are matched freely against the complete pathname. For example, **a/loop/** is equivalent to **a/. *loop. */** and would match **/dev/solooperation/lvol1**.

The following filter adds all discovered devices, which is the default behavior as there is no filter configured in the configuration file:

```
filter = [ "a/.*/" ]
```

The following filter removes the cdrom device in order to avoid delays if the drive contains no media:

```
filter = [ "r|/dev/cdrom|" ]
```

The following filter adds all loop and removes all other block devices:

```
filter = [ "a/loop.*/", "r/.*/" ]
```

The following filter adds all loop and IDE and removes all other block devices:

```
filter =[ "a|loop.*|", "a|/dev/hd.*|", "r|.*)" ]
```

The following filter adds just partition 8 on the first IDE drive and removes all other block devices:

```
filter = [ "a|^/dev/hda8$|", "r|.*)" ]
```



NOTE

When the **lvmetad** daemon is running, the **filter =** setting in the **/etc/lvm/lvm.conf** file does not apply when you execute the **pvscan --cache device** command. To filter devices, you need to use the **global_filter =** setting. Devices that fail the global filter are not opened by LVM and are never scanned. You may need to use a global filter, for example, when you use LVM devices in VMs and you do not want the contents of the devices in the VMs to be scanned by the physical host.

For more information on the **lvm.conf** file, see [Appendix B, The LVM Configuration Files](#) and the **lvm.conf(5)** man page.

4.6. ONLINE DATA RELOCATION

You can move data while the system is in use with the **pvmove** command.

The **pvmove** command breaks up the data to be moved into sections and creates a temporary mirror to move each section. For more information on the operation of the **pvmove** command, see the **pvmove(8)** man page.



NOTE

In order to perform a **pvmove** operation in a cluster, you should ensure that the **cmirror** package is installed and the **cmirrord** service is running.

The following command moves all allocated space off the physical volume **/dev/sdc1** to other free physical volumes in the volume group:

```
# pvmove /dev/sdc1
```

The following command moves just the extents of the logical volume **MyLV**.

```
# pvmove -n MyLV /dev/sdc1
```

Since the **pvmove** command can take a long time to execute, you may want to run the command in the background to avoid display of progress updates in the foreground. The following command moves all extents allocated to the physical volume **/dev/sdc1** over to **/dev/sdf1** in the background.

```
# pvmove -b /dev/sdc1 /dev/sdf1
```

The following command reports the progress of the **pvmove** command as a percentage at five second intervals.

```
# pvmove -i5 /dev/sdd1
```

4.7. ACTIVATING LOGICAL VOLUMES ON INDIVIDUAL NODES IN A CLUSTER

If you have LVM installed in a cluster environment, you may at times need to activate logical volumes exclusively on one node.

To activate logical volumes exclusively on one node, use the **lvchange -aey** command. Alternatively, you can use **lvchange -aly** command to activate logical volumes only on the local node but not exclusively. You can later activate them on additional nodes concurrently.

You can also activate logical volumes on individual nodes by using LVM tags, which are described in [Appendix D, LVM Object Tags](#). You can also specify activation of nodes in the configuration file, which is described in [Appendix B, The LVM Configuration Files](#).

4.8. CUSTOMIZED REPORTING FOR LVM

LVM provides a wide range of configuration and command line options to produce customized reports and to filter the report's output. For a full description of LVM reporting features and capabilities, see the **lvmreport(7)** man page.

You can produce concise and customizable reports of LVM objects with the **pvs**, **lvs**, and **vgs** commands. The reports that these commands generate include one line of output for each object. Each line contains an ordered list of fields of properties related to the object. There are five ways to select the objects to be reported: by physical volume, volume group, logical volume, physical volume segment, and logical volume segment.

As of the Red Hat Enterprise Linux 7.3 release, you can report information about physical volumes, volume groups, logical volumes, physical volume segments, and logical volume segments all at once with the **lvm fullreport** command. For information on this command and its capabilities, see the **lvm-fullreport(8)** man page.

As of the Red Hat Enterprise Linux 7.3 release, LVM supports log reports, which contain a log of operations, messages, and per-object status with complete object identification collected during LVM command execution. For an example of an LVM log report, see [Section 4.8.6, “Command Log Reporting \(Red Hat Enterprise Linux 7.3 and later\)”](#). For further information about the LVM log report, see the **lvmreport(7)** man page.

The following sections provide summary information on the use of the **pvs**, **lvs**, and **vgs** commands to customize a report:

- [Section 4.8.1, “Format Control”](#), which provides a summary of command arguments you can use to control the format of the report.
- [Section 4.8.2, “Object Display Fields”](#), which provides a list of the fields you can display for each LVM object.
- [Section 4.8.3, “Sorting LVM Reports”](#), which provides a summary of command arguments you can use to sort the generated report.
- [Section 4.8.4, “Specifying Units”](#), which provides instructions for specifying the units of the report output.

- [Section 4.8.5, “JSON Format Output \(Red Hat Enterprise Linux 7.3 and later\)”](#), which provides an example that specifies JSON format output (Red Hat Enterprise Linux 7.3 and later).
- [Section 4.8.6, “Command Log Reporting \(Red Hat Enterprise Linux 7.3 and later\)”](#), which provides an example of a command log.

4.8.1. Format Control

Whether you use the **pvs**, **lvs**, or **vgs** command determines the default set of fields displayed and the sort order. You can control the output of these commands with the following arguments:

- You can change what fields are displayed to something other than the default by using the **-o** argument. For example, the following output is the default display for the **pvs** command (which displays information about physical volumes).

```
# pvs
PV          VG      Fmt  Attr PSize  PFree
/dev/sdb1   new_vg  lvm2 a-   17.14G 17.14G
/dev/sdc1   new_vg  lvm2 a-   17.14G 17.09G
/dev/sdd1   new_vg  lvm2 a-   17.14G 17.14G
```

The following command displays only the physical volume name and size.

```
# pvs -o pv_name,pv_size
PV          PSize
/dev/sdb1   17.14G
/dev/sdc1   17.14G
/dev/sdd1   17.14G
```

- You can append a field to the output with the plus sign (+), which is used in combination with the **-o** argument.

The following example displays the UUID of the physical volume in addition to the default fields.

```
# pvs -o +pv_uuid
PV          VG      Fmt  Attr PSize  PFree  PV UUID
/dev/sdb1   new_vg  lvm2 a-   17.14G 17.14G onFF2w-1fLC-ughJ-D9eB-M7iv-6XqA-dqGeXY
/dev/sdc1   new_vg  lvm2 a-   17.14G 17.09G Joqlch-yWSj-kuEn-IdwM-01S9-X08M-mcpsVe
/dev/sdd1   new_vg  lvm2 a-   17.14G 17.14G yvfvZK-Cf31-j75k-dECm-0RZ3-0dGW-UqkCS
```

- Adding the **-v** argument to a command includes some extra fields. For example, the **pvs -v** command will display the **DevSize** and **PV UUID** fields in addition to the default fields.

```
# pvs -v
Scanning for physical volume names
PV          VG      Fmt  Attr PSize  PFree  DevSize PV UUID
/dev/sdb1   new_vg  lvm2 a-   17.14G 17.14G 17.14G onFF2w-1fLC-ughJ-D9eB-M7iv-6XqA-dqGeXY
/dev/sdc1   new_vg  lvm2 a-   17.14G 17.09G 17.14G Joqlch-yWSj-
```

```
kuEn-IdwM-01S9-X08M-mcpsVe
/dev/sdd1 new_vg lvm2 a- 17.14G 17.14G 17.14G yvfvZK-Cf31-
j75k-dECm-0RZ3-0dGW-tUqkCS
```

- The **--noheadings** argument suppresses the headings line. This can be useful for writing scripts.

The following example uses the **--noheadings** argument in combination with the **pv_name** argument, which will generate a list of all physical volumes.

```
# pvs --noheadings -o pv_name
/dev/sdb1
/dev/sdc1
/dev/sdd1
```

- The **--separator separator** argument uses *separator* to separate each field.

The following example separates the default output fields of the **pvs** command with an equals sign (=).

```
# pvs --separator =
PV=VG=Fmt=Attr=PSize=PFree
/dev/sdb1=new_vg=lvm2=a-=17.14G=17.14G
/dev/sdc1=new_vg=lvm2=a-=17.14G=17.09G
/dev/sdd1=new_vg=lvm2=a-=17.14G=17.14G
```

To keep the fields aligned when using the **separator** argument, use the **separator** argument in conjunction with the **--aligned** argument.

```
# pvs --separator = --aligned
PV          =VG      =Fmt =Attr=PSize =PFree
/dev/sdb1 =new_vg=lvm2=a- =17.14G=17.14G
/dev/sdc1 =new_vg=lvm2=a- =17.14G=17.09G
/dev/sdd1 =new_vg=lvm2=a- =17.14G=17.14G
```

You can use the **-P** argument of the **lvs** or **vgs** command to display information about a failed volume that would otherwise not appear in the output. For information on the output this argument yields, see [Section 6.2, “Displaying Information on Failed Devices”](#).

For a full listing of display arguments, see the **pvs(8)**, **vgs(8)** and **lvs(8)** man pages.

Volume group fields can be mixed with either physical volume (and physical volume segment) fields or with logical volume (and logical volume segment) fields, but physical volume and logical volume fields cannot be mixed. For example, the following command will display one line of output for each physical volume.

```
# vgs -o +pv_name
VG      #PV #LV #SN Attr   VSize  VFree  PV
new_vg   3   1   0 wz--n- 51.42G 51.37G /dev/sdc1
new_vg   3   1   0 wz--n- 51.42G 51.37G /dev/sdd1
new_vg   3   1   0 wz--n- 51.42G 51.37G /dev/sdb1
```

4.8.2. Object Display Fields

This section provides a series of tables that list the information you can display about the LVM objects with the **pvs**, **vgs**, and **lvs** commands.

For convenience, a field name prefix can be dropped if it matches the default for the command. For example, with the **pvs** command, **name** means **pv_name**, but with the **vgs** command, **name** is interpreted as **vg_name**.

Executing the following command is the equivalent of executing **pvs -o pv_free**.

```
# pvs -o free
PFree
17.14G
17.09G
17.14G
```



NOTE

The number of characters in the attribute fields in **pvs**, **vgs**, and **lvs** output may increase in later releases. The existing character fields will not change position, but new fields may be added to the end. You should take this into account when writing scripts that search for particular attribute characters, searching for the character based on its relative position to the beginning of the field, but not for its relative position to the end of the field. For example, to search for the character **p** in the ninth bit of the **lv_attr** field, you could search for the string `"^/.....p/"`, but you should not search for the string `"/*p$/"`.

The pvs Command

Table 4.3, “The pvs Command Display Fields” lists the display arguments of the **pvs** command, along with the field name as it appears in the header display and a description of the field.

Table 4.3. The pvs Command Display Fields

Argument	Header	Description
dev_size	DevSize	The size of the underlying device on which the physical volume was created
pe_start	1st PE	Offset to the start of the first physical extent in the underlying device
pv_attr	Attr	Status of the physical volume: (a)llocatable or e(x)ported.
pv_fmt	Fmt	The metadata format of the physical volume (lvm2 or lvm1)
pv_free	PFree	The free space remaining on the physical volume
pv_name	PV	The physical volume name
pv_pe_alloc_count	Alloc	Number of used physical extents
pv_pe_count	PE	Number of physical extents

Argument	Header	Description
pvseg_size	SSize	The segment size of the physical volume
pvseg_start	Start	The starting physical extent of the physical volume segment
pv_size	PSize	The size of the physical volume
pv_tags	PV Tags	LVM tags attached to the physical volume
pv_used	Used	The amount of space currently used on the physical volume
pv_uuid	PV UUID	The UUID of the physical volume

The **pvs** command displays the following fields by default: **pv_name**, **vg_name**, **pv_fmt**, **pv_attr**, **pv_size**, **pv_free**. The display is sorted by **pv_name**.

```
# pvs
PV          VG      Fmt  Attr PSize  PFree
/dev/sdb1   new_vg  lvm2 a-   17.14G 17.14G
/dev/sdc1   new_vg  lvm2 a-   17.14G 17.09G
/dev/sdd1   new_vg  lvm2 a-   17.14G 17.13G
```

Using the **-v** argument with the **pvs** command adds the following fields to the default display: **dev_size**, **pv_uuid**.

```
# pvs -v
Scanning for physical volume names
PV          VG      Fmt  Attr PSize  PFree  DevSize PV UUID
/dev/sdb1   new_vg  lvm2 a-   17.14G 17.14G  17.14G onFF2w-1fLC-ughJ-D9eB-
M7iv-6XqA-dqGeXY
/dev/sdc1   new_vg  lvm2 a-   17.14G 17.09G  17.14G Joqlch-yWSj-kuEn-IdwM-
01S9-X08M-mcpsVe
/dev/sdd1   new_vg  lvm2 a-   17.14G 17.13G  17.14G yvfvZK-Cf31-j75k-dECm-
0RZ3-0dGW-tUqkCS
```

You can use the **--segments** argument of the **pvs** command to display information about each physical volume segment. A segment is a group of extents. A segment view can be useful if you want to see whether your logical volume is fragmented.

The **pvs --segments** command displays the following fields by default: **pv_name**, **vg_name**, **pv_fmt**, **pv_attr**, **pv_size**, **pv_free**, **pvseg_start**, **pvseg_size**. The display is sorted by **pv_name** and **pvseg_size** within the physical volume.

```
# pvs --segments
PV          VG          Fmt  Attr PSize  PFree  Start SSize
/dev/hda2   VolGroup00  lvm2 a-   37.16G 32.00M    0  1172
/dev/hda2   VolGroup00  lvm2 a-   37.16G 32.00M  1172   16
/dev/hda2   VolGroup00  lvm2 a-   37.16G 32.00M  1188    1
/dev/sda1   vg          lvm2 a-   17.14G 16.75G    0   26
/dev/sda1   vg          lvm2 a-   17.14G 16.75G   26   24
```

```

/dev/sda1 vg          lvm2 a-   17.14G 16.75G    50    26
/dev/sda1 vg          lvm2 a-   17.14G 16.75G    76    24
/dev/sda1 vg          lvm2 a-   17.14G 16.75G   100    26
/dev/sda1 vg          lvm2 a-   17.14G 16.75G   126    24
/dev/sda1 vg          lvm2 a-   17.14G 16.75G   150    22
/dev/sda1 vg          lvm2 a-   17.14G 16.75G   172  4217
/dev/sdb1 vg          lvm2 a-   17.14G 17.14G     0  4389
/dev/sdc1 vg          lvm2 a-   17.14G 17.14G     0  4389
/dev/sdd1 vg          lvm2 a-   17.14G 17.14G     0  4389
/dev/sde1 vg          lvm2 a-   17.14G 17.14G     0  4389
/dev/sdf1 vg          lvm2 a-   17.14G 17.14G     0  4389
/dev/sdg1 vg          lvm2 a-   17.14G 17.14G     0  4389

```

You can use the **pvs -a** command to see devices detected by LVM that have not been initialized as LVM physical volumes.

```

# pvs -a
PV                               VG      Fmt  Attr PSize  PFree
/dev/VolGroup00/LogVol01        --      --    0      0
/dev/new_vg/lvol0               --      --    0      0
/dev/ram                        --      --    0      0
/dev/ram0                       --      --    0      0
/dev/ram2                       --      --    0      0
/dev/ram3                       --      --    0      0
/dev/ram4                       --      --    0      0
/dev/ram5                       --      --    0      0
/dev/ram6                       --      --    0      0
/dev/root                       --      --    0      0
/dev/sda                       --      --    0      0
/dev/sdb                       --      --    0      0
/dev/sdb1                      new_vg  lvm2 a-   17.14G 17.14G
/dev/sdc                       --      --    0      0
/dev/sdc1                      new_vg  lvm2 a-   17.14G 17.09G
/dev/sdd                       --      --    0      0
/dev/sdd1                      new_vg  lvm2 a-   17.14G 17.14G

```

The vgs Command

Table 4.4, “vgs Display Fields” lists the display arguments of the **vgs** command, along with the field name as it appears in the header display and a description of the field.

Table 4.4. vgs Display Fields

Argument	Header	Description
lv_count	#LV	The number of logical volumes the volume group contains
max_lv	MaxLV	The maximum number of logical volumes allowed in the volume group (0 if unlimited)
max_pv	MaxPV	The maximum number of physical volumes allowed in the volume group (0 if unlimited)

Argument	Header	Description
pv_count	#PV	The number of physical volumes that define the volume group
snap_count	#SN	The number of snapshots the volume group contains
vg_attr	Attr	Status of the volume group: (w)riteable, (r)eadonly, resi(z)eable, e(x)ported, (p)artial and (c)lustered.
vg_extent_count	#Ext	The number of physical extents in the volume group
vg_extent_size	Ext	The size of the physical extents in the volume group
vg_fmt	Fmt	The metadata format of the volume group (lvm2 or lvm1)
vg_free	VFree	Size of the free space remaining in the volume group
vg_free_count	Free	Number of free physical extents in the volume group
vg_name	VG	The volume group name
vg_seqno	Seq	Number representing the revision of the volume group
vg_size	VSize	The size of the volume group
vg_sysid	SYS ID	LVM1 System ID
vg_tags	VG Tags	LVM tags attached to the volume group
vg_uuid	VG UUID	The UUID of the volume group

The **vg**s command displays the following fields by default: **vg_name**, **pv_count**, **lv_count**, **snap_count**, **vg_attr**, **vg_size**, **vg_free**. The display is sorted by **vg_name**.

```
# vgs
VG      #PV #LV #SN Attr   VSize  VFree
new_vg   3   1   1 wz--n- 51.42G 51.36G
```

Using the **-v** argument with the **vg**s command adds the following fields to the default display: **vg_extent_size**, **vg_uuid**.

```
# vgs -v
Finding all volume groups
Finding volume group "new_vg"
VG      Attr   Ext   #PV #LV #SN VSize  VFree  VG UUID
new_vg wz--n- 4.00M   3   1   1 51.42G 51.36G jxQJ0a-ZKk0-0pM0-0118-
nlw0-wwqd-fD5D32
```

The lvs Command

Table 4.5, “lvs Display Fields” lists the display arguments of the **lvs** command, along with the field name as it appears in the header display and a description of the field.



NOTE

In later releases of Red Hat Enterprise Linux, the output of the **lvs** command may differ, with additional fields in the output. The order of the fields, however, will remain the same and any additional fields will appear at the end of the display.

Table 4.5. lvs Display Fields

Argument	Header	Description
chunksize chunk_size	Chunk	Unit size in a snapshot volume
copy_percent	Copy%	The synchronization percentage of a mirrored logical volume; also used when physical extents are being moved with the pv_move command
devices	Devices	The underlying devices that make up the logical volume: the physical volumes, logical volumes, and start physical extents and logical extents
lv_ancestors	Ancestors	(Red Hat Enterprise Linux 7.2 and later) For thin pool snapshots, the ancestors of the logical volume
lv_descendants	Descendants	(Red Hat Enterprise Linux 7.2 and later) For thin pool snapshots, the descendants of the logical volume
lv_attr	Attr	The status of the logical volume. The logical volume attribute bits are as follows: Bit 1: Volume type: (m)irrored, (M)irrored without initial sync, (o)rigin, (O)rigin with merging snapshot, (r)aid, (R)aid without initial sync, (s)napshot, merging (S)napshot, (p)vmove, (v)irtual, mirror or raid (i)mage, mirror or raid (I)mage out-of-sync, mirror (l)og device, under (c)onversion, thin (V)olume, (t)hin pool, (T)hin pool data, raid or thin pool m(e)tadata or pool metadata spare, Bit 2: Permissions: (w)riteable, (r)ead-only, (R)ead-only activation of non-read-only volume Bit 3: Allocation policy: (a)nywhere, (c)ontiguous, (i)nherited, c(l)ing, (n)ormal. This is capitalized if the volume is currently locked against allocation changes, for example while executing the pvmove command.

Argument	Header	Description (m)inor
		Bit 5: State: (a)ctive, (s)uspended, (l)invalid snapshot, invalid (S)uspended snapshot, snapshot (m)erge failed, suspended snapshot (M)erge failed, mapped (d)evice present without tables, mapped device present with (i)nactive table Bit 6: device (o)pen Bit 7: Target type: (m)irror, (r)aid, (s)napshot, (t)hin, (u)nknown, (v)irtual. This groups logical volumes related to the same kernel target together. So, for example, mirror images, mirror logs as well as mirrors themselves appear as (m) if they use the original device-mapper mirror kernel driver, whereas the raid equivalents using the md raid kernel driver all appear as (r). Snapshots using the original device-mapper driver appear as (s), whereas snapshots of thin volumes using the thin provisioning driver appear as (t). Bit 8: Newly-allocated data blocks are overwritten with blocks of (z)eroes before use. Bit 9: Volume Health: (p)artial, (r)efresh needed, (m)ismatches exist, (w)rite-mostly. (p)artial signifies that one or more of the Physical Volumes this Logical Volume uses is missing from the system. (r)efresh signifies that one or more of the Physical Volumes this RAID Logical Volume uses had suffered a write error. The write error could be due to a temporary failure of that Physical Volume or an indication that it is failing. The device should be refreshed or replaced. (m)ismatches signifies that the RAID logical volume has portions of the array that are not coherent. Inconsistencies are discovered by initiating a check operation on a RAID logical volume. (The scrubbing operations, check and repair, can be performed on a RAID Logical Volume by means of the lvchange command.) (w)rite-mostly signifies the devices in a RAID 1 logical volume that have been marked write-mostly. Bit 10: s(k)ip activation: this volume is flagged to be skipped during activation.
lv_kernel_major	KMaj	Actual major device number of the logical volume (-1 if inactive)
lv_kernel_minor	KMIN	Actual minor device number of the logical volume (-1 if inactive)
lv_major	Maj	The persistent major device number of the logical volume (-1 if not specified)
lv_minor	Min	The persistent minor device number of the logical volume (-1 if not specified)

Argument	Header	Description
lv_name	LV	The name of the logical volume
lv_size	LSize	The size of the logical volume
lv_tags	LV Tags	LVM tags attached to the logical volume
lv_uuid	LV UUID	The UUID of the logical volume.
mirror_log	Log	Device on which the mirror log resides
modules	Modules	Corresponding kernel device-mapper target necessary to use this logical volume
move_pv	Move	Source physical volume of a temporary logical volume created with the pvmove command
origin	Origin	The origin device of a snapshot volume
regionsize region_size	Region	The unit size of a mirrored logical volume
seg_count	#Seg	The number of segments in the logical volume
seg_size	SSize	The size of the segments in the logical volume
seg_start	Start	Offset of the segment in the logical volume
seg_tags	Seg Tags	LVM tags attached to the segments of the logical volume
segtype	Type	The segment type of a logical volume (for example: mirror, striped, linear)
snap_percent	Snap%	Current percentage of a snapshot volume that is in use
stripes	#Str	Number of stripes or mirrors in a logical volume
stripesize stripe_size	Stripe	Unit size of the stripe in a striped logical volume

The **lvs** command displays the following fields by default: **lv_name**, **vg_name**, **lv_attr**, **lv_size**, **origin**, **snap_percent**, **move_pv**, **mirror_log**, **copy_percent**, **convert_lv**. The default display is sorted by **vg_name** and **lv_name** within the volume group.

```
# lvs
LV          VG      Attr   LSize   Origin Snap%   Move Log Copy%   Convert
lvol0       new_vg  owi-a- 52.00M
newvgsnap1  new_vg  swi-a-  8.00M  lvol0    0.20
```

Using the **-v** argument with the **lvs** command adds the following fields to the default display: **seg_count**, **lv_major**, **lv_minor**, **lv_kernel_major**, **lv_kernel_minor**, **lv_uuid**.

```
# lvs -v
Finding all logical volumes
LV          VG      #Seg Attr   LSize   Maj Min KMaj KMin Origin Snap%
Move Copy%  Log Convert LV UUID
lvol0       new_vg   1 owi-a- 52.00M   -1 -1 253  3
LBy1Tz-sr23-0jsI-LT03-nHLC-y8XW-EhCl78
newvgsnap1  new_vg   1 swi-a-  8.00M   -1 -1 253  5    lvol0    0.20
1ye10U-1cIu-o79k-20h2-ZGF0-qCJm-CfbsIx
```

You can use the **--segments** argument of the **lvs** command to display information with default columns that emphasize the segment information. When you use the **segments** argument, the **seg** prefix is optional. The **lvs --segments** command displays the following fields by default: **lv_name**, **vg_name**, **lv_attr**, **stripes**, **segtype**, **seg_size**. The default display is sorted by **vg_name**, **lv_name** within the volume group, and **seg_start** within the logical volume. If the logical volumes were fragmented, the output from this command would show that.

```
# lvs --segments
LV          VG          Attr   #Str Type   SSize
LogVol00    VolGroup00 -wi-ao    1 linear 36.62G
LogVol01    VolGroup00 -wi-ao    1 linear 512.00M
lv          vg          -wi-a-    1 linear 104.00M
lv          vg          -wi-a-    1 linear 104.00M
lv          vg          -wi-a-    1 linear 104.00M
lv          vg          -wi-a-    1 linear  88.00M
```

Using the **-v** argument with the **lvs --segments** command adds the following fields to the default display: **seg_start**, **stripesize**, **chunksize**.

```
# lvs -v --segments
Finding all logical volumes
LV          VG      Attr   Start SSize   #Str Type   Stripe Chunk
lvol0       new_vg  owi-a-    0 52.00M    1 linear    0    0
newvgsnap1  new_vg  swi-a-    0  8.00M    1 linear    0  8.00K
```

The following example shows the default output of the **lvs** command on a system with one logical volume configured, followed by the default output of the **lvs** command with the **segments** argument specified.

```
# lvs
LV      VG      Attr   LSize   Origin Snap%   Move Log Copy%
lvol0  new_vg  -wi-a- 52.00M
```

```
# lvs --segments
LV      VG      Attr      #Str Type   SSize
lvol0   new_vg   -wi-a-    1 linear 52.00M
```

4.8.3. Sorting LVM Reports

Normally the entire output of the **lvs**, **vgs**, or **pvs** command has to be generated and stored internally before it can be sorted and columns aligned correctly. You can specify the **--unbuffered** argument to display unsorted output as soon as it is generated.

To specify an alternative ordered list of columns to sort on, use the **-O** argument of any of the reporting commands. It is not necessary to include these fields within the output itself.

The following example shows the output of the **pvs** command that displays the physical volume name, size, and free space.

```
# pvs -o pv_name,pv_size,pv_free
PV          PSize  PFree
/dev/sdb1   17.14G 17.14G
/dev/sdc1   17.14G 17.09G
/dev/sdd1   17.14G 17.14G
```

The following example shows the same output, sorted by the free space field.

```
# pvs -o pv_name,pv_size,pv_free -O pv_free
PV          PSize  PFree
/dev/sdc1   17.14G 17.09G
/dev/sdd1   17.14G 17.14G
/dev/sdb1   17.14G 17.14G
```

The following example shows that you do not need to display the field on which you are sorting.

```
# pvs -o pv_name,pv_size -O pv_free
PV          PSize
/dev/sdc1   17.14G
/dev/sdd1   17.14G
/dev/sdb1   17.14G
```

To display a reverse sort, precede a field you specify after the **-O** argument with the **-** character.

```
# pvs -o pv_name,pv_size,pv_free -O -pv_free
PV          PSize  PFree
/dev/sdd1   17.14G 17.14G
/dev/sdb1   17.14G 17.14G
/dev/sdc1   17.14G 17.09G
```

4.8.4. Specifying Units

To specify the unit for the LVM report display, use the **--units** argument of the report command. You can specify (b)ytes, (k)ilobytes, (m)egabytes, (g)igabytes, (t)erabytes, (e)xabytes, (p)etabytes, and (h)uman-readable. The default display is human-readable. You can override the default by setting the **units** parameter in the **global** section of the **lvm.conf** file.

The following example specifies the output of the **pvs** command in megabytes rather than the default gigabytes.

```
# pvs --units m
PV          VG      Fmt  Attr PSize      PFree
/dev/sda1   lvm2  --   17555.40M 17555.40M
/dev/sdb1   new_vg lvm2  a-   17552.00M 17552.00M
/dev/sdc1   new_vg lvm2  a-   17552.00M 17500.00M
/dev/sdd1   new_vg lvm2  a-   17552.00M 17552.00M
```

By default, units are displayed in powers of 2 (multiples of 1024). You can specify that units be displayed in multiples of 1000 by capitalizing the unit specification (B, K, M, G, T, H).

The following command displays the output as a multiple of 1024, the default behavior.

```
# pvs
PV          VG      Fmt  Attr PSize  PFree
/dev/sdb1   new_vg lvm2  a-   17.14G 17.14G
/dev/sdc1   new_vg lvm2  a-   17.14G 17.09G
/dev/sdd1   new_vg lvm2  a-   17.14G 17.14G
```

The following command displays the output as a multiple of 1000.

```
# pvs --units G
PV          VG      Fmt  Attr PSize  PFree
/dev/sdb1   new_vg lvm2  a-   18.40G 18.40G
/dev/sdc1   new_vg lvm2  a-   18.40G 18.35G
/dev/sdd1   new_vg lvm2  a-   18.40G 18.40G
```

You can also specify (s)ectors (defined as 512 bytes) or custom units.

The following example displays the output of the **pvs** command as a number of sectors.

```
# pvs --units s
PV          VG      Fmt  Attr PSize      PFree
/dev/sdb1   new_vg lvm2  a-   35946496S 35946496S
/dev/sdc1   new_vg lvm2  a-   35946496S 35840000S
/dev/sdd1   new_vg lvm2  a-   35946496S 35946496S
```

The following example displays the output of the **pvs** command in units of 4 MB.

```
# pvs --units 4m
PV          VG      Fmt  Attr PSize      PFree
/dev/sdb1   new_vg lvm2  a-   4388.00U 4388.00U
/dev/sdc1   new_vg lvm2  a-   4388.00U 4375.00U
/dev/sdd1   new_vg lvm2  a-   4388.00U 4388.00U
```

4.8.5. JSON Format Output (Red Hat Enterprise Linux 7.3 and later)

As of Red Hat Enterprise Linux 7.3, you can use the **--reportformat** option of the LVM display commands to display the output in JSON format.

The following example shows the output of the **lvs** in standard default format.

```
# lvs
  LV          VG          Attr          LSize    Pool Origin Data%  Meta%
Move Log Cpy%Sync Convert
  my_raid my_vg          Rwi-a-r--- 12.00m
100.00
  root      rhel_host-075 -wi-ao---- 6.67g
  swap      rhel_host-075 -wi-ao---- 820.00m
```

The following command shows the output of the same LVM configuration when you specify JSON format.

```
# lvs --reportformat json
{
  "report": [
    {
      "lv": [
        {
          "lv_name": "my_raid", "vg_name": "my_vg",
          "lv_attr": "Rwi-a-r---", "lv_size": "12.00m", "pool_lv": "", "origin": "",
          "data_percent": "", "metadata_percent": "", "move_pv": "", "mirror_log": "",
          "copy_percent": "100.00", "convert_lv": ""},
        {
          "lv_name": "root", "vg_name": "rhel_host-075",
          "lv_attr": "-wi-ao----", "lv_size": "6.67g", "pool_lv": "", "origin": "",
          "data_percent": "", "metadata_percent": "", "move_pv": "", "mirror_log": "",
          "copy_percent": "", "convert_lv": ""},
        {
          "lv_name": "swap", "vg_name": "rhel_host-075",
          "lv_attr": "-wi-ao----", "lv_size": "820.00m", "pool_lv": "", "origin": "",
          "data_percent": "", "metadata_percent": "", "move_pv": "", "mirror_log": "",
          "copy_percent": "", "convert_lv": ""}
      ]
    }
  ]
}
```

You can also set the report format as a configuration option in the `/etc/lvm/lvm.conf` file, using the **output_format** setting. The **--reportformat** setting of the command line, however, takes precedence over this setting.

4.8.6. Command Log Reporting (Red Hat Enterprise Linux 7.3 and later)

As of Red Hat Enterprise Linux 7.3, both report-oriented and processing-oriented LVM commands can report the command log if this is enabled with the **log/report_command_log** configuration setting. You can determine the set of fields to display and to sort by for this report.

The following examples configures LVM to generate a complete log report for LVM commands. In this example, you can see that both logical volumes **lvvol0** and **lvvol1** were successfully processed, as was the volume group **VG** that contains the volumes.

```
# lvmconfig --type full log/command_log_selection
command_log_selection="all"

# lvs
Logical Volume
=====
LV          LSize Cpy%Sync
```

```
lvvol1 4.00m 100.00
lvvol0 4.00m
```

```
Command Log
=====
```

Seq	LogType	Context	ObjType	ObjName	ObjGrp	Msg	Errno	RetCode
1	status	processing	lv	lvvol0	vg	success	0	1
2	status	processing	lv	lvvol1	vg	success	0	1
3	status	processing	vg	vg		success	0	1

```
# lvchange -an vg/lvol1
```

```
Command Log
=====
```

Seq	LogType	Context	ObjType	ObjName	ObjGrp	Msg	Errno	RetCode
1	status	processing	lv	lvvol1	vg	success	0	1
2	status	processing	vg	vg		success	0	1

For further information on configuring LVM reports and command logs, see the **lvmreport** man page.

CHAPTER 5. LVM CONFIGURATION EXAMPLES

This chapter provides some basic LVM configuration examples.

5.1. CREATING AN LVM LOGICAL VOLUME ON THREE DISKS

This example procedure creates an LVM logical volume called **new_logical_volume** that consists of the disks at **/dev/sda1**, **/dev/sdb1**, and **/dev/sdc1**.

1. To use disks in a volume group, label them as LVM physical volumes with the **pvcreate** command.



WARNING

This command destroys any data on **/dev/sda1**, **/dev/sdb1**, and **/dev/sdc1**.

```
# pvcreate /dev/sda1 /dev/sdb1 /dev/sdc1
Physical volume "/dev/sda1" successfully created
Physical volume "/dev/sdb1" successfully created
Physical volume "/dev/sdc1" successfully created
```

2. Create the a volume group that consists of the LVM physical volumes you have created. The following command creates the volume group **new_vol_group**.

```
# vgcreate new_vol_group /dev/sda1 /dev/sdb1 /dev/sdc1
Volume group "new_vol_group" successfully created
```

You can use the **vgs** command to display the attributes of the new volume group.

```
# vgs
VG                #PV #LV #SN Attr   VSize  VFree
new_vol_group     3   0   0 wz--n- 51.45G 51.45G
```

3. Create the logical volume from the volume group you have created. The following command creates the logical volume **new_logical_volume** from the volume group **new_vol_group**. This example creates a logical volume that uses 2 gigabytes of the volume group.

```
# lvcreate -L 2G -n new_logical_volume new_vol_group
Logical volume "new_logical_volume" created
```

4. Create a file system on the logical volume. The following command creates a GFS2 file system on the logical volume.

```
# mkfs.gfs2 -p lock_nolock -j 1
/dev/new_vol_group/new_logical_volume
This will destroy any data on /dev/new_vol_group/new_logical_volume.
```

```
Are you sure you want to proceed? [y/n] y

Device:                      /dev/new_vol_group/new_logical_volume
Blocksize:                   4096
Filesystem Size:             491460
Journals:                    1
Resource Groups:             8
Locking Protocol:            lock_nolock
Lock Table:

Syncing...
All Done
```

The following commands mount the logical volume and report the file system disk space usage.

```
# mount /dev/new_vol_group/new_logical_volume /mnt
# df
Filesystem          1K-blocks      Used Available Use% Mounted on
/dev/new_vol_group/new_logical_volume
                    1965840         20   1965820   1% /mnt
```

5.2. CREATING A STRIPED LOGICAL VOLUME

This example procedure creates an LVM striped logical volume called **striped_logical_volume** that stripes data across the disks at **/dev/sda1**, **/dev/sdb1**, and **/dev/sdc1**.

1. Label the disks you will use in the volume group as LVM physical volumes with the **pvcreeate** command.



WARNING

This command destroys any data on **/dev/sda1**, **/dev/sdb1**, and **/dev/sdc1**.

```
# pvcreate /dev/sda1 /dev/sdb1 /dev/sdc1
Physical volume "/dev/sda1" successfully created
Physical volume "/dev/sdb1" successfully created
Physical volume "/dev/sdc1" successfully created
```

2. Create the volume group **volgroup01**. The following command creates the volume group **volgroup01**.

```
# vgcreate volgroup01 /dev/sda1 /dev/sdb1 /dev/sdc1
Volume group "volgroup01" successfully created
```

You can use the **vgs** command to display the attributes of the new volume group.

■

```
# vgs
VG                #PV #LV #SN Attr   VSize  VFree
volgroup01        3   0   0 wz--n- 51.45G 51.45G
```

3. Create a striped logical volume from the volume group you have created. The following command creates the striped logical volume **striped_logical_volume** from the volume group **volgroup01**. This example creates a logical volume that is 2 gigabytes in size, with three stripes and a stripe size of 4 kilobytes.

```
# lvcreate -i 3 -I 4 -L 2G -n striped_logical_volume volgroup01
Rounding size (512 extents) up to stripe boundary size (513
extents)
Logical volume "striped_logical_volume" created
```

4. Create a file system on the striped logical volume. The following command creates a GFS2 file system on the logical volume.

```
# mkfs.gfs2 -p lock_nolock -j 1
/dev/volgroup01/striped_logical_volume
This will destroy any data on
/dev/volgroup01/striped_logical_volume.

Are you sure you want to proceed? [y/n] y

Device:                               /dev/volgroup01/striped_logical_volume
Blocksize:                            4096
Filesystem Size:                      492484
Journals:                             1
Resource Groups:                      8
Locking Protocol:                    lock_nolock
Lock Table:

Syncing...
All Done
```

The following commands mount the logical volume and report the file system disk space usage.

```
# mount /dev/volgroup01/striped_logical_volume /mnt
# df
Filesystem            1K-blocks      Used Available Use% Mounted on
/dev/mapper/VolGroup00-LogVol00
13902624    1656776    11528232    13% /
/dev/hda1        101086      10787      85080    12% /boot
tmpfs            127880         0      127880     0% /dev/shm
/dev/volgroup01/striped_logical_volume
1969936         20    1969916     1% /mnt
```

5.3. SPLITTING A VOLUME GROUP

In this example procedure, an existing volume group consists of three physical volumes. If there is enough unused space on the physical volumes, a new volume group can be created without adding new disks.

In the initial set up, the logical volume **mylv** is carved from the volume group **myvol**, which in turn consists of the three physical volumes, **/dev/sda1**, **/dev/sdb1**, and **/dev/sdc1**.

After completing this procedure, the volume group **myvg** will consist of **/dev/sda1** and **/dev/sdb1**. A second volume group, **yourvg**, will consist of **/dev/sdc1**.

1. Use the **pvscan** command to determine how much free space is currently available in the volume group.

```
# pvscan
PV /dev/sda1  VG myvg   lvm2 [17.15 GB / 0    free]
PV /dev/sdb1  VG myvg   lvm2 [17.15 GB / 12.15 GB free]
PV /dev/sdc1  VG myvg   lvm2 [17.15 GB / 15.80 GB free]
Total: 3 [51.45 GB] / in use: 3 [51.45 GB] / in no VG: 0 [0    ]
```

2. Move all the used physical extents in **/dev/sdc1** to **/dev/sdb1** with the **pvmove** command. The **pvmove** command can take a long time to execute.

```
# pvmove /dev/sdc1 /dev/sdb1
/dev/sdc1: Moved: 14.7%
/dev/sdc1: Moved: 30.3%
/dev/sdc1: Moved: 45.7%
/dev/sdc1: Moved: 61.0%
/dev/sdc1: Moved: 76.6%
/dev/sdc1: Moved: 92.2%
/dev/sdc1: Moved: 100.0%
```

After moving the data, you can see that all of the space on **/dev/sdc1** is free.

```
# pvscan
PV /dev/sda1  VG myvg   lvm2 [17.15 GB / 0    free]
PV /dev/sdb1  VG myvg   lvm2 [17.15 GB / 10.80 GB free]
PV /dev/sdc1  VG myvg   lvm2 [17.15 GB / 17.15 GB free]
Total: 3 [51.45 GB] / in use: 3 [51.45 GB] / in no VG: 0 [0    ]
```

3. To create the new volume group **yourvg**, use the **vgsplit** command to split the volume group **myvg**.

Before you can split the volume group, the logical volume must be inactive. If the file system is mounted, you must unmount the file system before deactivating the logical volume.

Deactivate the logical volumes with the **lvchange** command or the **vgchange** command. The following command deactivates the logical volume **mylv** and then splits the volume group **yourvg** from the volume group **myvg**, moving the physical volume **/dev/sdc1** into the new volume group **yourvg**.

```
# lvchange -a n /dev/myvg/mylv
# vgsplit myvg yourvg /dev/sdc1
Volume group "yourvg" successfully split from "myvg"
```

You can use the **vgs** command to see the attributes of the two volume groups.

```
# vgs
```

```

VG      #PV #LV #SN Attr   VSize  VFree
myvg    2   1   0 wz--n- 34.30G 10.80G
yourvg  1   0   0 wz--n- 17.15G 17.15G

```

- After creating the new volume group, create the new logical volume **yourlv**.

```

# lvcreate -L 5G -n yourlv yourvg
Logical volume "yourlv" created

```

- Create a file system on the new logical volume and mount it.

```

# mkfs.gfs2 -p lock_nolock -j 1 /dev/yourvg/yourlv
This will destroy any data on /dev/yourvg/yourlv.

Are you sure you want to proceed? [y/n] y

Device:                               /dev/yourvg/yourlv
Blocksize:                           4096
Filesystem Size:                      1277816
Journals:                             1
Resource Groups:                      20
Locking Protocol:                     lock_nolock
Lock Table:

Syncing...
All Done

# mount /dev/yourvg/yourlv /mnt

```

- Since you had to deactivate the logical volume **mylv**, you need to activate it again before you can mount it.

```

# lvchange -a y /dev/myvg/mylv

# mount /dev/myvg/mylv /mnt
# df

```

Filesystem	1K-blocks	Used	Available	Use%	Mounted on
/dev/yourvg/yourlv	24507776	32	24507744	1%	/mnt
/dev/myvg/mylv	24507776	32	24507744	1%	/mnt

5.4. REMOVING A DISK FROM A LOGICAL VOLUME

These example procedures show how you can remove a disk from an existing logical volume, either to replace the disk or to use the disk as part of a different volume. In order to remove a disk, you must first move the extents on the LVM physical volume to a different disk or set of disks.

5.4.1. Moving Extents to Existing Physical Volumes

In this example, the logical volume is distributed across four physical volumes in the volume group **myvg**.

```

# pvs -o+pv_used
PV      VG      Fmt  Attr  PSize  PFree  Used
/dev/sda1 myvg  lvm2 a-   17.15G 12.15G  5.00G

```

```

/dev/sdb1  myvg lvm2 a-   17.15G 12.15G  5.00G
/dev/sdc1  myvg lvm2 a-   17.15G 12.15G  5.00G
/dev/sdd1  myvg lvm2 a-   17.15G  2.15G 15.00G

```

This examples moves the extents off of **/dev/sdb1** so that it can be removed from the volume group.

1. If there are enough free extents on the other physical volumes in the volume group, you can execute the **pvmove** command on the device you want to remove with no other options and the extents will be distributed to the other devices.

```

# pvmove /dev/sdb1
/dev/sdb1: Moved: 2.0%
...
/dev/sdb1: Moved: 79.2%
...
/dev/sdb1: Moved: 100.0%

```

After the **pvmove** command has finished executing, the distribution of extents is as follows:

```

# pvs -o+pv_used
PV          VG   Fmt  Attr PSize  PFree  Used
/dev/sda1   myvg lvm2 a-   17.15G  7.15G 10.00G
/dev/sdb1   myvg lvm2 a-   17.15G 17.15G    0
/dev/sdc1   myvg lvm2 a-   17.15G 12.15G  5.00G
/dev/sdd1   myvg lvm2 a-   17.15G  2.15G 15.00G

```

2. Use the **vgreduce** command to remove the physical volume **/dev/sdb1** from the volume group.

```

# vgreduce myvg /dev/sdb1
Removed "/dev/sdb1" from volume group "myvg"
# pvs
PV          VG   Fmt  Attr PSize  PFree
/dev/sda1   myvg lvm2 a-   17.15G  7.15G
/dev/sdb1           lvm2 --   17.15G 17.15G
/dev/sdc1   myvg lvm2 a-   17.15G 12.15G
/dev/sdd1   myvg lvm2 a-   17.15G  2.15G

```

The disk can now be physically removed or allocated to other users.

5.4.2. Moving Extents to a New Disk

In this example, the logical volume is distributed across three physical volumes in the volume group **myvg** as follows:

```

# pvs -o+pv_used
PV          VG   Fmt  Attr PSize  PFree  Used
/dev/sda1   myvg lvm2 a-   17.15G  7.15G 10.00G
/dev/sdb1   myvg lvm2 a-   17.15G 15.15G  2.00G
/dev/sdc1   myvg lvm2 a-   17.15G 15.15G  2.00G

```

This example procedure moves the extents of **/dev/sdb1** to a new device, **/dev/sdd1**.

1. Create a new physical volume from **/dev/sdd1**.

```
# pvcreate /dev/sdd1
Physical volume "/dev/sdd1" successfully created
```

2. Add the new physical volume **/dev/sdd1** to the existing volume group **myvg**.

```
# vgextend myvg /dev/sdd1
Volume group "myvg" successfully extended
# pvs -o+pv_used
PV          VG      Fmt  Attr PSize  PFree  Used
/dev/sda1   myvg   lvm2 a-   17.15G  7.15G  10.00G
/dev/sdb1   myvg   lvm2 a-   17.15G  15.15G  2.00G
/dev/sdc1   myvg   lvm2 a-   17.15G  15.15G  2.00G
/dev/sdd1   myvg   lvm2 a-   17.15G  17.15G   0
```

3. Use the **pvmove** command to move the data from **/dev/sdb1** to **/dev/sdd1**.

```
# pvmove /dev/sdb1 /dev/sdd1
/dev/sdb1: Moved: 10.0%
...
/dev/sdb1: Moved: 79.7%
...
/dev/sdb1: Moved: 100.0%

# pvs -o+pv_used
PV          VG      Fmt  Attr PSize  PFree  Used
/dev/sda1   myvg   lvm2 a-   17.15G  7.15G  10.00G
/dev/sdb1   myvg   lvm2 a-   17.15G  17.15G   0
/dev/sdc1   myvg   lvm2 a-   17.15G  15.15G  2.00G
/dev/sdd1   myvg   lvm2 a-   17.15G  15.15G  2.00G
```

4. After you have moved the data off **/dev/sdb1**, you can remove it from the volume group.

```
# vgreduce myvg /dev/sdb1
Removed "/dev/sdb1" from volume group "myvg"
```

You can now reallocate the disk to another volume group or remove the disk from the system.

5.5. CREATING A MIRRORED LVM LOGICAL VOLUME IN A CLUSTER

Creating a mirrored LVM logical volume in a cluster requires the same commands and procedures as creating a mirrored LVM logical volume on a single node with a segment type of **mirror**. However, in order to create a mirrored LVM volume in a cluster:

- The cluster and cluster mirror infrastructure must be running
- The cluster must be quorate
- The locking type in the **lvm.conf** file must be set correctly to enable cluster locking and the **use_lvmetad** setting should be 0. Note, however, that in Red Hat Enterprise Linux 7 the **ocf:heartbeat:clvm** Pacemaker resource agent itself, as part of the start procedure, performs these tasks.

In Red Hat Enterprise Linux 7, clusters are managed through Pacemaker. Clustered LVM logical volumes are supported only in conjunction with Pacemaker clusters, and must be configured as cluster resources.

The following procedure creates a mirrored LVM volume in a cluster.

1. Install the cluster software and LVM packages, start the cluster software, and create the cluster. You must configure fencing for the cluster. The document *High Availability Add-On Administration* provides a sample procedure for creating a cluster and configuring fencing for the nodes in the cluster. The document *High Availability Add-On Reference* provides more detailed information about the components of cluster configuration.
2. In order to create a mirrored logical volume that is shared by all of the nodes in a cluster, the locking type must be set correctly in the **lvm.conf** file in every node of the cluster. By default, the locking type is set to local. To change this, execute the following command in each node of the cluster to enable clustered locking:

```
# /sbin/lvmconf --enable-cluster
```

3. Set up a **dlm** resource for the cluster. You create the resource as a cloned resource so that it will run on every node in the cluster.

```
# pcs resource create dlm ocf:pacemaker:controld op monitor
interval=30s on-fail=fence clone interleave=true ordered=true
```

4. Configure **clvmd** as a cluster resource. Just as for the **dlm** resource, you create the resource as a cloned resource so that it will run on every node in the cluster. Note that you must set the **with_cmirrord=true** parameter to enable the **cmirrord** daemon on all of the nodes that **clvmd** runs on.

```
# pcs resource create clvmd ocf:heartbeat:clvm with_cmirrord=true op
monitor interval=30s on-fail=fence clone interleave=true
ordered=true
```

If you have already configured a **clvmd** resource but did not specify the **with_cmirrord=true** parameter, you can update the resource to include the parameter with the following command.

```
# pcs resource update clvmd with_cmirrord=true
```

5. Set up **clvmd** and **dlm** dependency and start up order. **clvmd** must start after **dlm** and must run on the same node as **dlm**.

```
# pcs constraint order start dlm-clone then clvmd-clone
# pcs constraint colocation add clvmd-clone with dlm-clone
```

6. Create the mirror. The first step is creating the physical volumes. The following commands create three physical volumes. Two of the physical volumes will be used for the legs of the mirror, and the third physical volume will contain the mirror log.

```
# pvcreate /dev/sdb1
Physical volume "/dev/sdb1" successfully created
# pvcreate /dev/sdc1
```

```
Physical volume "/dev/sdc1" successfully created
# pvcreate /dev/sdd1
Physical volume "/dev/sdd1" successfully created
```

7. Create the volume group. This example creates a volume group **vg001** that consists of the three physical volumes that were created in the previous step.

```
# vgcreate vg001 /dev/sdb1 /dev/sdc1 /dev/sdd1
Clustered volume group "vg001" successfully created
```

Note that the output of the **vgcreate** command indicates that the volume group is clustered. You can verify that a volume group is clustered with the **vgs** command, which will show the volume group's attributes. If a volume group is clustered, it will show a **c** attribute.

```
# vgs vg001
VG      #PV #LV #SN Attr   VSize  VFree
vg001   3   0   0 wz--nc 68.97G 68.97G
```

8. Create the mirrored logical volume. This example creates the logical volume **mirrorlv** from the volume group **vg001**. This volume has one mirror leg. This example specifies which extents of the physical volume will be used for the logical volume.

```
# lvcreate --type mirror -l 1000 -m 1 vg001 -n mirrorlv /dev/sdb1:1-
1000 /dev/sdc1:1-1000 /dev/sdd1:0
Logical volume "mirrorlv" created
```

You can use the **lvs** command to display the progress of the mirror creation. The following example shows that the mirror is 47% synced, then 91% synced, then 100% synced when the mirror is complete.

```
# lvs vg001/mirrorlv
LV      VG      Attr      LSize Origin Snap%   Move Log
Copy%   Convert
mirrorlv vg001   mwi-a- 3.91G                vg001_mlog
47.00
# lvs vg001/mirrorlv
LV      VG      Attr      LSize Origin Snap%   Move Log
Copy%   Convert
mirrorlv vg001   mwi-a- 3.91G                vg001_mlog
91.00
# lvs vg001/mirrorlv
LV      VG      Attr      LSize Origin Snap%   Move Log
Copy%   Convert
mirrorlv vg001   mwi-a- 3.91G                vg001_mlog
100.00
```

The completion of the mirror is noted in the system log:

```
May 10 14:52:52 doc-07 [19402]: Monitoring mirror device vg001-
mirrorlv for events
May 10 14:55:00 doc-07 lvm[19402]: vg001-mirrorlv is now in-sync
```

9. You can use the **lvs** command with the **-o +devices** options to display the configuration of the mirror, including which devices make up the mirror legs. You can see that the logical volume in this example is composed of two linear images and one log.

```
# lvs -a -o +devices
LV          VG      Attr   LSize  Origin Snap%  Move
Log          Copy%   Convert Devices
mirrorlv     vg001    mwi-a- 3.91G
mirrorlv_mlog 100.00
mirrorlv_mimage_0(0),mirrorlv_mimage_1(0)
[mirrorlv_mimage_0] vg001    iwi-ao 3.91G
/dev/sdb1(1)
[mirrorlv_mimage_1] vg001    iwi-ao 3.91G
/dev/sdc1(1)
[mirrorlv_mlog]    vg001    lwi-ao 4.00M
/dev/sdd1(0)
```

You can use the **seg_pe_ranges** option of the **lvs** to display the data layout. You can use this option to verify that your layout is properly redundant. The output of this command displays PE ranges in the same format that the **lvcreate** and **lvresize** commands take as input.

```
# lvs -a -o +seg_pe_ranges --segments
PE Ranges
mirrorlv_mimage_0:0-999 mirrorlv_mimage_1:0-999
/dev/sdb1:1-1000
/dev/sdc1:1-1000
/dev/sdd1:0-0
```



NOTE

For information on recovering from the failure of one of the legs of an LVM mirrored volume, see [Section 6.3, “Recovering from LVM Mirror Failure”](#).

CHAPTER 6. LVM TROUBLESHOOTING

This chapter provides instructions for troubleshooting a variety of LVM issues.

6.1. TROUBLESHOOTING DIAGNOSTICS

If a command is not working as expected, you can gather diagnostics in the following ways:

- Use the **-v**, **-vv**, **-vvv**, or **-vvvv** argument of any command for increasingly verbose levels of output.
- If the problem is related to the logical volume activation, set **activation = 1** in the **log** section of the configuration file and run the command with the **-vvvv** argument. After you have finished examining this output be sure to reset this parameter to 0, to avoid possible problems with the machine locking during low memory situations.
- Run the **lvmdump** command, which provides an information dump for diagnostic purposes. For information, see the **lvmdump(8)** man page.
- Execute the **lvs -v**, **pvs -a**, or **dmsetup info -c** command for additional system information.
- Examine the last backup of the metadata in the **/etc/lvm/backup** file and archived versions in the **/etc/lvm/archive** file.
- Check the current configuration information by running the **lvmconfig** command.
- Check the **.cache** file in the **/etc/lvm** directory for a record of which devices have physical volumes on them.

6.2. DISPLAYING INFORMATION ON FAILED DEVICES

You can use the **-P** argument of the **lvs** or **vgs** command to display information about a failed volume that would otherwise not appear in the output. This argument permits some operations even though the metadata is not completely consistent internally. For example, if one of the devices that made up the volume group **vg** failed, the **vgs** command might show the following output.

```
# vgs -o +devices
Volume group "vg" not found
```

If you specify the **-P** argument of the **vgs** command, the volume group is still unusable but you can see more information about the failed device.

```
# vgs -P -o +devices
Partial mode. Incomplete volume groups will be activated read-only.
VG    #PV #LV #SN Attr   VSize VFree Devices
vg     9  2   0 rz-pn- 2.11T 2.07T unknown device(0)
vg     9  2   0 rz-pn- 2.11T 2.07T unknown device(5120),/dev/sda1(0)
```

In this example, the failed device caused both a linear and a striped logical volume in the volume group to fail. The **lvs** command without the **-P** argument shows the following output.

```
# lvs -a -o +devices
Volume group "vg" not found
```

Using the **-P** argument shows the logical volumes that have failed.

```
# lvs -P -a -o +devices
Partial mode. Incomplete volume groups will be activated read-only.
LV      VG      Attr   LSize   Origin Snap%   Move Log Copy%  Devices
linear  vg      -wi-a- 20.00G                                     unknown
device(0)
stripe  vg      -wi-a- 20.00G                                     unknown
device(5120),/dev/sda1(0)
```

The following examples show the output of the **pvs** and **lvs** commands with the **-P** argument specified when a leg of a mirrored logical volume has failed.

```
# vgs -a -o +devices -P
Partial mode. Incomplete volume groups will be activated read-only.
VG      #PV #LV #SN Attr   VSize VFree Devices
corey    4   4   0 rz-pnc 1.58T 1.34T
my_mirror_mimage_0(0),my_mirror_mimage_1(0)
corey    4   4   0 rz-pnc 1.58T 1.34T /dev/sdd1(0)
corey    4   4   0 rz-pnc 1.58T 1.34T unknown device(0)
corey    4   4   0 rz-pnc 1.58T 1.34T /dev/sdb1(0)
```

```
# lvs -a -o +devices -P
Partial mode. Incomplete volume groups will be activated read-only.
LV      VG      Attr   LSize   Origin Snap%   Move Log Copy%  Devices
my_mirror      corey mwi-a- 120.00G
my_mirror_mlog 1.95 my_mirror_mimage_0(0),my_mirror_mimage_1(0)
[my_mirror_mimage_0] corey iwi-ao 120.00G
unknown device(0)
[my_mirror_mimage_1] corey iwi-ao 120.00G
/dev/sdb1(0)
[my_mirror_mlog]      corey lwi-ao   4.00M
/dev/sdd1(0)
```

6.3. RECOVERING FROM LVM MIRROR FAILURE

This section provides an example of recovering from a situation where one leg of an LVM mirrored volume fails because the underlying device for a physical volume goes down and the **mirror_log_fault_policy** parameter is set to **remove**. This requires that you manually rebuild the mirror. For information on setting the **mirror_log_fault_policy** parameter, see [Section 4.4.4.1, “Mirrored Logical Volume Failure Policy”](#).

When a mirror leg fails, LVM converts the mirrored volume into a linear volume, which continues to operate as before but without the mirrored redundancy. At that point, you can add a new disk device to the system to use as a replacement physical device and rebuild the mirror.

The following command creates the physical volumes which will be used for the mirror.

```
# pvcreate /dev/sd[abcdefgh][12]
```

```
Physical volume "/dev/sda1" successfully created
Physical volume "/dev/sda2" successfully created
Physical volume "/dev/sdb1" successfully created
Physical volume "/dev/sdb2" successfully created
Physical volume "/dev/sdc1" successfully created
Physical volume "/dev/sdc2" successfully created
Physical volume "/dev/sdd1" successfully created
Physical volume "/dev/sdd2" successfully created
Physical volume "/dev/sde1" successfully created
Physical volume "/dev/sde2" successfully created
Physical volume "/dev/sdf1" successfully created
Physical volume "/dev/sdf2" successfully created
Physical volume "/dev/sgd1" successfully created
Physical volume "/dev/sgd2" successfully created
Physical volume "/dev/sdh1" successfully created
Physical volume "/dev/sdh2" successfully created
```

The following commands creates the volume group **vg** and the mirrored volume **groupfs**.

```
# vgcreate vg /dev/sd[abcdefgh][12]
Volume group "vg" successfully created
# lvcreate -L 750M -n groupfs -m 1 vg /dev/sda1 /dev/sdb1 /dev/sdc1
Rounding up size to full physical extent 752.00 MB
Logical volume "groupfs" created
```

You can use the **lvs** command to verify the layout of the mirrored volume and the underlying devices for the mirror leg and the mirror log. Note that in the first example the mirror is not yet completely synced; you should wait until the **Copy%** field displays 100.00 before continuing.

```
# lvs -a -o +devices
LV          VG      Attr      LSize   Origin Snap%   Move Log
Copy% Devices
groupfs      vg      mwi-a-    752.00M                      groupfs_mlog
21.28 groupfs_mimage_0(0),groupfs_mimage_1(0)
[groupfs_mimage_0] vg      iwi-ao    752.00M
/dev/sda1(0)
[groupfs_mimage_1] vg      iwi-ao    752.00M
/dev/sdb1(0)
[groupfs_mlog]   vg      lwi-ao     4.00M
/dev/sdc1(0)

# lvs -a -o +devices
LV          VG      Attr      LSize   Origin Snap%   Move Log
Copy% Devices
groupfs      vg      mwi-a-    752.00M                      groupfs_mlog
100.00 groupfs_mimage_0(0),groupfs_mimage_1(0)
[groupfs_mimage_0] vg      iwi-ao    752.00M
/dev/sda1(0)
[groupfs_mimage_1] vg      iwi-ao    752.00M
/dev/sdb1(0)
[groupfs_mlog]   vg      lwi-ao     4.00M      i
/dev/sdc1(0)
```

In this example, the primary leg of the mirror `/dev/sda1` fails. Any write activity to the mirrored volume causes LVM to detect the failed mirror. When this occurs, LVM converts the mirror into a single linear volume. In this case, to trigger the conversion, we execute a `dd` command

```
# dd if=/dev/zero of=/dev/vg/groupfs count=10
10+0 records in
10+0 records out
```

You can use the `lvs` command to verify that the device is now a linear device. Because of the failed disk, I/O errors occur.

```
# lvs -a -o +devices
/dev/sda1: read failed after 0 of 2048 at 0: Input/output error
LV      VG      Attr   LSize   Origin Snap%   Move Log Copy%  Devices
groupfs vg      -wi-a- 752.00M                               /dev/sdb1(0)
```

At this point you should still be able to use the logical volume, but there will be no mirror redundancy.

To rebuild the mirrored volume, you replace the broken drive and recreate the physical volume. If you use the same disk rather than replacing it with a new one, you will see "inconsistent" warnings when you run the `pvcreate` command. You can prevent that warning from appearing by executing the `vgreduce --removemissing` command.

```
# pvcreate /dev/sdi[12]
Physical volume "/dev/sdi1" successfully created
Physical volume "/dev/sdi2" successfully created

# pvscan
PV /dev/sdb1   VG vg    lvm2 [67.83 GB / 67.10 GB free]
PV /dev/sdb2   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdc1   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdc2   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdd1   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdd2   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sde1   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sde2   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdf1   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdf2   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdg1   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdg2   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdh1   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdh2   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdi1   VG vg    lvm2 [603.94 GB]
PV /dev/sdi2   VG vg    lvm2 [603.94 GB]
Total: 16 [2.11 TB] / in use: 14 [949.65 GB] / in no VG: 2 [1.18 TB]
```

Next you extend the original volume group with the new physical volume.

```
# vgextend vg /dev/sdi[12]
Volume group "vg" successfully extended

# pvscan
PV /dev/sdb1   VG vg    lvm2 [67.83 GB / 67.10 GB free]
PV /dev/sdb2   VG vg    lvm2 [67.83 GB / 67.83 GB free]
```

```

PV /dev/sdc1   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdc2   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdd1   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdd2   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sde1   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sde2   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdf1   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdf2   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdg1   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdg2   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdh1   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdh2   VG vg    lvm2 [67.83 GB / 67.83 GB free]
PV /dev/sdi1   VG vg    lvm2 [603.93 GB / 603.93 GB free]
PV /dev/sdi2   VG vg    lvm2 [603.93 GB / 603.93 GB free]
Total: 16 [2.11 TB] / in use: 16 [2.11 TB] / in no VG: 0 [0 ]

```

Convert the linear volume back to its original mirrored state.

```

# lvconvert -m 1 /dev/vg/groupfs /dev/sdi1 /dev/sdb1 /dev/sdc1
Logical volume mirror converted.

```

You can use the **lvs** command to verify that the mirror is restored.

```

# lvs -a -o +devices
LV          VG      Attr      LSize   Origin Snap%  Move Log
Copy% Devices
groupfs     vg      mwi-a-   752.00M                      groupfs_mlog
68.62 groupfs_mimage_0(0),groupfs_mimage_1(0)
[groupfs_mimage_0] vg      iwi-ao   752.00M
/dev/sdb1(0)
[groupfs_mimage_1] vg      iwi-ao   752.00M
/dev/sdi1(0)
[groupfs_mlog]   vg      lwi-ao    4.00M
/dev/sdc1(0)

```

6.4. RECOVERING PHYSICAL VOLUME METADATA

If the volume group metadata area of a physical volume is accidentally overwritten or otherwise destroyed, you will get an error message indicating that the metadata area is incorrect, or that the system was unable to find a physical volume with a particular UUID. You may be able to recover the data from the physical volume by writing a new metadata area on the physical volume specifying the same UUID as the lost metadata.



WARNING

You should not attempt this procedure with a working LVM logical volume. You will lose your data if you specify the incorrect UUID.

The following example shows the sort of output you may see if the metadata area is missing or corrupted.

```
# lvs -a -o +devices
Couldn't find device with uuid 'FmGRh3-zhok-iVI8-7qTD-S5BI-MAEN-NYM5Sk'.
Couldn't find all physical volumes for volume group VG.
Couldn't find device with uuid 'FmGRh3-zhok-iVI8-7qTD-S5BI-MAEN-NYM5Sk'.
Couldn't find all physical volumes for volume group VG.
...
```

You may be able to find the UUID for the physical volume that was overwritten by looking in the `/etc/lvm/archive` directory. Look in the file `VolumeGroupName_xxxx.vg` for the last known valid archived LVM metadata for that volume group.

Alternately, you may find that deactivating the volume and setting the **partial** (`-P`) argument will enable you to find the UUID of the missing corrupted physical volume.

```
# vgchange -an --partial
Partial mode. Incomplete volume groups will be activated read-only.
Couldn't find device with uuid 'FmGRh3-zhok-iVI8-7qTD-S5BI-MAEN-NYM5Sk'.
Couldn't find device with uuid 'FmGRh3-zhok-iVI8-7qTD-S5BI-MAEN-NYM5Sk'.
...
```

Use the `--uuid` and `--restorefile` arguments of the **pvcreate** command to restore the physical volume. The following example labels the `/dev/sdh1` device as a physical volume with the UUID indicated above, **FmGRh3-zhok-iVI8-7qTD-S5BI-MAEN-NYM5Sk**. This command restores the physical volume label with the metadata information contained in **VG_00050.vg**, the most recent good archived metadata for the volume group. The **restorefile** argument instructs the **pvcreate** command to make the new physical volume compatible with the old one on the volume group, ensuring that the new metadata will not be placed where the old physical volume contained data (which could happen, for example, if the original **pvcreate** command had used the command line arguments that control metadata placement, or if the physical volume was originally created using a different version of the software that used different defaults). The **pvcreate** command overwrites only the LVM metadata areas and does not affect the existing data areas.

```
# pvcreate --uuid "FmGRh3-zhok-iVI8-7qTD-S5BI-MAEN-NYM5Sk" --restorefile
/etc/lvm/archive/VG_00050.vg /dev/sdh1
Physical volume "/dev/sdh1" successfully created
```

You can then use the **vgcfgrestore** command to restore the volume group's metadata.

```
# vgcfgrestore VG
Restored volume group VG
```

You can now display the logical volumes.

```
# lvs -a -o +devices
LV      VG      Attr   LSize   Origin Snap%   Move Log Copy%  Devices
stripe VG      -wi--- 300.00G                                     /dev/sdh1
(0),/dev/sda1(0)
stripe VG      -wi--- 300.00G                                     /dev/sdh1
(34728),/dev/sdb1(0)
```

The following commands activate the volumes and display the active volumes.

```
# lvchange -ay /dev/VG/stripe
# lvs -a -o +devices
  LV      VG      Attr   LSize   Origin Snap%   Move Log Copy%  Devices
  stripe VG      -wi-a- 300.00G                                     /dev/sdh1
(0),/dev/sda1(0)
  stripe VG      -wi-a- 300.00G                                     /dev/sdh1
(34728),/dev/sdb1(0)
```

If the on-disk LVM metadata takes as least as much space as what overrode it, this command can recover the physical volume. If what overrode the metadata went past the metadata area, the data on the volume may have been affected. You might be able to use the **fsck** command to recover that data.

6.5. REPLACING A MISSING PHYSICAL VOLUME

If a physical volume fails or otherwise needs to be replaced, you can label a new physical volume to replace the one that has been lost in the existing volume group by following the same procedure as you would for recovering physical volume metadata, described in [Section 6.4, “Recovering Physical Volume Metadata”](#). You can use the **--partial** and **--verbose** arguments of the **vgdisplay** command to display the UUIDs and sizes of any physical volumes that are no longer present. If you wish to substitute another physical volume of the same size, you can use the **pvccreate** command with the **--restorefile** and **--uuid** arguments to initialize a new device with the same UUID as the missing physical volume. You can then use the **vgcfgrestore** command to restore the volume group's metadata.

6.6. REMOVING LOST PHYSICAL VOLUMES FROM A VOLUME GROUP

If you lose a physical volume, you can activate the remaining physical volumes in the volume group with the **--partial** argument of the **vgchange** command. You can remove all the logical volumes that used that physical volume from the volume group with the **--removemissing** argument of the **vgreduce** command.

You should run the **vgreduce** command with the **--test** argument to verify what you will be destroying.

Like most LVM operations, the **vgreduce** command is reversible if you immediately use the **vgcfgrestore** command to restore the volume group metadata to its previous state. For example, if you used the **--removemissing** argument of the **vgreduce** command without the **--test** argument and find you have removed logical volumes you wanted to keep, you can still replace the physical volume and use another **vgcfgrestore** command to return the volume group to its previous state.

6.7. INSUFFICIENT FREE EXTENTS FOR A LOGICAL VOLUME

You may get the error message "Insufficient free extents" when creating a logical volume when you think you have enough extents based on the output of the **vgdisplay** or **vgs** commands. This is because these commands round figures to 2 decimal places to provide human-readable output. To specify exact size, use free physical extent count instead of a multiple of bytes to determine the size of the logical volume.

The **vgdisplay** command, by default, includes this line of output that indicates the free physical extents.

```
# vgdisplay
```

```

--- Volume group ---
...
Free  PE / Size          8780 / 34.30 GB

```

Alternately, you can use the **vg_free_count** and **vg_extent_count** arguments of the **vgs** command to display the free extents and the total number of extents.

```

# vgs -o +vg_free_count,vg_extent_count
VG      #PV #LV #SN Attr   VSize  VFree  Free #Ext
testvg   2   0   0 wz--n- 34.30G 34.30G 8780 8780

```

With 8780 free physical extents, you can enter the following command, using the lower-case **l** argument to use extents instead of bytes:

```
# lvcreate -l 8780 -n testlv testvg
```

This uses all the free extents in the volume group.

```

# vgs -o +vg_free_count,vg_extent_count
VG      #PV #LV #SN Attr   VSize  VFree  Free #Ext
testvg   2   1   0 wz--n- 34.30G    0    0 8780

```

Alternately, you can extend the logical volume to use a percentage of the remaining free space in the volume group by using the **-l** argument of the **lvcreate** command. For information, see [Section 4.4.1](#), “Creating Linear Logical Volumes”.

6.8. DUPLICATE PV WARNINGS FOR MULTIPATHED DEVICES

When using LVM with multipathed storage, some LVM commands (such as **vgs** or **lvchange**) may display messages such as the following when listing a volume group or logical volume.

```

Found duplicate PV GDjTZf7Y03GJHjteq0wrye2dcSCjdaUi: using /dev/dm-5 not
/dev/sdd
Found duplicate PV GDjTZf7Y03GJHjteq0wrye2dcSCjdaUi: using /dev/emcpowerb
not /dev/sde
Found duplicate PV GDjTZf7Y03GJHjteq0wrye2dcSCjdaUi: using /dev/sddlmb
not /dev/sdf

```

After providing information on the root cause for these warnings, this section describes how to address this issue in the following two cases.

- The two devices displayed in the output are both single paths to the same device
- The two devices displayed in the output are both multipath maps

6.8.1. Root Cause of Duplicate PV Warning

With a default configuration, LVM commands will scan for devices in **/dev** and check every resulting device for LVM metadata. This is caused by the default filter in the **/etc/lvm/lvm.conf**, which is as follows:

```
filter = [ "a/*/" ]
```

When using Device Mapper Multipath or other multipath software such as EMC PowerPath or Hitachi Dynamic Link Manager (HDLM), each path to a particular logical unit number (LUN) is registered as a different SCSI device, such as `/dev/sdb` or `/dev/sdc`. The multipath software will then create a new device that maps to those individual paths, such as `/dev/mapper/mpath1` or `/dev/mapper/mpatha` for Device Mapper Multipath, `/dev/emcpowera` for EMC PowerPath, or `/dev/sddlmaab` for Hitachi HDLM. Since each LUN has multiple device nodes in `/dev` that point to the same underlying data, they all contain the same LVM metadata and thus LVM commands will find the same metadata multiple times and report them as duplicates.

These duplicate messages are only warnings and do not mean the LVM operation has failed. Rather, they are alerting the user that only one of the devices has been used as a physical volume and the others are being ignored. If the messages indicate the incorrect device is being chosen or if the warnings are disruptive to users, then a filter can be applied to search only the necessary devices for physical volumes, and to leave out any underlying paths to multipath devices.

6.8.2. Duplicate Warnings for Single Paths

The following example shows a duplicate PV warning in which the duplicate devices displayed are both single paths to the same device. In this case, both `/dev/sdd` and `/dev/sdf` can be found under the same multipath map in the output to the `multipath -ll` command.

```
Found duplicate PV GDjTZf7Y03GJHjteq0wrye2dcSCjdaUi: using **/dev/sdd**
not **/dev/sdf**
```

To prevent this warning from appearing, you can configure a filter in the `/etc/lvm/lvm.conf` file to restrict the devices that LVM will search for metadata. The filter is a list of patterns that will be applied to each device found by a scan of `/dev` (or the directory specified by the `dir` keyword in the `/etc/lvm/lvm.conf` file). Patterns are regular expressions delimited by any character and preceded by `a` (for accept) or `r` (for reject). The list is traversed in order, and the first regex that matches a device determines if the device will be accepted or rejected (ignored). Devices that don't match any patterns are accepted. For general information on LVM filters, see [Section 4.5, "Controlling LVM Device Scans with Filters"](#).

The filter you configure should include all devices that need to be checked for LVM metadata, such as the local hard drive with the root volume group on it and any multipathed devices. By rejecting the underlying paths to a multipath device (such as `/dev/sdb`, `/dev/sdd`, and so on) you can avoid these duplicate PV warnings, since each unique metadata area will only be found once on the multipath device itself.

The following examples show filters that will avoid duplicate PV warnings due to multiple storage paths being available.

- This filter accepts the second partition on the first hard drive (`/dev/sda` and any device-mapper-multipath devices, while rejecting everything else.

```
filter = [ "a|/dev/sda2$|", "a|/dev/mapper/mpath.*|", "r|.*)" ]
```

- This filter accepts all HP SmartArray controllers and any EMC PowerPath devices.

```
filter = [ "a|/dev/cciss/.*)" , "a|/dev/emcpower.*|", "r|.*)" ]
```

- This filter accepts any partitions on the first IDE drive and any multipath devices.

```
filter = [ "a|/dev/hda.*|", "a|/dev/mapper/mpath.*|", "r|.*)" ]
```

**NOTE**

When adding a new filter to the `/etc/lvm/lvm.conf` file, ensure that the original filter is either commented out with a `#` or is removed.

Once a filter has been configured and the `/etc/lvm/lvm.conf` file has been saved, check the output of these commands to ensure that no physical volumes or volume groups are missing.

```
# pvscan
# vgscan
```

You can also test a filter on the fly, without modifying the `/etc/lvm/lvm.conf` file, by adding the `--config` argument to the LVM command, as in the following example.

```
# lvs --config 'devices{ filter = [ "a|/dev/emcpower.*|", "r|.*)" ] }'
```

**NOTE**

Testing filters using the `--config` argument will not make permanent changes to the server's configuration. Make sure to include the working filter in the `/etc/lvm/lvm.conf` file after testing.

After configuring an LVM filter, it is recommended that you rebuild the `initrd` device with the `dracut` command so that only the necessary devices are scanned upon reboot.

6.8.3. Duplicate Warnings for Multipath Maps

The following examples show a duplicate PV warning for two devices that are both multipath maps. In these examples we are not looking at two different paths, but two different devices.

```
Found duplicate PV GDjTZf7Y03GJHjtec0wrye2dcSCjdaUi: using
**/dev/mapper/mpatha** not **/dev/mapper/mpathc**
```

```
Found duplicate PV GDjTZf7Y03GJHjtec0wrye2dcSCjdaUi: using
**/dev/emcpowera** not **/dev/emcpowerh**
```

This situation is more serious than duplicate warnings for devices that are both single paths to the same device, since these warnings often mean that the machine has been presented devices which it should not be seeing (for example, LUN clones or mirrors). In this case, unless you have a clear idea of what devices should be removed from the machine, the situation could be unrecoverable. It is recommended that you contact Red Hat Technical Support to address this issue.

APPENDIX A. THE DEVICE MAPPER

The Device Mapper is a kernel driver that provides a framework for volume management. It provides a generic way of creating mapped devices, which may be used as logical volumes. It does not specifically know about volume groups or metadata formats.

The Device Mapper provides the foundation for a number of higher-level technologies. In addition to LVM, Device-Mapper multipath and the **dmraid** command use the Device Mapper. The application interface to the Device Mapper is the **ioctl** system call. The user interface is the **dmsetup** command.

LVM logical volumes are activated using the Device Mapper. Each logical volume is translated into a mapped device. Each segment translates into a line in the mapping table that describes the device. The Device Mapper supports a variety of mapping targets, including linear mapping, striped mapping, and error mapping. For example, two disks may be concatenated into one logical volume with a pair of linear mappings, one for each disk. When LVM creates a volume, it creates an underlying device-mapper device that can be queried with the **dmsetup** command. For information about the format of devices in a mapping table, see [Section A.1, “Device Table Mappings”](#). For information about using the **dmsetup** command to query a device, see [Section A.2, “The dmsetup Command”](#).

A.1. DEVICE TABLE MAPPINGS

A mapped device is defined by a table that specifies how to map each range of logical sectors of the device using a supported Device Table mapping. The table for a mapped device is constructed from a list of lines of the form:

```
start length mapping [mapping_parameters...]
```

In the first line of a Device Mapper table, the **start** parameter must equal 0. The **start + length** parameters on one line must equal the **start** on the next line. Which mapping parameters are specified in a line of the mapping table depends on which **mapping** type is specified on the line.

Sizes in the Device Mapper are always specified in sectors (512 bytes).

When a device is specified as a mapping parameter in the Device Mapper, it can be referenced by the device name in the filesystem (for example, **/dev/hda**) or by the major and minor numbers in the format **major:minor**. The major:minor format is preferred because it avoids pathname lookups.

The following shows a sample mapping table for a device. In this table there are four linear targets:

```
0 35258368 linear 8:48 65920
35258368 35258368 linear 8:32 65920
70516736 17694720 linear 8:16 17694976
88211456 17694720 linear 8:16 256
```

The first 2 parameters of each line are the segment starting block and the length of the segment. The next keyword is the mapping target, which in all of the cases in this example is **linear**. The rest of the line consists of the parameters for a **linear** target.

The following subsections describe the format of the following mappings:

- linear
- striped

- mirror
- snapshot and snapshot-origin
- error
- zero
- multipath
- crypt

A.1.1. The linear Mapping Target

A linear mapping target maps a continuous range of blocks onto another block device. The format of a linear target is as follows:

```
start length linear device offset
```

start

starting block in virtual device

length

length of this segment

device

block device, referenced by the device name in the filesystem or by the major and minor numbers in the format ***major:minor***

offset

starting offset of the mapping on the device

The following example shows a linear target with a starting block in the virtual device of 0, a segment length of 1638400, a major:minor number pair of 8:2, and a starting offset for the device of 41146992.

```
0 1638400 linear 8:2 41156992
```

The following example shows a linear target with the device parameter specified as the device ***/dev/hda***.

```
0 20971520 linear /dev/hda 384
```

A.1.2. The striped Mapping Target

The striped mapping target supports striping across physical devices. It takes as arguments the number of stripes and the striping chunk size followed by a list of pairs of device name and sector. The format of a striped target is as follows:

```
start length striped #stripes chunk_size device1 offset1 ... deviceN  
offsetN
```

There is one set of **device** and **offset** parameters for each stripe.

start

starting block in virtual device

length

length of this segment

#stripes

number of stripes for the virtual device

chunk_size

number of sectors written to each stripe before switching to the next; must be power of 2 at least as big as the kernel page size

device

block device, referenced by the device name in the filesystem or by the major and minor numbers in the format **major:minor**.

offset

starting offset of the mapping on the device

The following example shows a striped target with three stripes and a chunk size of 128:

```
| 0 73728 striped 3 128 8:9 384 8:8 384 8:7 9789824
```

0

starting block in virtual device

73728

length of this segment

striped 3 128

stripe across three devices with chunk size of 128 blocks

8:9

major:minor numbers of first device

384

starting offset of the mapping on the first device

8:8

major:minor numbers of second device

384

starting offset of the mapping on the second device

8:7

major:minor numbers of third device

9789824

starting offset of the mapping on the third device

The following example shows a striped target for 2 stripes with 256 KiB chunks, with the device parameters specified by the device names in the file system rather than by the major and minor numbers.

```
0 65536 striped 2 512 /dev/hda 0 /dev/hdb 0
```

A.1.3. The mirror Mapping Target

The mirror mapping target supports the mapping of a mirrored logical device. The format of a mirrored target is as follows:

```
start length mirror log_type #logargs logarg1 ... logargN #devs device1  
offset1 ... deviceN offsetN
```

start

starting block in virtual device

length

length of this segment

log_type

The possible log types and their arguments are as follows:

core

The mirror is local and the mirror log is kept in core memory. This log type takes 1 - 3 arguments:

```
regionsize [[no]sync] [block_on_error]
```

disk

The mirror is local and the mirror log is kept on disk. This log type takes 2 - 4 arguments:

```
logdevice regionsize [[no]sync] [block_on_error]
```

clustered_core

The mirror is clustered and the mirror log is kept in core memory. This log type takes 2 - 4 arguments:

```
regionsize UUID [[no]sync] [block_on_error]
```

clustered_disk

The mirror is clustered and the mirror log is kept on disk. This log type takes 3 - 5 arguments:

```
logdevice regionsize UUID [[no]sync] [block_on_error]
```

LVM maintains a small log which it uses to keep track of which regions are in sync with the mirror or mirrors. The *regionsize* argument specifies the size of these regions.

In a clustered environment, the *UUID* argument is a unique identifier associated with the mirror log device so that the log state can be maintained throughout the cluster.

The optional **[no]sync** argument can be used to specify the mirror as "in-sync" or "out-of-sync". The **block_on_error** argument is used to tell the mirror to respond to errors rather than ignoring them.

#log_args

number of log arguments that will be specified in the mapping

logargs

the log arguments for the mirror; the number of log arguments provided is specified by the **#log-args** parameter and the valid log arguments are determined by the **log_type** parameter.

#devs

the number of legs in the mirror; a device and an offset is specified for each leg

device

block device for each mirror leg, referenced by the device name in the filesystem or by the major and minor numbers in the format **major:minor**. A block device and offset is specified for each mirror leg, as indicated by the **#devs** parameter.

offset

starting offset of the mapping on the device. A block device and offset is specified for each mirror leg, as indicated by the **#devs** parameter.

The following example shows a mirror mapping target for a clustered mirror with a mirror log kept on disk.

```
0 52428800 mirror clustered_disk 4 253:2 1024 UUID block_on_error 3 253:3
0 253:4 0 253:5 0
```

0

starting block in virtual device

52428800

length of this segment

mirror clustered_disk

mirror target with a log type specifying that mirror is clustered and the mirror log is maintained on disk

4

4 mirror log arguments will follow

253:2

major:minor numbers of log device

1024

region size the mirror log uses to keep track of what is in sync

UUID

UUID of mirror log device to maintain log information throughout a cluster

block_on_error

mirror should respond to errors

3

number of legs in mirror

253:3 0 253:4 0 253:5 0

major:minor numbers and offset for devices constituting each leg of mirror

A.1.4. The snapshot and snapshot-origin Mapping Targets

When you create the first LVM snapshot of a volume, four Device Mapper devices are used:

1. A device with a **linear** mapping containing the original mapping table of the source volume.
2. A device with a **linear** mapping used as the copy-on-write (COW) device for the source volume; for each write, the original data is saved in the COW device of each snapshot to keep its visible content unchanged (until the COW device fills up).
3. A device with a **snapshot** mapping combining #1 and #2, which is the visible snapshot volume.
4. The "original" volume (which uses the device number used by the original source volume), whose table is replaced by a "snapshot-origin" mapping from device #1.

A fixed naming scheme is used to create these devices. For example, you might use the following commands to create an LVM volume named **base** and a snapshot volume named **snap** based on that volume.

```
# lvcreate -L 1G -n base volumeGroup
# lvcreate -L 100M --snapshot -n snap volumeGroup/base
```

This yields four devices, which you can view with the following commands:

```
# dmsetup table|grep volumeGroup
volumeGroup-base-real: 0 2097152 linear 8:19 384
volumeGroup-snap-cow: 0 204800 linear 8:19 2097536
volumeGroup-snap: 0 2097152 snapshot 254:11 254:12 P 16
volumeGroup-base: 0 2097152 snapshot-origin 254:11

# ls -lL /dev/mapper/volumeGroup-*
brw----- 1 root root 254, 11 29 ago 18:15 /dev/mapper/volumeGroup-base-real
brw----- 1 root root 254, 12 29 ago 18:15 /dev/mapper/volumeGroup-snap-
```

```
COW  
brw----- 1 root root 254, 13 29 ago 18:15 /dev/mapper/volumeGroup-snap  
brw----- 1 root root 254, 10 29 ago 18:14 /dev/mapper/volumeGroup-base
```

The format for the **snapshot-origin** target is as follows:

```
start length snapshot-origin origin
```

start

starting block in virtual device

length

length of this segment

origin

base volume of snapshot

The **snapshot-origin** will normally have one or more snapshots based on it. Reads will be mapped directly to the backing device. For each write, the original data will be saved in the COW device of each snapshot to keep its visible content unchanged until the COW device fills up.

The format for the **snapshot** target is as follows:

```
start length snapshot origin COW-device P|N chunksize
```

start

starting block in virtual device

length

length of this segment

origin

base volume of snapshot

COW-device

device on which changed chunks of data are stored

P|N

P (Persistent) or N (Not persistent); indicates whether the snapshot will survive after reboot. For transient snapshots (N) less metadata must be saved on disk; they can be kept in memory by the kernel.

chunksize

size in sectors of changed chunks of data that will be stored on the COW device

The following example shows a **snapshot-origin** target with an origin device of 254:11.

```
0 2097152 snapshot-origin 254:11
```

The following example shows a **snapshot** target with an origin device of 254:11 and a COW device of 254:12. This snapshot device is persistent across reboots and the chunk size for the data stored on the COW device is 16 sectors.

```
0 2097152 snapshot 254:11 254:12 P 16
```

A.1.5. The error Mapping Target

With an error mapping target, any I/O operation to the mapped sector fails.

An error mapping target can be used for testing. To test how a device behaves in failure, you can create a device mapping with a bad sector in the middle of a device, or you can swap out the leg of a mirror and replace the leg with an error target.

An error target can be used in place of a failing device, as a way of avoiding timeouts and retries on the actual device. It can serve as an intermediate target while you rearrange LVM metadata during failures.

The **error** mapping target takes no additional parameters besides the *start* and *length* parameters.

The following example shows an **error** target.

```
0 65536 error
```

A.1.6. The zero Mapping Target

The **zero** mapping target is a block device equivalent of **/dev/zero**. A read operation to this mapping returns blocks of zeros. Data written to this mapping is discarded, but the write succeeds. The **zero** mapping target takes no additional parameters besides the *start* and *length* parameters.

The following example shows a **zero** target for a 16Tb Device.

```
0 65536 zero
```

A.1.7. The multipath Mapping Target

The multipath mapping target supports the mapping of a multipathed device. The format for the **multipath** target is as follows:

```
start length multipath #features [feature1 ... featureN] #handlerargs
[handlerarg1 ... handlerargN] #pathgroups pathgroup pathgroupargs1 ...
pathgroupargsN
```

There is one set of **pathgroupargs** parameters for each path group.

start

starting block in virtual device

length

length of this segment

#features

The number of multipath features, followed by those features. If this parameter is zero, then there is no **feature** parameter and the next device mapping parameter is **#handlerargs**. Currently there is one supported feature that can be set with the **features** attribute in the **multipath.conf** file, **queue_if_no_path**. This indicates that this multipathed device is currently set to queue I/O operations if there is no path available.

In the following example, the **no_path_retry** attribute in the **multipath.conf** file has been set to queue I/O operations only until all paths have been marked as failed after a set number of attempts have been made to use the paths. In this case, the mapping appears as follows until all the path checkers have failed the specified number of checks.

```
0 71014400 multipath 1 queue_if_no_path 0 2 1 round-robin 0 2 1 66:128 \
1000 65:64 1000 round-robin 0 2 1 8:0 1000 67:192 1000
```

After all the path checkers have failed the specified number of checks, the mapping would appear as follows.

```
0 71014400 multipath 0 0 2 1 round-robin 0 2 1 66:128 1000 65:64 1000 \
round-robin 0 2 1 8:0 1000 67:192 1000
```

#handlerargs

The number of hardware handler arguments, followed by those arguments. A hardware handler specifies a module that will be used to perform hardware-specific actions when switching path groups or handling I/O errors. If this is set to 0, then the next parameter is **#pathgroups**.

#pathgroups

The number of path groups. A path group is the set of paths over which a multipathed device will load balance. There is one set of **pathgroupargs** parameters for each path group.

pathgroup

The next path group to try.

pathgroupsargs

Each path group consists of the following arguments:

```
pathselector #selectorargs #paths #pathargs device1 ioreqs1 ... deviceN
ioreqsN
```

There is one set of path arguments for each path in the path group.

pathselector

Specifies the algorithm in use to determine what path in this path group to use for the next I/O operation.

#selectorargs

The number of path selector arguments which follow this argument in the multipath mapping. Currently, the value of this argument is always 0.

#paths

The number of paths in this path group.

#pathargs

The number of path arguments specified for each path in this group. Currently this number is always 1, the **ioreqs** argument.

device

The block device number of the path, referenced by the major and minor numbers in the format **major:minor**

ioreqs

The number of I/O requests to route to this path before switching to the next path in the current group.

Figure A.1, “Multipath Mapping Target” shows the format of a multipath target with two path groups.




Figure A.1. Multipath Mapping Target

The following example shows a pure failover target definition for the same multipath device. In this target there are four path groups, with only one open path per path group so that the multipathed device will use only one path at a time.

```
0 71014400 multipath 0 0 4 1 round-robin 0 1 1 66:112 1000 \
round-robin 0 1 1 67:176 1000 round-robin 0 1 1 68:240 1000 \
round-robin 0 1 1 65:48 1000
```

The following example shows a full spread (multibus) target definition for the same multipathed device. In this target there is only one path group, which includes all of the paths. In this setup, multipath spreads the load evenly to all of the paths.

```
0 71014400 multipath 0 0 1 1 round-robin 0 4 1 66:112 1000 \
67:176 1000 68:240 1000 65:48 1000
```

For further information about multipathing, see the *DM Multipath* manual.

A.1.8. The crypt Mapping Target

The **crypt** target encrypts the data passing through the specified device. It uses the kernel Crypto API.

The format for the **crypt** target is as follows:

```
start length crypt cipher key IV-offset device offset
```

start

starting block in virtual device

length

length of this segment

cipher

Cipher consists of **cipher[-chainmode]-ivmode[:iv options]**.

cipher

Ciphers available are listed in **/proc/crypto** (for example, **aes**).

chainmode

Always use **cbc**. Do not use **ebc**; it does not use an initial vector (IV).

ivmode[:iv options]

IV is an initial vector used to vary the encryption. The IV mode is **plain** or **essiv:hash**. An **ivmode** of **-plain** uses the sector number (plus IV offset) as the IV. An **ivmode** of **-essiv** is an enhancement avoiding a watermark weakness.

key

Encryption key, supplied in hex

IV-offset

Initial Vector (IV) offset

device

block device, referenced by the device name in the filesystem or by the major and minor numbers in the format **major:minor**

offset

starting offset of the mapping on the device

The following is an example of a **crypt** target.

```
0 2097152 crypt aes-plain 0123456789abcdef0123456789abcdef 0 /dev/hda 0
```

A.2. THE DMSETUP COMMAND

The **dmsetup** command is a command line wrapper for communication with the Device Mapper. For general system information about LVM devices, you may find the **info**, **ls**, **status**, and **deps** options of the **dmsetup** command to be useful, as described in the following subsections.

For information about additional options and capabilities of the **dmsetup** command, see the **dmsetup(8)** man page.

A.2.1. The dmsetup info Command

The **dmsetup info device** command provides summary information about Device Mapper devices. If you do not specify a device name, the output is information about all of the currently configured Device Mapper devices. If you specify a device, then this command yields information for that device only.

The **dmsetup info** command provides information in the following categories:

Name

The name of the device. An LVM device is expressed as the volume group name and the logical volume name separated by a hyphen. A hyphen in the original name is translated to two hyphens. During standard LVM operations, you should not use the name of an LVM device in this format to specify an LVM device directly, but instead you should use the *vg/lv* alternative.

State

Possible device states are **SUSPENDED**, **ACTIVE**, and **READ-ONLY**. The **dmsetup suspend** command sets a device state to **SUSPENDED**. When a device is suspended, all I/O operations to that device stop. The **dmsetup resume** command restores a device state to **ACTIVE**.

Read Ahead

The number of data blocks that the system reads ahead for any open file on which read operations are ongoing. By default, the kernel chooses a suitable value automatically. You can change this value with the **--readahead** option of the **dmsetup** command.

Tables present

Possible states for this category are **LIVE** and **INACTIVE**. An **INACTIVE** state indicates that a table has been loaded which will be swapped in when a **dmsetup resume** command restores a device state to **ACTIVE**, at which point the table's state becomes **LIVE**. For information, see the **dmsetup** man page.

Open count

The open reference count indicates how many times the device is opened. A **mount** command opens a device.

Event number

The current number of events received. Issuing a **dmsetup wait n** command allows you to wait for the n'th event, blocking the call until it is received.

Major, minor

Major and minor device number.

Number of targets

The number of segments that make up a device. For example, a linear device spanning 3 disks would have 3 targets. A linear device composed of the beginning and end of a disk, but not the middle would have 2 targets.

UUID

UUID of the device.

The following example shows partial output for the **dmsetup info** command.

```
# dmsetup info
Name:                testgfsvg-testgfslv1
State:               ACTIVE
Read Ahead:         256
Tables present:     LIVE
Open count:         0
Event number:       0
Major, minor:       253, 2
Number of targets:  2
UUID: LVM-K528WUGQgPadNXycFrrf9LnPlUMswgkCkpgPIgYzSvigM7SfewCypddNSwtNzc2N
...
Name:                VolGroup00-LogVol00
State:               ACTIVE
Read Ahead:         256
Tables present:     LIVE
Open count:         1
Event number:       0
Major, minor:       253, 0
Number of targets:  1
UUID: LVM-t0cS1kqFV9drb0X1Vr8sxeYP0tqcrpdegyqj5lZxe45JMGlmvtqLmbLpBcenh2L3
```

A.2.2. The dmsetup ls Command

You can list the device names of mapped devices with the **dmsetup ls** command. You can list devices that have at least one target of a specified type with the **dmsetup ls --target *target_type*** command. For other options of the **dmsetup ls** command, see the **dmsetup** man page.

The following example shows the command to list the device names of currently configured mapped devices.

```
# dmsetup ls
testgfsvg-testgfslv3    (253:4)
testgfsvg-testgfslv2    (253:3)
testgfsvg-testgfslv1    (253:2)
VolGroup00-LogVol01     (253:1)
VolGroup00-LogVol00     (253:0)
```

The following example shows the command to list the device names of currently configured mirror mappings.

```
# dmsetup ls --target mirror
lock_stress-grant--02.1722    (253, 34)
lock_stress-grant--01.1720    (253, 18)
lock_stress-grant--03.1718    (253, 52)
```

```
lock_stress-grant--02.1716      (253, 40)
lock_stress-grant--03.1713      (253, 47)
lock_stress-grant--02.1709      (253, 23)
lock_stress-grant--01.1707      (253, 8)
lock_stress-grant--01.1724      (253, 14)
lock_stress-grant--03.1711      (253, 27)
```

LVM configurations that are stacked on multipath or other device mapper devices can be complex to sort out. The **dmsetup ls** command provides a **--tree** option that displays dependencies between devices as a tree, as in the following example.

```
# dmsetup ls --tree
vgtest-lvmir (253:13)
├─vgtest-lvmir_mimage_1 (253:12)
│   └─mpathep1 (253:8)
│       └─mpathe (253:5)
│           ├── (8:112)
│           └─ (8:64)
├─vgtest-lvmir_mimage_0 (253:11)
│   └─mpathcp1 (253:3)
│       └─mpathc (253:2)
│           ├── (8:32)
│           └─ (8:16)
└─vgtest-lvmir_mlog (253:4)
    └─mpathfp1 (253:10)
        └─mpathf (253:6)
            ├── (8:128)
            └─ (8:80)
```

A.2.3. The dmsetup status Command

The **dmsetup status device** command provides status information for each target in a specified device. If you do not specify a device name, the output is information about all of the currently configured Device Mapper devices. You can list the status only of devices that have at least one target of a specified type with the **dmsetup status --target target_type** command.

The following example shows the command to list the status of the targets in all currently configured mapped devices.

```
# dmsetup status
testgfsvg-testgfs1v3: 0 312352768 linear
testgfsvg-testgfs1v2: 0 312352768 linear
testgfsvg-testgfs1v1: 0 312352768 linear
testgfsvg-testgfs1v1: 312352768 50331648 linear
VolGroup00-LogVol01: 0 4063232 linear
VolGroup00-LogVol00: 0 151912448 linear
```

A.2.4. The dmsetup deps Command

The **dmsetup deps device** command provides a list of (major, minor) pairs for devices referenced by the mapping table for the specified device. If you do not specify a device name, the output is information about all of the currently configured Device Mapper devices.

The following example shows the command to list the dependencies of all currently configured mapped devices.

```
# dmsetup deps
testgfsvg-testgfs1v3: 1 dependencies      : (8, 16)
testgfsvg-testgfs1v2: 1 dependencies      : (8, 16)
testgfsvg-testgfs1v1: 1 dependencies      : (8, 16)
VolGroup00-LogVol01: 1 dependencies      : (8, 2)
VolGroup00-LogVol00: 1 dependencies      : (8, 2)
```

The following example shows the command to list the dependencies only of the device **lock_stress-grant--02.1722**:

```
# dmsetup deps lock_stress-grant--02.1722
3 dependencies      : (253, 33) (253, 32) (253, 31)
```

A.3. DEVICE MAPPER SUPPORT FOR THE UDEV DEVICE MANAGER

The primary role of the **udev** device manager is to provide a dynamic way of setting up nodes in the **/dev** directory. The creation of these nodes is directed by the application of **udev** rules in user space. These rules are processed on **udev** events sent from the kernel directly as a result of adding, removing or changing particular devices. This provides a convenient and central mechanism for hotplugging support.

Besides creating the actual nodes, the **udev** device manager is able to create symbolic links which you can name. This provides you the freedom to choose their own customized naming and directory structure in the **/dev** directory, if needed.

Each **udev** event contains basic information about the device being processed, such as its name, the subsystem it belongs to, the device's type, its major and minor number used, and the type of the event. Given that, and having the possibility of accessing all the information found in the **/sys** directory that is also accessible within **udev** rules, you are able to utilize simple filters based on this information and run the rules conditionally based on this information.

The **udev** device manager also provides a centralized way of setting up the nodes' permissions. You can easily add a customized set of rules to define the permissions for any device specified by any bit of information that is available while processing the event.

It is also possible to add program hooks in **udev** rules directly. The **udev** device manager can call these programs to provide further processing that is needed to handle the event. Also, the program can export environment variables as a result of this processing. Any results given can be used further in the rules as a supplementary source of information.

Any software using the **udev** library is able to receive and process **udev** events with all the information that is available, so the processing is not bound to the **udev** daemon only.

A.3.1. udev Integration with the Device Mapper

The Device Mapper provides direct support for **udev** integration. This synchronizes the Device Mapper with all **udev** processing related to Device Mapper devices, including LVM devices. The synchronization is needed since the rule application in the **udev** daemon is a form of parallel processing with the program that is the source of the device's changes (such as **dmsetup** and LVM). Without this support, it

was a common problem for a user to try to remove a device that was still open and processed by **udev** rules as a result of a previous change event; this was particularly common when there was a very short time between changes for that device.

Red Hat Enterprise Linux provides officially supported **udev** rules for Device Mapper devices in general and for LVM as well. [Table A.1, “udev Rules for Device-Mapper Devices”](#) summarizes these rules, which are installed in `/lib/udev/rules.d`.

Table A.1. udev Rules for Device-Mapper Devices

Filename	Description
10-dm.rules	Contains general Device Mapper rules and creates the symlinks in <code>/dev/mapper</code> with a <code>/dev/dm-N</code> target where N is a number assigned dynamically to a device by the kernel (<code>/dev/dm-N</code> is a node) NOTE: <code>/dev/dm-N</code> nodes should <i>never</i> be used in scripts to access the device since the N number is assigned dynamically and changes with the sequence of how devices are activated. Therefore, true names in the <code>/dev/mapper</code> directory should be used. This layout is to support udev requirements of how nodes/symlinks should be created.
11-dm-lvm.rules	Contains rules applied for LVM devices and creates the symlinks for the volume group's logical volumes. The symlinks are created in the <code>/dev/vgname</code> directory with a <code>/dev/dm-N</code> target. NOTE: To be consistent with the standard for naming all future rules for Device Mapper subsystems, udev rules should follow the format 11-dm-subsystem_name.rules. Any libdevmapper users providing udev rules as well should follow this standard.
13-dm-disk.rules	Contains rules to be applied for all Device Mapper devices in general and creates symlinks in the <code>/dev/disk/by-id</code> and the <code>/dev/disk/by-uuid</code> directories.
95-dm-notify.rules	Contains the rule to notify the waiting process using libdevmapper (just like LVM and dmsetup). The notification is done after all previous rules are applied, to ensure any udev processing is complete. Notified process is then resumed.
69-dm-lvm-metad.rules	Contains a hook to trigger an LVM scan on any newly appeared block device in the system and do any LVM autoactivation if possible. This supports the lvmetad daemon, which is set with use_lvmetad=1 in the lvm.conf file. The lvmetad daemon and autoactivation are not supported in a clustered environment.

You can add additional customized permission rules by means of the **12-dm-permissions.rules** file. This file is *not* installed in the `/lib/udev/rules` directory; it is found in the

`/usr/share/doc/device-mapper-version` directory. The `12-dm-permissions.rules` file is a template containing hints for how to set the permissions, based on some matching rules given as an example; the file contains examples for some common situations. You can edit this file and place it manually in the `/etc/udev/rules.d` directory where it will survive updates, so the settings will remain.

These rules set all basic variables that could be used by any other rules while processing the events.

The following variables are set in `10-dm.rules`:

- **DM_NAME**: Device Mapper device name
- **DM_UUID**: Device Mapper device UUID
- **DM_SUSPENDED**: the suspended state of Device Mapper device
- **DM_UDEV_RULES_VSN**: `udev` rules version (this is primarily for all other rules to check that previously mentioned variables are set directly by official Device Mapper rules)

The following variables are set in `11-dm-lvm.rules`:

- **DM_LV_NAME**: logical volume name
- **DM_VG_NAME**: volume group name
- **DM_LV_LAYER**: LVM layer name

All these variables can be used in the `12-dm-permissions.rules` file to define a permission for specific Device Mapper devices, as documented in the `12-dm-permissions.rules` file.

A.3.2. Commands and Interfaces that Support udev

Table A.2, “[dmsetup Commands to Support udev](#)” summarizes the `dmsetup` commands that support `udev` integration.

Table A.2. dmsetup Commands to Support udev

Command	Description
<code>dmsetup udevcomplete</code>	Used to notify that udev has completed processing the rules and unlocks waiting process (called from within udev rules in <code>95-dm-notify.rules</code>).
<code>dmsetup udevcomplete_all</code>	Used for debugging purposes to manually unlock all waiting processes.
<code>dmsetup udevcookies</code>	Used for debugging purposes, to show all existing cookies (system-wide semaphores).
<code>dmsetup udevcreatecookie</code>	Used to create a cookie (semaphore) manually. This is useful to run more processes under one synchronization resource.
<code>dmsetup udevreleasecookie</code>	Used to wait for all udev processing related to all processes put under that one synchronization cookie.

The **dmsetup** options that support **udev** integration are as follows.

--udevcookie

Needs to be defined for all **dmsetup** processes we would like to add into a **udev** transaction. It is used in conjunction with **udevcreatecookie** and **udevreleasecookie**:

```
COOKIE=$(dmsetup udevcreatecookie)
dmsetup command --udevcookie $COOKIE ....
dmsetup command --udevcookie $COOKIE ....
....
dmsetup command --udevcookie $COOKIE ....
dmsetup udevreleasecookie --udevcookie $COOKIE
```

Besides using the **--udevcookie** option, you can just export the variable into an environment of the process:

```
export DM_UDEV_COOKIE=$(dmsetup udevcreatecookie)
dmsetup command ...
dmsetup command ...
...
dmsetup command ...
```

--noudevrules

Disables **udev** rules. Nodes/symlinks will be created by **libdevmapper** itself (the old way). This option is for debugging purposes, if **udev** does not work correctly.

--noudevsync

Disables **udev** synchronization. This is also for debugging purposes.

For more information on the **dmsetup** command and its options, see the **dmsetup(8)** man page.

The LVM commands support the following options that support **udev** integration:

- **--noudevrules**: as for the **dmsetup** command, disables **udev** rules.
- **--noudevsync**: as for the **dmsetup** command, disables **udev** synchronization.

The **lvm.conf** file includes the following options that support **udev** integration:

- **udev_rules**: enables/disables **udev_rules** for all LVM2 commands globally.
- **udev_sync**: enables/disables **udev** synchronization for all LVM commands globally.

For more information on the **lvm.conf** file options, see the inline comments in the **lvm.conf** file.

APPENDIX B. THE LVM CONFIGURATION FILES

LVM supports multiple configuration files. At system startup, the **lvm.conf** configuration file is loaded from the directory specified by the environment variable **LVM_SYSTEM_DIR**, which is set to **/etc/lvm** by default.

The **lvm.conf** file can specify additional configuration files to load. Settings in later files override settings from earlier ones. To display the settings in use after loading all the configuration files, execute the **lvmconfig** command.

For information on loading additional configuration files, see [Section D.2, “Host Tags”](#).

B.1. THE LVM CONFIGURATION FILES

The following files are used for LVM configuration:

/etc/lvm/lvm.conf

Central configuration file read by the tools.

etc/lvm/lvm_hosttag.conf

For each host tag, an extra configuration file is read if it exists: **lvm_hosttag.conf**. If that file defines new tags, then further configuration files will be appended to the list of files to read in. For information on host tags, see [Section D.2, “Host Tags”](#).

In addition to the LVM configuration files, a system running LVM includes the following files that affect LVM system setup:

/etc/lvm/cache/.cache

Device name filter cache file (configurable).

/etc/lvm/backup/

Directory for automatic volume group metadata backups (configurable).

/etc/lvm/archive/

Directory for automatic volume group metadata archives (configurable with regard to directory path and archive history depth).

/var/lock/lvm/

In single-host configuration, lock files to prevent parallel tool runs from corrupting the metadata; in a cluster, cluster-wide DLM is used.

B.2. THE LVMCONFIG COMMAND

You can display the current LVM configuration, or save the configuration to a file, with the **lvmconfig** command. There are a variety of features that the **lvmconfig** command provides, including the following;

- You can dump the current lvm configuration merged with any tag configuration files.

- You can dump all current configuration settings for which the values differ from the defaults.
- You can dump all new configuration settings introduced in the current LVM version, in a specific LVM version.
- You can dump all configuration settings that can be customized in a profile, either in their entirety or separately for command and metadata profiles. For information on LVM profiles see [Section B.3, “LVM Profiles”](#).
- You can dump only the configuration settings for a specific version of LVM.
- You can validate the current configuration.

For a full list of supported features and information on specifying the **lvmconfig** options, see the **lvmconfig** man page.

B.3. LVM PROFILES

An LVM profile is a set of selected customizable configuration settings that can be used to achieve certain characteristics in various environments or uses. Normally, the name of the profile should reflect that environment or use. An LVM profile overrides existing configuration.

There are two groups of LVM profiles that LVM recognizes: *command* profiles and *metadata* profiles.

- A command profile is used to override selected configuration settings at the global LVM command level. The profile is applied at the beginning of LVM command execution and it is used throughout the time of the LVM command execution. You apply a command profile by specifying the **--commandprofile *ProfileName*** option when executing an LVM command.
- A metadata profile is used to override selected configuration settings at the volume group/logical volume level. It is applied independently for each volume group/logical volume that is being processed. As such, each volume group/logical volume can store the profile name used in its metadata so that next time the volume group/logical volume is processed, the profile is applied automatically. If the volume group and any of its logical volumes have different profiles defined, the profile defined for the logical volume is preferred.
 - You can attach a metadata profile to a volume group or logical volume by specifying the **--metadataprofile *ProfileName*** option when you create the volume group or logical volume with the **vgcreate** or **lvcreate** command.
 - You can attach or detach a metadata profile to an existing volume group or logical volume by specifying the **--metadataprofile *ProfileName*** or the **--detachprofile** option of the **lvchange** or **vgchange** command.
 - You can specify the **-o vg_profile** and **-o lv_profile** output options of the **vgs** and **lvs** commands to display the metadata profile currently attached to a volume group or a logical volume.

The set of options allowed for command profiles and the set of options allowed for metadata profiles are mutually exclusive. The settings that belong to either of these two sets cannot be mixed together and the LVM tools will reject such profiles.

LVM provides a few predefined configuration profiles. The LVM profiles are stored in the **/etc/lvm/profile** directory by default. This location can be changed by using the **profile_dir** setting in the **/etc/lvm/lvm.conf** file. Each profile configuration is stored in *ProfileName*.profile file in

the **profile** directory. When referencing the profile in an LVM command, the **.profile** suffix is omitted.

You can create additional profiles with different values. For this purpose, LVM provides the **command_profile_template.profile** file (for command profiles) and the **metadata_profile_template.profile** file (for metadata profiles) which contain all settings that are customizable by profiles of each type. You can copy these template profiles and edit them as needed.

Alternatively, you can use the **lvmconfig** command to generate a new profile for a given section of the profile file for either profile type. The following command creates a new command profile named *ProfileName.profile* consisting of the settings in *section*.

```
lvmconfig --file ProfileName.profile --type profilable-command section
```

The following command creates a new metadata profile named *ProfileName.profile* consisting of the settings in *section*.

```
lvmconfig --file ProfileName.profile --type profilable-metadata section
```

If the section is not specified, all settings that can be customized by a profile are reported.

B.4. SAMPLE LVM.CONF FILE

The following is a sample **lvm.conf** configuration file. Your configuration file may differ slightly from this one.



NOTE

You can generate an **lvm.conf** file with all of the default values set and with the comments included by running the following command:

```
lvmconfig --type default --withcomments
```

```
# This is an example configuration file for the LVM2 system.
# It contains the default settings that would be used if there was no
# /etc/lvm/lvm.conf file.
#
# Refer to 'man lvm.conf' for further information including the file
# layout.
#
# Refer to 'man lvm.conf' for information about how settings configured in
# this file are combined with built-in values and command line options to
# arrive at the final values used by LVM.
#
# Refer to 'man lvmconfig' for information about displaying the built-in
# and configured values used by LVM.
#
# If a default value is set in this file (not commented out), then a
# new version of LVM using this file will continue using that value,
# even if the new version of LVM changes the built-in default value.
#
```

```

# To put this file in a different directory and override /etc/lvm set
# the environment variable LVM_SYSTEM_DIR before running the tools.
#
# N.B. Take care that each setting only appears once if uncommenting
# example settings in this file.

# Configuration section config.
# How LVM configuration settings are handled.
config {

    # Configuration option config/checks.
    # If enabled, any LVM configuration mismatch is reported.
    # This implies checking that the configuration key is understood by
    # LVM and that the value of the key is the proper type. If disabled,
    # any configuration mismatch is ignored and the default value is used
    # without any warning (a message about the configuration key not being
    # found is issued in verbose mode only).
    checks = 1

    # Configuration option config/abort_on_errors.
    # Abort the LVM process if a configuration mismatch is found.
    abort_on_errors = 0

    # Configuration option config/profile_dir.
    # Directory where LVM looks for configuration profiles.
    profile_dir = "/etc/lvm/profile"
}

# Configuration section devices.
# How LVM uses block devices.
devices {

    # Configuration option devices/dir.
    # Directory in which to create volume group device nodes.
    # Commands also accept this as a prefix on volume group names.
    # This configuration option is advanced.
    dir = "/dev"

    # Configuration option devices/scan.
    # Directories containing device nodes to use with LVM.
    # This configuration option is advanced.
    scan = [ "/dev" ]

    # Configuration option devices/obtain_device_list_from_udev.
    # Obtain the list of available devices from udev.
    # This avoids opening or using any inapplicable non-block devices or
    # subdirectories found in the udev directory. Any device node or
    # symlink not managed by udev in the udev directory is ignored. This
    # setting applies only to the udev-managed device directory; other
    # directories will be scanned fully. LVM needs to be compiled with
    # udev support for this setting to apply.
    obtain_device_list_from_udev = 1

    # Configuration option devices/external_device_info_source.
    # Select an external device information source.

```

```

# Some information may already be available in the system and LVM can
# use this information to determine the exact type or use of devices it
# processes. Using an existing external device information source can
# speed up device processing as LVM does not need to run its own native
# routines to acquire this information. For example, this information
# is used to drive LVM filtering like MD component detection, multipath
# component detection, partition detection and others.
#
# Accepted values:
#   none
#       No external device information source is used.
#   udev
#       Reuse existing udev database records. Applicable only if LVM is
#       compiled with udev support.
#
external_device_info_source = "none"

# Configuration option devices/preferred_names.
# Select which path name to display for a block device.
# If multiple path names exist for a block device, and LVM needs to
# display a name for the device, the path names are matched against
# each item in this list of regular expressions. The first match is
# used. Try to avoid using un-descriptive /dev/dm-N names, if present.
# If no preferred name matches, or if preferred_names are not defined,
# the following built-in preferences are applied in order until one
# produces a preferred name:
# Prefer names with path prefixes in the order of:
# /dev/mapper, /dev/disk, /dev/dm-*, /dev/block.
# Prefer the name with the least number of slashes.
# Prefer a name that is a symlink.
# Prefer the path with least value in lexicographical order.
#
# Example
# preferred_names = [ "^/dev/mpath/", "^/dev/mapper/mpath", "^/dev/[hs]d"
]
#
preferred_names = [ "^/dev/mpath/", "^/dev/mapper/mpath", "^/dev/[hs]d" ]

# Configuration option devices/filter.
# Limit the block devices that are used by LVM commands.
# This is a list of regular expressions used to accept or reject block
# device path names. Each regex is delimited by a vertical bar '|'
# (or any character) and is preceded by 'a' to accept the path, or
# by 'r' to reject the path. The first regex in the list to match the
# path is used, producing the 'a' or 'r' result for the device.
# When multiple path names exist for a block device, if any path name
# matches an 'a' pattern before an 'r' pattern, then the device is
# accepted. If all the path names match an 'r' pattern first, then the
# device is rejected. Unmatching path names do not affect the accept
# or reject decision. If no path names for a device match a pattern,
# then the device is accepted. Be careful mixing 'a' and 'r' patterns,
# as the combination might produce unexpected results (test changes.)
# Run vgscan after changing the filter to regenerate the cache.
# See the use_lvmetad comment for a special case regarding filters.
#
# Example

```

```

# Accept every block device:
# filter = [ "a|.*|" ]
# Reject the cdrom drive:
# filter = [ "r|/dev/cdrom|" ]
# Work with just loopback devices, e.g. for testing:
# filter = [ "a|loop|", "r|.*|" ]
# Accept all loop devices and ide drives except hdc:
# filter = [ "a|loop|", "r|/dev/hdc|", "a|/dev/ide|", "r|.*|" ]
# Use anchors to be very specific:
# filter = [ "a|^/dev/hda8$", "r|.*|" ]
#
# This configuration option has an automatic default value.
# filter = [ "a|.*|" ]

# Configuration option devices/global_filter.
# Limit the block devices that are used by LVM system components.
# Because devices/filter may be overridden from the command line, it is
# not suitable for system-wide device filtering, e.g. udev and lvmetad.
# Use global_filter to hide devices from these LVM system components.
# The syntax is the same as devices/filter. Devices rejected by
# global_filter are not opened by LVM.
# This configuration option has an automatic default value.
# global_filter = [ "a|.*|" ]

# Configuration option devices/cache_dir.
# Directory in which to store the device cache file.
# The results of filtering are cached on disk to avoid rescanning dud
# devices (which can take a very long time). By default this cache is
# stored in a file named .cache. It is safe to delete this file; the
# tools regenerate it. If obtain_device_list_from_udev is enabled, the
# list of devices is obtained from udev and any existing .cache file
# is removed.
cache_dir = "/etc/lvm/cache"

# Configuration option devices/cache_file_prefix.
# A prefix used before the .cache file name. See devices/cache_dir.
cache_file_prefix = ""

# Configuration option devices/write_cache_state.
# Enable/disable writing the cache file. See devices/cache_dir.
write_cache_state = 1

# Configuration option devices/types.
# List of additional acceptable block device types.
# These are of device type names from /proc/devices, followed by the
# maximum number of partitions.
#
# Example
# types = [ "fd", 16 ]
#
# This configuration option is advanced.
# This configuration option does not have a default value defined.

# Configuration option devices/sysfs_scan.
# Restrict device scanning to block devices appearing in sysfs.
# This is a quick way of filtering out block devices that are not

```

```
# present on the system. sysfs must be part of the kernel and mounted.)
sysfs_scan = 1

# Configuration option devices/multipath_component_detection.
# Ignore devices that are components of DM multipath devices.
multipath_component_detection = 1

# Configuration option devices/md_component_detection.
# Ignore devices that are components of software RAID (md) devices.
md_component_detection = 1

# Configuration option devices/fw_raid_component_detection.
# Ignore devices that are components of firmware RAID devices.
# LVM must use an external_device_info_source other than none for this
# detection to execute.
fw_raid_component_detection = 0

# Configuration option devices/md_chunk_alignment.
# Align PV data blocks with md device's stripe-width.
# This applies if a PV is placed directly on an md device.
md_chunk_alignment = 1

# Configuration option devices/default_data_alignment.
# Default alignment of the start of a PV data area in MB.
# If set to 0, a value of 64KiB will be used.
# Set to 1 for 1MiB, 2 for 2MiB, etc.
# This configuration option has an automatic default value.
# default_data_alignment = 1

# Configuration option devices/data_alignment_detection.
# Detect PV data alignment based on sysfs device information.
# The start of a PV data area will be a multiple of minimum_io_size or
# optimal_io_size exposed in sysfs. minimum_io_size is the smallest
# request the device can perform without incurring a read-modify-write
# penalty, e.g. MD chunk size. optimal_io_size is the device's
# preferred unit of receiving I/O, e.g. MD stripe width.
# minimum_io_size is used if optimal_io_size is undefined (0).
# If md_chunk_alignment is enabled, that detects the optimal_io_size.
# This setting takes precedence over md_chunk_alignment.
data_alignment_detection = 1

# Configuration option devices/data_alignment.
# Alignment of the start of a PV data area in KiB.
# If a PV is placed directly on an md device and md_chunk_alignment or
# data_alignment_detection are enabled, then this setting is ignored.
# Otherwise, md_chunk_alignment and data_alignment_detection are
# disabled if this is set. Set to 0 to use the default alignment or the
# page size, if larger.
data_alignment = 0

# Configuration option devices/data_alignment_offset_detection.
# Detect PV data alignment offset based on sysfs device information.
# The start of a PV aligned data area will be shifted by the
# alignment_offset exposed in sysfs. This offset is often 0, but may
# be non-zero. Certain 4KiB sector drives that compensate for windows
# partitioning will have an alignment_offset of 3584 bytes (sector 7
```

```

# is the lowest aligned logical block, the 4KiB sectors start at
# LBA -1, and consequently sector 63 is aligned on a 4KiB boundary).
# pvcreate --dataalignmentoffset will skip this detection.
data_alignment_offset_detection = 1

# Configuration option devices/ignore_suspended_devices.
# Ignore DM devices that have I/O suspended while scanning devices.
# Otherwise, LVM waits for a suspended device to become accessible.
# This should only be needed in recovery situations.
ignore_suspended_devices = 0

# Configuration option devices/ignore_lvm_mirrors.
# Do not scan 'mirror' LVs to avoid possible deadlocks.
# This avoids possible deadlocks when using the 'mirror' segment type.
# This setting determines whether LVs using the 'mirror' segment type
# are scanned for LVM labels. This affects the ability of mirrors to
# be used as physical volumes. If this setting is enabled, it is
# impossible to create VGs on top of mirror LVs, i.e. to stack VGs on
# mirror LVs. If this setting is disabled, allowing mirror LVs to be
# scanned, it may cause LVM processes and I/O to the mirror to become
# blocked. This is due to the way that the mirror segment type handles
# failures. In order for the hang to occur, an LVM command must be run
# just after a failure and before the automatic LVM repair process
# takes place, or there must be failures in multiple mirrors in the
# same VG at the same time with write failures occurring moments before
# a scan of the mirror's labels. The 'mirror' scanning problems do not
# apply to LVM RAID types like 'raid1' which handle failures in a
# different way, making them a better choice for VG stacking.
ignore_lvm_mirrors = 1

# Configuration option devices/disable_after_error_count.
# Number of I/O errors after which a device is skipped.
# During each LVM operation, errors received from each device are
# counted. If the counter of a device exceeds the limit set here,
# no further I/O is sent to that device for the remainder of the
# operation. Setting this to 0 disables the counters altogether.
disable_after_error_count = 0

# Configuration option devices/require_restorefile_with_uuid.
# Allow use of pvcreate --uuid without requiring --restorefile.
require_restorefile_with_uuid = 1

# Configuration option devices/pv_min_size.
# Minimum size in KiB of block devices which can be used as PVs.
# In a clustered environment all nodes must use the same value.
# Any value smaller than 512KiB is ignored. The previous built-in
# value was 512.
pv_min_size = 2048

# Configuration option devices/issue_discards.
# Issue discards to PVs that are no longer used by an LV.
# Discards are sent to an LV's underlying physical volumes when the LV
# is no longer using the physical volumes' space, e.g. lvremove,
# lvreduce. Discards inform the storage that a region is no longer
# used. Storage that supports discards advertise the protocol-specific
# way discards should be issued by the kernel (TRIM, UNMAP, or

```

```

# WRITE SAME with UNMAP bit set). Not all storage will support or
# benefit from discards, but SSDs and thinly provisioned LUNs
# generally do. If enabled, discards will only be issued if both the
# storage and kernel provide support.
issue_discards = 0

# Configuration option devices/allow_changes_with_duplicate_pvs.
# Allow VG modification while a PV appears on multiple devices.
# When a PV appears on multiple devices, LVM attempts to choose the
# best device to use for the PV. If the devices represent the same
# underlying storage, the choice has minimal consequence. If the
# devices represent different underlying storage, the wrong choice
# can result in data loss if the VG is modified. Disabling this
# setting is the safest option because it prevents modifying a VG
# or activating LVs in it while a PV appears on multiple devices.
# Enabling this setting allows the VG to be used as usual even with
# uncertain devices.
allow_changes_with_duplicate_pvs = 0
}

# Configuration section allocation.
# How LVM selects space and applies properties to LVs.
allocation {

# Configuration option allocation/cling_tag_list.
# Advise LVM which PVs to use when searching for new space.
# When searching for free space to extend an LV, the 'cling' allocation
# policy will choose space on the same PVs as the last segment of the
# existing LV. If there is insufficient space and a list of tags is
# defined here, it will check whether any of them are attached to the
# PVs concerned and then seek to match those PV tags between existing
# extents and new extents.
#
# Example
# Use the special tag "@" as a wildcard to match any PV tag:
# cling_tag_list = [ "@" ]
# LVs are mirrored between two sites within a single VG, and
# PVs are tagged with either @site1 or @site2 to indicate where
# they are situated:
# cling_tag_list = [ "@site1", "@site2" ]
#
# This configuration option does not have a default value defined.

# Configuration option allocation/maximise_cling.
# Use a previous allocation algorithm.
# Changes made in version 2.02.85 extended the reach of the 'cling'
# policies to detect more situations where data can be grouped onto
# the same disks. This setting can be used to disable the changes
# and revert to the previous algorithm.
maximise_cling = 1

# Configuration option allocation/use_blkid_wiping.
# Use blkid to detect existing signatures on new PVs and LVs.
# The blkid library can detect more signatures than the native LVM
# detection code, but may take longer. LVM needs to be compiled with
# blkid wiping support for this setting to apply. LVM native detection

```

```

# code is currently able to recognize: MD device signatures,
# swap signature, and LUKS signatures. To see the list of signatures
# recognized by blkid, check the output of the 'blkid -k' command.
use_blkid_wiping = 1

# Configuration option allocation/wipe_signatures_when_zeroing_new_lvs.
# Look for and erase any signatures while zeroing a new LV.
# The --wipesignatures option overrides this setting.
# Zeroing is controlled by the -Z/--zero option, and if not specified,
# zeroing is used by default if possible. Zeroing simply overwrites the
# first 4KiB of a new LV with zeroes and does no signature detection or
# wiping. Signature wiping goes beyond zeroing and detects exact types
# and positions of signatures within the whole LV. It provides a
# cleaner LV after creation as all known signatures are wiped. The LV
# is not claimed incorrectly by other tools because of old signatures
# from previous use. The number of signatures that LVM can detect
# depends on the detection code that is selected (see
# use_blkid_wiping.) Wiping each detected signature must be confirmed.
# When this setting is disabled, signatures on new LVs are not detected
# or erased unless the --wipesignatures option is used directly.
wipe_signatures_when_zeroing_new_lvs = 1

# Configuration option allocation/mirror_logs_require_separate_pvs.
# Mirror logs and images will always use different PVs.
# The default setting changed in version 2.02.85.
mirror_logs_require_separate_pvs = 0

# Configuration option allocation/raid_stripe_all_devices.
# Stripe across all PVs when RAID stripes are not specified.
# If enabled, all PVs in the VG or on the command line are used for
raid0/4/5/6/10
# when the command does not specify the number of stripes to use.
# This was the default behaviour until release 2.02.162.
# This configuration option has an automatic default value.
# raid_stripe_all_devices = 0

# Configuration option
allocation/cache_pool_metadata_require_separate_pvs.
# Cache pool metadata and data will always use different PVs.
cache_pool_metadata_require_separate_pvs = 0

# Configuration option allocation/cache_mode.
# The default cache mode used for new cache.
#
# Accepted values:
#   writethrough
#     Data blocks are immediately written from the cache to disk.
#   writeback
#     Data blocks are written from the cache back to disk after some
#     delay to improve performance.
#
# This setting replaces allocation/cache_pool_cachemode.
# This configuration option has an automatic default value.
# cache_mode = "writethrough"

# Configuration option allocation/cache_policy.

```

```
# The default cache policy used for new cache volume.
# Since kernel 4.2 the default policy is smq (Stochastic multique),
# otherwise the older mq (Multiqueue) policy is selected.
# This configuration option does not have a default value defined.

# Configuration section allocation/cache_settings.
# Settings for the cache policy.
# See documentation for individual cache policies for more info.
# This configuration section has an automatic default value.
# cache_settings {
# }

# Configuration option allocation/cache_pool_chunk_size.
# The minimal chunk size in KiB for cache pool volumes.
# Using a chunk_size that is too large can result in wasteful use of
# the cache, where small reads and writes can cause large sections of
# an LV to be mapped into the cache. However, choosing a chunk_size
# that is too small can result in more overhead trying to manage the
# numerous chunks that become mapped into the cache. The former is
# more of a problem than the latter in most cases, so the default is
# on the smaller end of the spectrum. Supported values range from
# 32KiB to 1GiB in multiples of 32.
# This configuration option does not have a default value defined.

# Configuration option
allocation/thin_pool_metadata_require_separate_pvs.
# Thin pool metadata and data will always use different PVs.
thin_pool_metadata_require_separate_pvs = 0

# Configuration option allocation/thin_pool_zero.
# Thin pool data chunks are zeroed before they are first used.
# Zeroing with a larger thin pool chunk size reduces performance.
# This configuration option has an automatic default value.
# thin_pool_zero = 1

# Configuration option allocation/thin_pool_discards.
# The discards behaviour of thin pool volumes.
#
# Accepted values:
#   ignore
#   nopassdown
#   passdown
#
# This configuration option has an automatic default value.
# thin_pool_discards = "passdown"

# Configuration option allocation/thin_pool_chunk_size_policy.
# The chunk size calculation policy for thin pool volumes.
#
# Accepted values:
#   generic
#       If thin_pool_chunk_size is defined, use it. Otherwise, calculate
#       the chunk size based on estimation and device hints exposed in
#       sysfs - the minimum_io_size. The chunk size is always at least
#       64KiB.
#   performance
```

```

# If thin_pool_chunk_size is defined, use it. Otherwise, calculate
# the chunk size for performance based on device hints exposed in
# sysfs - the optimal_io_size. The chunk size is always at least
# 512KiB.
#
# This configuration option has an automatic default value.
# thin_pool_chunk_size_policy = "generic"

# Configuration option allocation/thin_pool_chunk_size.
# The minimal chunk size in KiB for thin pool volumes.
# Larger chunk sizes may improve performance for plain thin volumes,
# however using them for snapshot volumes is less efficient, as it
# consumes more space and takes extra time for copying. When unset,
# lvm tries to estimate chunk size starting from 64KiB. Supported
# values are in the range 64KiB to 1GiB.
# This configuration option does not have a default value defined.

# Configuration option allocation/physical_extent_size.
# Default physical extent size in KiB to use for new VGs.
# This configuration option has an automatic default value.
# physical_extent_size = 4096
}

# Configuration section log.
# How LVM log information is reported.
log {

# Configuration option log/report_command_log.
# Enable or disable LVM log reporting.
# If enabled, LVM will collect a log of operations, messages,
# per-object return codes with object identification and associated
# error numbers (errno) during LVM command processing. Then the
# log is either reported solely or in addition to any existing
# reports, depending on LVM command used. If it is a reporting command
# (e.g. pvs, vgs, lvs, lvm fullreport), then the log is reported in
# addition to any existing reports. Otherwise, there's only log report
# on output. For all applicable LVM commands, you can request that
# the output has only log report by using --logonly command line
# option. Use log/command_log_cols and log/command_log_sort settings
# to define fields to display and sort fields for the log report.
# You can also use log/command_log_selection to define selection
# criteria used each time the log is reported.
# This configuration option has an automatic default value.
# report_command_log = 0

# Configuration option log/command_log_sort.
# List of columns to sort by when reporting command log.
# See <lvm command> --logonly --configreport log -o help
# for the list of possible fields.
# This configuration option has an automatic default value.
# command_log_sort = "log_seq_num"

# Configuration option log/command_log_cols.
# List of columns to report when reporting command log.
# See <lvm command> --logonly --configreport log -o help
# for the list of possible fields.

```

```

# This configuration option has an automatic default value.
# command_log_cols =
"log_seq_num,log_type,log_context,log_object_type,log_object_name,log_object_id,log_object_group,log_object_group_id,log_message,log_errno,log_ret_code"

# Configuration option log/command_log_selection.
# Selection criteria used when reporting command log.
# You can define selection criteria that are applied each
# time log is reported. This way, it is possible to control the
# amount of log that is displayed on output and you can select
# only parts of the log that are important for you. To define
# selection criteria, use fields from log report. See also
# <lvm command> --logonly --configreport log -S help for the
# list of possible fields and selection operators. You can also
# define selection criteria for log report on command line directly
# using <lvm command> --configreport log -S <selection criteria>
# which has precedence over log/command_log_selection setting.
# For more information about selection criteria in general, see
# lvm(8) man page.
# This configuration option has an automatic default value.
# command_log_selection = "!(log_type=status && message=success)"

# Configuration option log/verbose.
# Controls the messages sent to stdout or stderr.
verbose = 0

# Configuration option log/silent.
# Suppress all non-essential messages from stdout.
# This has the same effect as -qq. When enabled, the following commands
# still produce output: dumpconfig, lvdisplay, lvmdiskscan, lvs, pvck,
# pvdisplay, pvs, version, vgcfgrestore -l, vgdisplay, vgs.
# Non-essential messages are shifted from log level 4 to log level 5
# for syslog and lvm2_log_fn purposes.
# Any 'yes' or 'no' questions not overridden by other arguments are
# suppressed and default to 'no'.
silent = 0

# Configuration option log/syslog.
# Send log messages through syslog.
syslog = 1

# Configuration option log/file.
# Write error and debug log messages to a file specified here.
# This configuration option does not have a default value defined.

# Configuration option log/overwrite.
# Overwrite the log file each time the program is run.
overwrite = 0

# Configuration option log/level.
# The level of log messages that are sent to the log file or syslog.
# There are 6 syslog-like log levels currently in use: 2 to 7 inclusive.
# 7 is the most verbose (LOG_DEBUG).
level = 0

```

```

# Configuration option log/indent.
# Indent messages according to their severity.
indent = 1

# Configuration option log/command_names.
# Display the command name on each line of output.
command_names = 0

# Configuration option log/prefix.
# A prefix to use before the log message text.
# (After the command name, if selected).
# Two spaces allows you to see/grep the severity of each message.
# To make the messages look similar to the original LVM tools use:
# indent = 0, command_names = 1, prefix = " -- "
prefix = "  "

# Configuration option log/activation.
# Log messages during activation.
# Don't use this in low memory situations (can deadlock).
activation = 0

# Configuration option log/debug_classes.
# Select log messages by class.
# Some debugging messages are assigned to a class and only appear in
# debug output if the class is listed here. Classes currently
# available: memory, devices, activation, allocation, lvmetad,
# metadata, cache, locking, lvmpolld. Use "all" to see everything.
debug_classes = [ "memory", "devices", "activation", "allocation",
"lvmetad", "metadata", "cache", "locking", "lvmpolld", "dbus" ]
}

# Configuration section backup.
# How LVM metadata is backed up and archived.
# In LVM, a 'backup' is a copy of the metadata for the current system,
# and an 'archive' contains old metadata configurations. They are
# stored in a human readable text format.
backup {

# Configuration option backup/backup.
# Maintain a backup of the current metadata configuration.
# Think very hard before turning this off!
backup = 1

# Configuration option backup/backup_dir.
# Location of the metadata backup files.
# Remember to back up this directory regularly!
backup_dir = "/etc/lvm/backup"

# Configuration option backup/archive.
# Maintain an archive of old metadata configurations.
# Think very hard before turning this off.
archive = 1

# Configuration option backup/archive_dir.
# Location of the metadata archive files.
# Remember to back up this directory regularly!

```

```
archive_dir = "/etc/lvm/archive"

# Configuration option backup/retain_min.
# Minimum number of archives to keep.
retain_min = 10

# Configuration option backup/retain_days.
# Minimum number of days to keep archive files.
retain_days = 30
}

# Configuration section shell.
# Settings for running LVM in shell (readline) mode.
shell {

    # Configuration option shell/history_size.
    # Number of lines of history to store in ~/.lvm_history.
    history_size = 100
}

# Configuration section global.
# Miscellaneous global LVM settings.
global {

    # Configuration option global/umask.
    # The file creation mask for any files and directories created.
    # Interpreted as octal if the first digit is zero.
    umask = 077

    # Configuration option global/test.
    # No on-disk metadata changes will be made in test mode.
    # Equivalent to having the -t option on every command.
    test = 0

    # Configuration option global/units.
    # Default value for --units argument.
    units = "h"

    # Configuration option global/si_unit_consistency.
    # Distinguish between powers of 1024 and 1000 bytes.
    # The LVM commands distinguish between powers of 1024 bytes,
    # e.g. KiB, MiB, GiB, and powers of 1000 bytes, e.g. KB, MB, GB.
    # If scripts depend on the old behaviour, disable this setting
    # temporarily until they are updated.
    si_unit_consistency = 1

    # Configuration option global/suffix.
    # Display unit suffix for sizes.
    # This setting has no effect if the units are in human-readable form
    # (global/units = "h") in which case the suffix is always displayed.
    suffix = 1

    # Configuration option global/activation.
    # Enable/disable communication with the kernel device-mapper.
    # Disable to use the tools to manipulate LVM metadata without
    # activating any logical volumes. If the device-mapper driver
```

```

# is not present in the kernel, disabling this should suppress
# the error messages.
activation = 1

# Configuration option global/fallback_to_lvm1.
# Try running LVM1 tools if LVM cannot communicate with DM.
# This option only applies to 2.4 kernels and is provided to help
# switch between device-mapper kernels and LVM1 kernels. The LVM1
# tools need to be installed with .lvm1 suffices, e.g. vgscan.lvm1.
# They will stop working once the lvm2 on-disk metadata format is used.
# This configuration option has an automatic default value.
# fallback_to_lvm1 = 1

# Configuration option global/format.
# The default metadata format that commands should use.
# The -M 1|2 option overrides this setting.
#
# Accepted values:
#   lvm1
#   lvm2
#
# This configuration option has an automatic default value.
# format = "lvm2"

# Configuration option global/format_libraries.
# Shared libraries that process different metadata formats.
# If support for LVM1 metadata was compiled as a shared library use
# format_libraries = "liblvm2format1.so"
# This configuration option does not have a default value defined.

# Configuration option global/segment_libraries.
# This configuration option does not have a default value defined.

# Configuration option global/proc.
# Location of proc filesystem.
# This configuration option is advanced.
proc = "/proc"

# Configuration option global/etc.
# Location of /etc system configuration directory.
etc = "/etc"

# Configuration option global/locking_type.
# Type of locking to use.
#
# Accepted values:
#   0
#   Turns off locking. Warning: this risks metadata corruption if
#   commands run concurrently.
#   1
#   LVM uses local file-based locking, the standard mode.
#   2
#   LVM uses the external shared library locking_library.
#   3
#   LVM uses built-in clustered locking with clvmd.
#   This is incompatible with lvmetad. If use_lvmetad is enabled,

```

```
# LVM prints a warning and disables lvmetad use.
# 4
# LVM uses read-only locking which forbids any operations that
# might change metadata.
# 5
# Offers dummy locking for tools that do not need any locks.
# You should not need to set this directly; the tools will select
# when to use it instead of the configured locking_type.
# Do not use lvmetad or the kernel device-mapper driver with this
# locking type. It is used by the --readonly option that offers
# read-only access to Volume Group metadata that cannot be locked
# safely because it belongs to an inaccessible domain and might be
# in use, for example a virtual machine image or a disk that is
# shared by a clustered machine.
#
locking_type = 3

# Configuration option global/wait_for_locks.
# When disabled, fail if a lock request would block.
wait_for_locks = 1

# Configuration option global/fallback_to_clustered_locking.
# Attempt to use built-in cluster locking if locking_type 2 fails.
# If using external locking (type 2) and initialisation fails, with
# this enabled, an attempt will be made to use the built-in clustered
# locking. Disable this if using a customised locking_library.
fallback_to_clustered_locking = 1

# Configuration option global/fallback_to_local_locking.
# Use locking_type 1 (local) if locking_type 2 or 3 fail.
# If an attempt to initialise type 2 or type 3 locking failed, perhaps
# because cluster components such as clvmd are not running, with this
# enabled, an attempt will be made to use local file-based locking
# (type 1). If this succeeds, only commands against local VGs will
# proceed. VGs marked as clustered will be ignored.
fallback_to_local_locking = 1

# Configuration option global/locking_dir.
# Directory to use for LVM command file locks.
# Local non-LV directory that holds file-based locks while commands are
# in progress. A directory like /tmp that may get wiped on reboot is OK.
locking_dir = "/run/lock/lvm"

# Configuration option global/prioritise_write_locks.
# Allow quicker VG write access during high volume read access.
# When there are competing read-only and read-write access requests for
# a volume group's metadata, instead of always granting the read-only
# requests immediately, delay them to allow the read-write requests to
# be serviced. Without this setting, write access may be stalled by a
# high volume of read-only requests. This option only affects
# locking_type 1 viz. local file-based locking.
prioritise_write_locks = 1

# Configuration option global/library_dir.
# Search this directory first for shared libraries.
# This configuration option does not have a default value defined.
```

```

# Configuration option global/locking_library.
# The external locking library to use for locking_type 2.
# This configuration option has an automatic default value.
# locking_library = "liblvm2clusterlock.so"

# Configuration option global/abort_on_internal_errors.
# Abort a command that encounters an internal error.
# Treat any internal errors as fatal errors, aborting the process that
# encountered the internal error. Please only enable for debugging.
abort_on_internal_errors = 0

# Configuration option global/detect_internal_vg_cache_corruption.
# Internal verification of VG structures.
# Check if CRC matches when a parsed VG is used multiple times. This
# is useful to catch unexpected changes to cached VG structures.
# Please only enable for debugging.
detect_internal_vg_cache_corruption = 0

# Configuration option global/metadata_read_only.
# No operations that change on-disk metadata are permitted.
# Additionally, read-only commands that encounter metadata in need of
# repair will still be allowed to proceed exactly as if the repair had
# been performed (except for the unchanged vg_seqno). Inappropriate
# use could mess up your system, so seek advice first!
metadata_read_only = 0

# Configuration option global/mirror_segtype_default.
# The segment type used by the short mirroring option -m.
# The --type mirror|raid1 option overrides this setting.
#
# Accepted values:
#   mirror
#       The original RAID1 implementation from LVM/DM. It is
#       characterized by a flexible log solution (core, disk, mirrored),
#       and by the necessity to block I/O while handling a failure.
#       There is an inherent race in the dmeventd failure handling logic
#       with snapshots of devices using this type of RAID1 that in the
#       worst case could cause a deadlock. (Also see
#       devices/ignore_lvm_mirrors.)
#   raid1
#       This is a newer RAID1 implementation using the MD RAID1
#       personality through device-mapper. It is characterized by a
#       lack of log options. (A log is always allocated for every
#       device and they are placed on the same device as the image,
#       so no separate devices are required.) This mirror
#       implementation does not require I/O to be blocked while
#       handling a failure. This mirror implementation is not
#       cluster-aware and cannot be used in a shared (active/active)
#       fashion in a cluster.
#
mirror_segtype_default = "raid1"

# Configuration option global/raid10_segtype_default.
# The segment type used by the -i -m combination.
# The --type raid10|mirror option overrides this setting.

```

```

# The --stripes/-i and --mirrors/-m options can both be specified
# during the creation of a logical volume to use both striping and
# mirroring for the LV. There are two different implementations.
#
# Accepted values:
#   raid10
#       LVM uses MD's RAID10 personality through DM. This is the
#       preferred option.
#   mirror
#       LVM layers the 'mirror' and 'stripe' segment types. The layering
#       is done by creating a mirror LV on top of striped sub-LVs,
#       effectively creating a RAID 0+1 array. The layering is suboptimal
#       in terms of providing redundancy and performance.
#
raid10_segtype_default = "raid10"

# Configuration option global/sparse_segtype_default.
# The segment type used by the -V -L combination.
# The --type snapshot|thin option overrides this setting.
# The combination of -V and -L options creates a sparse LV. There are
# two different implementations.
#
# Accepted values:
#   snapshot
#       The original snapshot implementation from LVM/DM. It uses an old
#       snapshot that mixes data and metadata within a single COW
#       storage volume and performs poorly when the size of stored data
#       passes hundreds of MB.
#   thin
#       A newer implementation that uses thin provisioning. It has a
#       bigger minimal chunk size (64KiB) and uses a separate volume for
#       metadata. It has better performance, especially when more data
#       is used. It also supports full snapshots.
#
sparse_segtype_default = "thin"

# Configuration option global/lvdisplay_shows_full_device_path.
# Enable this to reinstate the previous lvdisplay name format.
# The default format for displaying LV names in lvdisplay was changed
# in version 2.02.89 to show the LV name and path separately.
# Previously this was always shown as /dev/vgname/lvname even when that
# was never a valid path in the /dev filesystem.
# This configuration option has an automatic default value.
# lvdisplay_shows_full_device_path = 0

# Configuration option global/use_lvmetad.
# Use lvmetad to cache metadata and reduce disk scanning.
# When enabled (and running), lvmetad provides LVM commands with VG
# metadata and PV state. LVM commands then avoid reading this
# information from disks which can be slow. When disabled (or not
# running), LVM commands fall back to scanning disks to obtain VG
# metadata. lvmetad is kept updated via udev rules which must be set
# up for LVM to work correctly. (The udev rules should be installed
# by default.) Without a proper udev setup, changes in the system's
# block device configuration will be unknown to LVM, and ignored
# until a manual 'pvscan --cache' is run. If lvmetad was running

```

```

# while use_lvmetad was disabled, it must be stopped, use_lvmetad
# enabled, and then started. When using lvmetad, LV activation is
# switched to an automatic, event-based mode. In this mode, LVs are
# activated based on incoming udev events that inform lvmetad when
# PVs appear on the system. When a VG is complete (all PVs present),
# it is auto-activated. The auto_activation_volume_list setting
# controls which LVs are auto-activated (all by default.)
# When lvmetad is updated (automatically by udev events, or directly
# by pvscan --cache), devices/filter is ignored and all devices are
# scanned by default. lvmetad always keeps unfiltered information
# which is provided to LVM commands. Each LVM command then filters
# based on devices/filter. This does not apply to other, non-regexp,
# filtering settings: component filters such as multipath and MD
# are checked during pvscan --cache. To filter a device and prevent
# scanning from the LVM system entirely, including lvmetad, use
# devices/global_filter.
    use_lvmetad = 0

# Configuration option global/lvmetad_update_wait_time.
# The number of seconds a command will wait for lvmetad update to finish.
# After waiting for this period, a command will not use lvmetad, and
# will revert to disk scanning.
# This configuration option has an automatic default value.
# lvmetad_update_wait_time = 10

# Configuration option global/use_lvmlckd.
# Use lvmlckd for locking among hosts using LVM on shared storage.
# Applicable only if LVM is compiled with lockd support in which
# case there is also lvmlckd(8) man page available for more
# information.
use_lvmlckd = 0

# Configuration option global/lvmlckd_lock_retries.
# Retry lvmlckd lock requests this many times.
# Applicable only if LVM is compiled with lockd support
# This configuration option has an automatic default value.
# lvmlckd_lock_retries = 3

# Configuration option global/sanlock_lv_extend.
# Size in MiB to extend the internal LV holding sanlock locks.
# The internal LV holds locks for each LV in the VG, and after enough
# LVs have been created, the internal LV needs to be extended. lvcreate
# will automatically extend the internal LV when needed by the amount
# specified here. Setting this to 0 disables the automatic extension
# and can cause lvcreate to fail. Applicable only if LVM is compiled
# with lockd support
# This configuration option has an automatic default value.
# sanlock_lv_extend = 256

# Configuration option global/thin_check_executable.
# The full path to the thin_check command.
# LVM uses this command to check that a thin metadata device is in a
# usable state. When a thin pool is activated and after it is
# deactivated, this command is run. Activation will only proceed if
# the command has an exit status of 0. Set to "" to skip this check.
# (Not recommended.) Also see thin_check_options.

```

```
# (See package device-mapper-persistent-data or thin-provisioning-tools)
# This configuration option has an automatic default value.
# thin_check_executable = "/usr/sbin/thin_check"

# Configuration option global/thin_dump_executable.
# The full path to the thin_dump command.
# LVM uses this command to dump thin pool metadata.
# (See package device-mapper-persistent-data or thin-provisioning-tools)
# This configuration option has an automatic default value.
# thin_dump_executable = "/usr/sbin/thin_dump"

# Configuration option global/thin_repair_executable.
# The full path to the thin_repair command.
# LVM uses this command to repair a thin metadata device if it is in
# an unusable state. Also see thin_repair_options.
# (See package device-mapper-persistent-data or thin-provisioning-tools)
# This configuration option has an automatic default value.
# thin_repair_executable = "/usr/sbin/thin_repair"

# Configuration option global/thin_check_options.
# List of options passed to the thin_check command.
# With thin_check version 2.1 or newer you can add the option
# --ignore-non-fatal-errors to let it pass through ignorable errors
# and fix them later. With thin_check version 3.2 or newer you should
# include the option --clear-needs-check-flag.
# This configuration option has an automatic default value.
# thin_check_options = [ "-q", "--clear-needs-check-flag" ]

# Configuration option global/thin_repair_options.
# List of options passed to the thin_repair command.
# This configuration option has an automatic default value.
# thin_repair_options = [ "" ]

# Configuration option global/thin_disabled_features.
# Features to not use in the thin driver.
# This can be helpful for testing, or to avoid using a feature that is
# causing problems. Features include: block_size, discards,
# discards_non_power_2, external_origin, metadata_resize,
# external_origin_extend, error_if_no_space.
#
# Example
# thin_disabled_features = [ "discards", "block_size" ]
#
# This configuration option does not have a default value defined.

# Configuration option global/cache_disabled_features.
# Features to not use in the cache driver.
# This can be helpful for testing, or to avoid using a feature that is
# causing problems. Features include: policy_mq, policy_smq.
#
# Example
# cache_disabled_features = [ "policy_smq" ]
#
# This configuration option does not have a default value defined.

# Configuration option global/cache_check_executable.
```

```

# The full path to the cache_check command.
# LVM uses this command to check that a cache metadata device is in a
# usable state. When a cached LV is activated and after it is
# deactivated, this command is run. Activation will only proceed if the
# command has an exit status of 0. Set to "" to skip this check.
# (Not recommended.) Also see cache_check_options.
# (See package device-mapper-persistent-data or thin-provisioning-tools)
# This configuration option has an automatic default value.
# cache_check_executable = "/usr/sbin/cache_check"

# Configuration option global/cache_dump_executable.
# The full path to the cache_dump command.
# LVM uses this command to dump cache pool metadata.
# (See package device-mapper-persistent-data or thin-provisioning-tools)
# This configuration option has an automatic default value.
# cache_dump_executable = "/usr/sbin/cache_dump"

# Configuration option global/cache_repair_executable.
# The full path to the cache_repair command.
# LVM uses this command to repair a cache metadata device if it is in
# an unusable state. Also see cache_repair_options.
# (See package device-mapper-persistent-data or thin-provisioning-tools)
# This configuration option has an automatic default value.
# cache_repair_executable = "/usr/sbin/cache_repair"

# Configuration option global/cache_check_options.
# List of options passed to the cache_check command.
# With cache_check version 5.0 or newer you should include the option
# --clear-needs-check-flag.
# This configuration option has an automatic default value.
# cache_check_options = [ "-q", "--clear-needs-check-flag" ]

# Configuration option global/cache_repair_options.
# List of options passed to the cache_repair command.
# This configuration option has an automatic default value.
# cache_repair_options = [ "" ]

# Configuration option global/system_id_source.
# The method LVM uses to set the local system ID.
# Volume Groups can also be given a system ID (by vgcreate, vgchange,
# or vgimport.) A VG on shared storage devices is accessible only to
# the host with a matching system ID. See 'man lvmsystemid' for
# information on limitations and correct usage.
#
# Accepted values:
#   none
#       The host has no system ID.
#   lvmlocal
#       Obtain the system ID from the system_id setting in the 'local'
#       section of an lvm configuration file, e.g. lvmlocal.conf.
#   uname
#       Set the system ID from the hostname (uname) of the system.
#       System IDs beginning localhost are not permitted.
#   machineid
#       Use the contents of the machine-id file to set the system ID.
#       Some systems create this file at installation time.

```

```

# See 'man machine-id' and global/etc.
# file
# Use the contents of another file (system_id_file) to set the
# system ID.
#
system_id_source = "none"

# Configuration option global/system_id_file.
# The full path to the file containing a system ID.
# This is used when system_id_source is set to 'file'.
# Comments starting with the character # are ignored.
# This configuration option does not have a default value defined.

# Configuration option global/use_lvmpolld.
# Use lvmpolld to supervise long running LVM commands.
# When enabled, control of long running LVM commands is transferred
# from the original LVM command to the lvmpolld daemon. This allows
# the operation to continue independent of the original LVM command.
# After lvmpolld takes over, the LVM command displays the progress
# of the ongoing operation. lvmpolld itself runs LVM commands to
# manage the progress of ongoing operations. lvmpolld can be used as
# a native systemd service, which allows it to be started on demand,
# and to use its own control group. When this option is disabled, LVM
# commands will supervise long running operations by forking themselves.
# Applicable only if LVM is compiled with lvmpolld support.
use_lvmpolld = 1

# Configuration option global/notify_dbus.
# Enable D-Bus notification from LVM commands.
# When enabled, an LVM command that changes PVs, changes VG metadata,
# or changes the activation state of an LV will send a notification.
notify_dbus = 1
}

# Configuration section activation.
activation {

# Configuration option activation/checks.
# Perform internal checks of libdevmapper operations.
# Useful for debugging problems with activation. Some of the checks may
# be expensive, so it's best to use this only when there seems to be a
# problem.
checks = 0

# Configuration option activation/udev_sync.
# Use udev notifications to synchronize udev and LVM.
# The --nodevsysync option overrides this setting.
# When disabled, LVM commands will not wait for notifications from
# udev, but continue irrespective of any possible udev processing in
# the background. Only use this if udev is not running or has rules
# that ignore the devices LVM creates. If enabled when udev is not
# running, and LVM processes are waiting for udev, run the command
# 'dmsetup udevcomplete_all' to wake them up.
udev_sync = 1

# Configuration option activation/udev_rules.

```

```

# Use udev rules to manage LV device nodes and symlinks.
# When disabled, LVM will manage the device nodes and symlinks for
# active LVs itself. Manual intervention may be required if this
# setting is changed while LVs are active.
udev_rules = 1

# Configuration option activation/verify_udev_operations.
# Use extra checks in LVM to verify udev operations.
# This enables additional checks (and if necessary, repairs) on entries
# in the device directory after udev has completed processing its
# events. Useful for diagnosing problems with LVM/udev interactions.
verify_udev_operations = 0

# Configuration option activation/retry_deactivation.
# Retry failed LV deactivation.
# If LV deactivation fails, LVM will retry for a few seconds before
# failing. This may happen because a process run from a quick udev rule
# temporarily opened the device.
retry_deactivation = 1

# Configuration option activation/missing_stripe_filler.
# Method to fill missing stripes when activating an incomplete LV.
# Using 'error' will make inaccessible parts of the device return I/O
# errors on access. You can instead use a device path, in which case,
# that device will be used in place of missing stripes. Using anything
# other than 'error' with mirrored or snapshotted volumes is likely to
# result in data corruption.
# This configuration option is advanced.
missing_stripe_filler = "error"

# Configuration option activation/use_linear_target.
# Use the linear target to optimize single stripe LVs.
# When disabled, the striped target is used. The linear target is an
# optimised version of the striped target that only handles a single
# stripe.
use_linear_target = 1

# Configuration option activation/reserved_stack.
# Stack size in KiB to reserve for use while devices are suspended.
# Insufficient reserve risks I/O deadlock during device suspension.
reserved_stack = 64

# Configuration option activation/reserved_memory.
# Memory size in KiB to reserve for use while devices are suspended.
# Insufficient reserve risks I/O deadlock during device suspension.
reserved_memory = 8192

# Configuration option activation/process_priority.
# Nice value used while devices are suspended.
# Use a high priority so that LVs are suspended
# for the shortest possible time.
process_priority = -18

# Configuration option activation/volume_list.
# Only LVs selected by this list are activated.
# If this list is defined, an LV is only activated if it matches an

```

```

# entry in this list. If this list is undefined, it imposes no limits
# on LV activation (all are allowed).
#
# Accepted values:
#   vgname
#       The VG name is matched exactly and selects all LVs in the VG.
#   vgname/lvname
#       The VG name and LV name are matched exactly and selects the LV.
#   @tag
#       Selects an LV if the specified tag matches a tag set on the LV
#       or VG.
#   @*
#       Selects an LV if a tag defined on the host is also set on the LV
#       or VG. See tags/hosttags. If any host tags exist but volume_list
#       is not defined, a default single-entry list containing '@*'
#       is assumed.
#
# Example
# volume_list = [ "vg1", "vg2/lvol1", "@tag1", "@*" ]
#
# This configuration option does not have a default value defined.

# Configuration option activation/auto_activation_volume_list.
# Only LVs selected by this list are auto-activated.
# This list works like volume_list, but it is used only by
# auto-activation commands. It does not apply to direct activation
# commands. If this list is defined, an LV is only auto-activated
# if it matches an entry in this list. If this list is undefined, it
# imposes no limits on LV auto-activation (all are allowed.) If this
# list is defined and empty, i.e. "[]", then no LVs are selected for
# auto-activation. An LV that is selected by this list for
# auto-activation, must also be selected by volume_list (if defined)
# before it is activated. Auto-activation is an activation command that
# includes the 'a' argument: --activate ay or -a ay. The 'a' (auto)
# argument for auto-activation is meant to be used by activation
# commands that are run automatically by the system, as opposed to LVM
# commands run directly by a user. A user may also use the 'a' flag
# directly to perform auto-activation. Also see pvscan(8) for more
# information about auto-activation.
#
# Accepted values:
#   vgname
#       The VG name is matched exactly and selects all LVs in the VG.
#   vgname/lvname
#       The VG name and LV name are matched exactly and selects the LV.
#   @tag
#       Selects an LV if the specified tag matches a tag set on the LV
#       or VG.
#   @*
#       Selects an LV if a tag defined on the host is also set on the LV
#       or VG. See tags/hosttags. If any host tags exist but volume_list
#       is not defined, a default single-entry list containing '@*'
#       is assumed.
#
# Example
# auto_activation_volume_list = [ "vg1", "vg2/lvol1", "@tag1", "@*" ]

```

```

#
# This configuration option does not have a default value defined.

# Configuration option activation/read_only_volume_list.
# LVs in this list are activated in read-only mode.
# If this list is defined, each LV that is to be activated is checked
# against this list, and if it matches, it is activated in read-only
# mode. This overrides the permission setting stored in the metadata,
# e.g. from --permission rw.
#
# Accepted values:
#   vname
#       The VG name is matched exactly and selects all LVs in the VG.
#   vname/lvname
#       The VG name and LV name are matched exactly and selects the LV.
#   @tag
#       Selects an LV if the specified tag matches a tag set on the LV
#       or VG.
#   @*
#       Selects an LV if a tag defined on the host is also set on the LV
#       or VG. See tags/hosttags. If any host tags exist but volume_list
#       is not defined, a default single-entry list containing '@*'
#       is assumed.
#
# Example
# read_only_volume_list = [ "vg1", "vg2/lvol1", "@tag1", "@*" ]
#
# This configuration option does not have a default value defined.

# Configuration option activation/raid_region_size.
# Size in KiB of each raid or mirror synchronization region.
# For raid or mirror segment types, this is the amount of data that is
# copied at once when initializing, or moved at once by pvmove.
raid_region_size = 512

# Configuration option activation/error_when_full.
# Return errors if a thin pool runs out of space.
# The --errorwhenfull option overrides this setting.
# When enabled, writes to thin LVs immediately return an error if the
# thin pool is out of data space. When disabled, writes to thin LVs
# are queued if the thin pool is out of space, and processed when the
# thin pool data space is extended. New thin pools are assigned the
# behavior defined here.
# This configuration option has an automatic default value.
# error_when_full = 0

# Configuration option activation/readahead.
# Setting to use when there is no readahead setting in metadata.
#
# Accepted values:
#   none
#       Disable readahead.
#   auto
#       Use default value chosen by kernel.
#
readahead = "auto"

```

```

# Configuration option activation/raid_fault_policy.
# Defines how a device failure in a RAID LV is handled.
# This includes LVs that have the following segment types:
# raid1, raid4, raid5*, and raid6*.
# If a device in the LV fails, the policy determines the steps
# performed by dmeventd automatically, and the steps performed by the
# manual command lvconvert --repair --use-policies.
# Automatic handling requires dmeventd to be monitoring the LV.
#
# Accepted values:
#   warn
#       Use the system log to warn the user that a device in the RAID LV
#       has failed. It is left to the user to run lvconvert --repair
#       manually to remove or replace the failed device. As long as the
#       number of failed devices does not exceed the redundancy of the LV
#       (1 device for raid4/5, 2 for raid6), the LV will remain usable.
#   allocate
#       Attempt to use any extra physical volumes in the VG as spares and
#       replace faulty devices.
#
raid_fault_policy = "warn"

# Configuration option activation/mirror_image_fault_policy.
# Defines how a device failure in a 'mirror' LV is handled.
# An LV with the 'mirror' segment type is composed of mirror images
# (copies) and a mirror log. A disk log ensures that a mirror LV does
# not need to be re-synced (all copies made the same) every time a
# machine reboots or crashes. If a device in the LV fails, this policy
# determines the steps performed by dmeventd automatically, and the steps
# performed by the manual command lvconvert --repair --use-policies.
# Automatic handling requires dmeventd to be monitoring the LV.
#
# Accepted values:
#   remove
#       Simply remove the faulty device and run without it. If the log
#       device fails, the mirror would convert to using an in-memory log.
#       This means the mirror will not remember its sync status across
#       crashes/reboots and the entire mirror will be re-synced. If a
#       mirror image fails, the mirror will convert to a non-mirrored
#       device if there is only one remaining good copy.
#   allocate
#       Remove the faulty device and try to allocate space on a new
#       device to be a replacement for the failed device. Using this
#       policy for the log is fast and maintains the ability to remember
#       sync state through crashes/reboots. Using this policy for a
#       mirror device is slow, as it requires the mirror to resynchronize
#       the devices, but it will preserve the mirror characteristic of
#       the device. This policy acts like 'remove' if no suitable device
#       and space can be allocated for the replacement.
#   allocate_anywhere
#       Not yet implemented. Useful to place the log device temporarily
#       on the same physical volume as one of the mirror images. This
#       policy is not recommended for mirror devices since it would break
#       the redundant nature of the mirror. This policy acts like
#       'remove' if no suitable device and space can be allocated for the

```

```

# replacement.
#
mirror_image_fault_policy = "remove"

# Configuration option activation/mirror_log_fault_policy.
# Defines how a device failure in a 'mirror' log LV is handled.
# The mirror_image_fault_policy description for mirrored LVs also
# applies to mirrored log LVs.
mirror_log_fault_policy = "allocate"

# Configuration option activation/snapshot_autoextend_threshold.
# Auto-extend a snapshot when its usage exceeds this percent.
# Setting this to 100 disables automatic extension.
# The minimum value is 50 (a smaller value is treated as 50.)
# Also see snapshot_autoextend_percent.
# Automatic extension requires dmeventd to be monitoring the LV.
#
# Example
# Using 70% autoextend threshold and 20% autoextend size, when a 1G
# snapshot exceeds 700M, it is extended to 1.2G, and when it exceeds
# 840M, it is extended to 1.44G:
# snapshot_autoextend_threshold = 70
#
snapshot_autoextend_threshold = 100

# Configuration option activation/snapshot_autoextend_percent.
# Auto-extending a snapshot adds this percent extra space.
# The amount of additional space added to a snapshot is this
# percent of its current size.
#
# Example
# Using 70% autoextend threshold and 20% autoextend size, when a 1G
# snapshot exceeds 700M, it is extended to 1.2G, and when it exceeds
# 840M, it is extended to 1.44G:
# snapshot_autoextend_percent = 20
#
snapshot_autoextend_percent = 20

# Configuration option activation/thin_pool_autoextend_threshold.
# Auto-extend a thin pool when its usage exceeds this percent.
# Setting this to 100 disables automatic extension.
# The minimum value is 50 (a smaller value is treated as 50.)
# Also see thin_pool_autoextend_percent.
# Automatic extension requires dmeventd to be monitoring the LV.
#
# Example
# Using 70% autoextend threshold and 20% autoextend size, when a 1G
# thin pool exceeds 700M, it is extended to 1.2G, and when it exceeds
# 840M, it is extended to 1.44G:
# thin_pool_autoextend_threshold = 70
#
thin_pool_autoextend_threshold = 100

# Configuration option activation/thin_pool_autoextend_percent.
# Auto-extending a thin pool adds this percent extra space.
# The amount of additional space added to a thin pool is this

```

```

# percent of its current size.
#
# Example
# Using 70% autoextend threshold and 20% autoextend size, when a 1G
# thin pool exceeds 700M, it is extended to 1.2G, and when it exceeds
# 840M, it is extended to 1.44G:
# thin_pool_autoextend_percent = 20
#
thin_pool_autoextend_percent = 20

# Configuration option activation/mlock_filter.
# Do not mlock these memory areas.
# While activating devices, I/O to devices being (re)configured is
# suspended. As a precaution against deadlocks, LVM pins memory it is
# using so it is not paged out, and will not require I/O to reread.
# Groups of pages that are known not to be accessed during activation
# do not need to be pinned into memory. Each string listed in this
# setting is compared against each line in /proc/self/maps, and the
# pages corresponding to lines that match are not pinned. On some
# systems, locale-archive was found to make up over 80% of the memory
# used by the process.
#
# Example
# mlock_filter = [ "locale/locale-archive", "gconv/gconv-modules.cache" ]
#
# This configuration option is advanced.
# This configuration option does not have a default value defined.

# Configuration option activation/use_mlockall.
# Use the old behavior of mlockall to pin all memory.
# Prior to version 2.02.62, LVM used mlockall() to pin the whole
# process's memory while activating devices.
use_mlockall = 0

# Configuration option activation/monitoring.
# Monitor LVs that are activated.
# The --ignoremonitoring option overrides this setting.
# When enabled, LVM will ask dmeventd to monitor activated LVs.
monitoring = 1

# Configuration option activation/polling_interval.
# Check pvmove or lvconvert progress at this interval (seconds).
# When pvmove or lvconvert must wait for the kernel to finish
# synchronising or merging data, they check and report progress at
# intervals of this number of seconds. If this is set to 0 and there
# is only one thing to wait for, there are no progress reports, but
# the process is awoken immediately once the operation is complete.
polling_interval = 15

# Configuration option activation/auto_set_activation_skip.
# Set the activation skip flag on new thin snapshot LVs.
# The --setactivationskip option overrides this setting.
# An LV can have a persistent 'activation skip' flag. The flag causes
# the LV to be skipped during normal activation. The lvchange/vgchange
# -K option is required to activate LVs that have the activation skip
# flag set. When this setting is enabled, the activation skip flag is

```

```

# set on new thin snapshot LVs.
# This configuration option has an automatic default value.
# auto_set_activation_skip = 1

# Configuration option activation/activation_mode.
# How LVs with missing devices are activated.
# The --activationmode option overrides this setting.
#
# Accepted values:
#   complete
#       Only allow activation of an LV if all of the Physical Volumes it
#       uses are present. Other PVs in the Volume Group may be missing.
#   degraded
#       Like complete, but additionally RAID LVs of segment type raid1,
#       raid4, raid5, raid6 and raid10 will be activated if there is no
#       data loss, i.e. they have sufficient redundancy to present the
#       entire addressable range of the Logical Volume.
#   partial
#       Allows the activation of any LV even if a missing or failed PV
#       could cause data loss with a portion of the LV inaccessible.
#       This setting should not normally be used, but may sometimes
#       assist with data recovery.
#
activation_mode = "degraded"

# Configuration option activation/lock_start_list.
# Locking is started only for VGs selected by this list.
# The rules are the same as those for volume_list.
# This configuration option does not have a default value defined.

# Configuration option activation/auto_lock_start_list.
# Locking is auto-started only for VGs selected by this list.
# The rules are the same as those for auto_activation_volume_list.
# This configuration option does not have a default value defined.
}

# Configuration section metadata.
# This configuration section has an automatic default value.
# metadata {

# Configuration option metadata/check_pv_device_sizes.
# Check device sizes are not smaller than corresponding PV sizes.
# If device size is less than corresponding PV size found in metadata,
# there is always a risk of data loss. If this option is set, then LVM
# issues a warning message each time it finds that the device size is
# less than corresponding PV size. You should not disable this unless
# you are absolutely sure about what you are doing!
# This configuration option is advanced.
# This configuration option has an automatic default value.
# check_pv_device_sizes = 1

# Configuration option metadata/record_lvs_history.
# When enabled, LVM keeps history records about removed LVs in
# metadata. The information that is recorded in metadata for
# historical LVs is reduced when compared to original
# information kept in metadata for live LVs. Currently, this

```

```
# feature is supported for thin and thin snapshot LVs only.
# This configuration option has an automatic default value.
# record_lvs_history = 0

# Configuration option metadata/lvs_history_retention_time.
# Retention time in seconds after which a record about individual
# historical logical volume is automatically destroyed.
# A value of 0 disables this feature.
# This configuration option has an automatic default value.
# lvs_history_retention_time = 0

# Configuration option metadata/pvmetadacopies.
# Number of copies of metadata to store on each PV.
# The --pvmetadacopies option overrides this setting.
#
# Accepted values:
# 2
#     Two copies of the VG metadata are stored on the PV, one at the
#     front of the PV, and one at the end.
# 1
#     One copy of VG metadata is stored at the front of the PV.
# 0
#     No copies of VG metadata are stored on the PV. This may be
#     useful for VGs containing large numbers of PVs.
#
# This configuration option is advanced.
# This configuration option has an automatic default value.
# pvmetadacopies = 1

# Configuration option metadata/vgmetadacopies.
# Number of copies of metadata to maintain for each VG.
# The --vgmetadacopies option overrides this setting.
# If set to a non-zero value, LVM automatically chooses which of the
# available metadata areas to use to achieve the requested number of
# copies of the VG metadata. If you set a value larger than the the
# total number of metadata areas available, then metadata is stored in
# them all. The value 0 (unmanaged) disables this automatic management
# and allows you to control which metadata areas are used at the
# individual PV level using pvchange --metadataignore y|n.
# This configuration option has an automatic default value.
# vgmetadacopies = 0

# Configuration option metadata/pvmetadatasize.
# Approximate number of sectors to use for each metadata copy.
# VGs with large numbers of PVs or LVs, or VGs containing complex LV
# structures, may need additional space for VG metadata. The metadata
# areas are treated as circular buffers, so unused space becomes filled
# with an archive of the most recent previous versions of the metadata.
# This configuration option has an automatic default value.
# pvmetadatasize = 255

# Configuration option metadata/pvmetadataignore.
# Ignore metadata areas on a new PV.
# The --metadataignore option overrides this setting.
# If metadata areas on a PV are ignored, LVM will not store metadata
# in them.
```

```

# This configuration option is advanced.
# This configuration option has an automatic default value.
# pvmetadataignore = 0

# Configuration option metadata/stripesize.
# This configuration option is advanced.
# This configuration option has an automatic default value.
# stripesize = 64

# Configuration option metadata/dirs.
# Directories holding live copies of text format metadata.
# These directories must not be on logical volumes!
# It's possible to use LVM with a couple of directories here,
# preferably on different (non-LV) filesystems, and with no other
# on-disk metadata (pvmetadatacopies = 0). Or this can be in addition
# to on-disk metadata areas. The feature was originally added to
# simplify testing and is not supported under low memory situations -
# the machine could lock up. Never edit any files in these directories
# by hand unless you are absolutely sure you know what you are doing!
# Use the supplied toolset to make changes (e.g. vgcfgrestore).
#
# Example
# dirs= [ "/etc/lvm/metadata", "/mnt/disk2/lvm/metadata2" ]
#
# This configuration option is advanced.
# This configuration option does not have a default value defined.
# }

# Configuration section report.
# LVM report command output formatting.
# This configuration section has an automatic default value.
# report {

# Configuration option report/output_format.
# Format of LVM command's report output.
# If there is more than one report per command, then the format
# is applied for all reports. You can also change output format
# directly on command line using --reportformat option which
# has precedence over log/output_format setting.
# Accepted values:
#   basic
#       Original format with columns and rows. If there is more than
#       one report per command, each report is prefixed with report's
#       name for identification.
#   json
#       JSON format.
# This configuration option has an automatic default value.
# output_format = "basic"

# Configuration option report/compact_output.
# Do not print empty values for all report fields.
# If enabled, all fields that don't have a value set for any of the
# rows reported are skipped and not printed. Compact output is
# applicable only if report/buffered is enabled. If you need to
# compact only specified fields, use compact_output=0 and define
# report/compact_output_cols configuration setting instead.

```

```
# This configuration option has an automatic default value.
# compact_output = 0

# Configuration option report/compact_output_cols.
# Do not print empty values for specified report fields.
# If defined, specified fields that don't have a value set for any
# of the rows reported are skipped and not printed. Compact output
# is applicable only if report/buffered is enabled. If you need to
# compact all fields, use compact_output=1 instead in which case
# the compact_output_cols setting is then ignored.
# This configuration option has an automatic default value.
# compact_output_cols = ""

# Configuration option report/aligned.
# Align columns in report output.
# This configuration option has an automatic default value.
# aligned = 1

# Configuration option report/buffered.
# Buffer report output.
# When buffered reporting is used, the report's content is appended
# incrementally to include each object being reported until the report
# is flushed to output which normally happens at the end of command
# execution. Otherwise, if buffering is not used, each object is
# reported as soon as its processing is finished.
# This configuration option has an automatic default value.
# buffered = 1

# Configuration option report/headings.
# Show headings for columns on report.
# This configuration option has an automatic default value.
# headings = 1

# Configuration option report/separator.
# A separator to use on report after each field.
# This configuration option has an automatic default value.
# separator = " "

# Configuration option report/list_item_separator.
# A separator to use for list items when reported.
# This configuration option has an automatic default value.
# list_item_separator = ","

# Configuration option report/prefixes.
# Use a field name prefix for each field reported.
# This configuration option has an automatic default value.
# prefixes = 0

# Configuration option report/quoted.
# Quote field values when using field name prefixes.
# This configuration option has an automatic default value.
# quoted = 1

# Configuration option report/columns_as_rows.
# Output each column as a row.
# If set, this also implies report/prefixes=1.
```

```

# This configuration option has an automatic default value.
# columns_as_rows = 0

# Configuration option report/binary_values_as_numeric.
# Use binary values 0 or 1 instead of descriptive literal values.
# For columns that have exactly two valid values to report
# (not counting the 'unknown' value which denotes that the
# value could not be determined).
# This configuration option has an automatic default value.
# binary_values_as_numeric = 0

# Configuration option report/time_format.
# Set time format for fields reporting time values.
# Format specification is a string which may contain special character
# sequences and ordinary character sequences. Ordinary character
# sequences are copied verbatim. Each special character sequence is
# introduced by the '%' character and such sequence is then
# substituted with a value as described below.
#
# Accepted values:
# %a
#     The abbreviated name of the day of the week according to the
#     current locale.
# %A
#     The full name of the day of the week according to the current
#     locale.
# %b
#     The abbreviated month name according to the current locale.
# %B
#     The full month name according to the current locale.
# %c
#     The preferred date and time representation for the current
#     locale (alt E)
# %C
#     The century number (year/100) as a 2-digit integer. (alt E)
# %d
#     The day of the month as a decimal number (range 01 to 31).
#     (alt 0)
# %D
#     Equivalent to %m/%d/%y. (For Americans only. Americans should
#     note that in other countries %d/%m/%y is rather common. This
#     means that in international context this format is ambiguous and
#     should not be used.
# %e
#     Like %d, the day of the month as a decimal number, but a leading
#     zero is replaced by a space. (alt 0)
# %E
#     Modifier: use alternative local-dependent representation if
#     available.
# %F
#     Equivalent to %Y-%m-%d (the ISO 8601 date format).
# %G
#     The ISO 8601 week-based year with century as a decimal number.
#     The 4-digit year corresponding to the ISO week number (see %V).
#     This has the same format and value as %Y, except that if the
#     ISO week number belongs to the previous or next year, that year

```

```
#      is used instead.
# %g
#      Like %G, but without century, that is, with a 2-digit year
#      (00-99).
# %h
#      Equivalent to %b.
# %H
#      The hour as a decimal number using a 24-hour clock
#      (range 00 to 23). (alt 0)
# %I
#      The hour as a decimal number using a 12-hour clock
#      (range 01 to 12). (alt 0)
# %j
#      The day of the year as a decimal number (range 001 to 366).
# %k
#      The hour (24-hour clock) as a decimal number (range 0 to 23);
#      single digits are preceded by a blank. (See also %H.)
# %l
#      The hour (12-hour clock) as a decimal number (range 1 to 12);
#      single digits are preceded by a blank. (See also %I.)
# %m
#      The month as a decimal number (range 01 to 12). (alt 0)
# %M
#      The minute as a decimal number (range 00 to 59). (alt 0)
# %O
#      Modifier: use alternative numeric symbols.
# %p
#      Either "AM" or "PM" according to the given time value,
#      or the corresponding strings for the current locale. Noon is
#      treated as "PM" and midnight as "AM".
# %P
#      Like %p but in lowercase: "am" or "pm" or a corresponding
#      string for the current locale.
# %r
#      The time in a.m. or p.m. notation. In the POSIX locale this is
#      equivalent to %I:%M:%S %p.
# %R
#      The time in 24-hour notation (%H:%M). For a version including
#      the seconds, see %T below.
# %s
#      The number of seconds since the Epoch,
#      1970-01-01 00:00:00 +0000 (UTC)
# %S
#      The second as a decimal number (range 00 to 60). (The range is
#      up to 60 to allow for occasional leap seconds.) (alt 0)
# %t
#      A tab character.
# %T
#      The time in 24-hour notation (%H:%M:%S).
# %u
#      The day of the week as a decimal, range 1 to 7, Monday being 1.
#      See also %w. (alt 0)
# %U
#      The week number of the current year as a decimal number,
#      range 00 to 53, starting with the first Sunday as the first
#      day of week 01. See also %V and %W. (alt 0)
```

```

# %V
#   The ISO 8601 week number of the current year as a decimal number,
#   range 01 to 53, where week 1 is the first week that has at least
#   4 days in the new year. See also %U and %W. (alt 0)
# %w
#   The day of the week as a decimal, range 0 to 6, Sunday being 0.
#   See also %u. (alt 0)
# %W
#   The week number of the current year as a decimal number,
#   range 00 to 53, starting with the first Monday as the first day
#   of week 01. (alt 0)
# %x
#   The preferred date representation for the current locale without
#   the time. (alt E)
# %X
#   The preferred time representation for the current locale without
#   the date. (alt E)
# %y
#   The year as a decimal number without a century (range 00 to 99).
#   (alt E, alt 0)
# %Y
#   The year as a decimal number including the century. (alt E)
# %Z
#   The +hhmm or -hhmm numeric timezone (that is, the hour and minute
#   offset from UTC).
# %Z
#   The timezone name or abbreviation.
# %%
#   A literal '%' character.
#
# This configuration option has an automatic default value.
# time_format = "%Y-%m-%d %T %Z"

# Configuration option report/devtypes_sort.
# List of columns to sort by when reporting 'lvm devtypes' command.
# See 'lvm devtypes -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# devtypes_sort = "devtype_name"

# Configuration option report/devtypes_cols.
# List of columns to report for 'lvm devtypes' command.
# See 'lvm devtypes -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# devtypes_cols =
"devtype_name,devtype_max_partitions,devtype_description"

# Configuration option report/devtypes_cols_verbose.
# List of columns to report for 'lvm devtypes' command in verbose mode.
# See 'lvm devtypes -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# devtypes_cols_verbose =
"devtype_name,devtype_max_partitions,devtype_description"

# Configuration option report/lvs_sort.
# List of columns to sort by when reporting 'lvs' command.
# See 'lvs -o help' for the list of possible fields.

```

```
# This configuration option has an automatic default value.
# lvs_sort = "vg_name,lv_name"

# Configuration option report/lvs_cols.
# List of columns to report for 'lvs' command.
# See 'lvs -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# lvs_cols =
"lv_name,vg_name,lv_attr,lv_size,pool_lv,origin,data_percent,metadata_percent,move_pv,mirror_log,copy_percent,convert_lv"

# Configuration option report/lvs_cols_verbose.
# List of columns to report for 'lvs' command in verbose mode.
# See 'lvs -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# lvs_cols_verbose =
"lv_name,vg_name,seg_count,lv_attr,lv_size,lv_major,lv_minor,lv_kernel_major,lv_kernel_minor,pool_lv,origin,data_percent,metadata_percent,move_pv,copy_percent,mirror_log,convert_lv,lv_uuid,lv_profile"

# Configuration option report/vgs_sort.
# List of columns to sort by when reporting 'vgs' command.
# See 'vgs -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# vgs_sort = "vg_name"

# Configuration option report/vgs_cols.
# List of columns to report for 'vgs' command.
# See 'vgs -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# vgs_cols =
"vg_name,pv_count,lv_count,snap_count,vg_attr,vg_size,vg_free"

# Configuration option report/vgs_cols_verbose.
# List of columns to report for 'vgs' command in verbose mode.
# See 'vgs -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# vgs_cols_verbose =
"vg_name,vg_attr,vg_extent_size,pv_count,lv_count,snap_count,vg_size,vg_free,vg_uuid,vg_profile"

# Configuration option report/pvs_sort.
# List of columns to sort by when reporting 'pvs' command.
# See 'pvs -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# pvs_sort = "pv_name"

# Configuration option report/pvs_cols.
# List of columns to report for 'pvs' command.
# See 'pvs -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# pvs_cols = "pv_name,vg_name,pv_fmt,pv_attr,pv_size,pv_free"

# Configuration option report/pvs_cols_verbose.
# List of columns to report for 'pvs' command in verbose mode.
# See 'pvs -o help' for the list of possible fields.
```

```

# This configuration option has an automatic default value.
# pvs_cols_verbose =
"pv_name,vg_name,pv_fmt,pv_attr,pv_size,pv_free,dev_size,pv_uuid"

# Configuration option report/secs_sort.
# List of columns to sort by when reporting 'lvs --segments' command.
# See 'lvs --segments -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# secs_sort = "vg_name,lv_name,seg_start"

# Configuration option report/secs_cols.
# List of columns to report for 'lvs --segments' command.
# See 'lvs --segments -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# secs_cols = "lv_name,vg_name,lv_attr,stripes,segtype,seg_size"

# Configuration option report/secs_cols_verbose.
# List of columns to report for 'lvs --segments' command in verbose mode.
# See 'lvs --segments -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# secs_cols_verbose =
"lv_name,vg_name,lv_attr,seg_start,seg_size,stripes,segtype,stripesize,chunksize"

# Configuration option report/pvsecs_sort.
# List of columns to sort by when reporting 'pvs --segments' command.
# See 'pvs --segments -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# pvsecs_sort = "pv_name,pvseg_start"

# Configuration option report/pvsecs_cols.
# List of columns to sort by when reporting 'pvs --segments' command.
# See 'pvs --segments -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# pvsecs_cols =
"pv_name,vg_name,pv_fmt,pv_attr,pv_size,pv_free,pvseg_start,pvseg_size"

# Configuration option report/pvsecs_cols_verbose.
# List of columns to sort by when reporting 'pvs --segments' command in verbose mode.
# See 'pvs --segments -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# pvsecs_cols_verbose =
"pv_name,vg_name,pv_fmt,pv_attr,pv_size,pv_free,pvseg_start,pvseg_size,lv_name,seg_start_pe,segtype,seg_pe_ranges"

# Configuration option report/vgs_cols_full.
# List of columns to report for lvm fullreport's 'vgs' subreport.
# See 'vgs -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# vgs_cols_full = "vg_all"

# Configuration option report/pvs_cols_full.
# List of columns to report for lvm fullreport's 'vgs' subreport.
# See 'pvs -o help' for the list of possible fields.
# This configuration option has an automatic default value.

```

```
# pvs_cols_full = "pv_all"

# Configuration option report/lvs_cols_full.
# List of columns to report for lvm fullreport's 'lvs' subreport.
# See 'lvs -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# lvs_cols_full = "lv_all"

# Configuration option report/pvsegs_cols_full.
# List of columns to report for lvm fullreport's 'pvseg' subreport.
# See 'pvs --segments -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# pvsegs_cols_full = "pvseg_all,pv_uuid,lv_uuid"

# Configuration option report/segs_cols_full.
# List of columns to report for lvm fullreport's 'seg' subreport.
# See 'lvs --segments -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# segs_cols_full = "seg_all,lv_uuid"

# Configuration option report/vgs_sort_full.
# List of columns to sort by when reporting lvm fullreport's 'vgs'
subreport.
# See 'vgs -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# vgs_sort_full = "vg_name"

# Configuration option report/pvs_sort_full.
# List of columns to sort by when reporting lvm fullreport's 'vgs'
subreport.
# See 'pvs -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# pvs_sort_full = "pv_name"

# Configuration option report/lvs_sort_full.
# List of columns to sort by when reporting lvm fullreport's 'lvs'
subreport.
# See 'lvs -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# lvs_sort_full = "vg_name,lv_name"

# Configuration option report/pvsegs_sort_full.
# List of columns to sort by when reporting for lvm fullreport's 'pvseg'
subreport.
# See 'pvs --segments -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# pvsegs_sort_full = "pv_uuid,pvseg_start"

# Configuration option report/segs_sort_full.
# List of columns to sort by when reporting lvm fullreport's 'seg'
subreport.
# See 'lvs --segments -o help' for the list of possible fields.
# This configuration option has an automatic default value.
# segs_sort_full = "lv_uuid,seg_start"

# Configuration option report/mark_hidden_devices.
```

```

# Use brackets [] to mark hidden devices.
# This configuration option has an automatic default value.
# mark_hidden_devices = 1

# Configuration option report/two_word_unknown_device.
# Use the two words 'unknown device' in place of '[unknown]'.
# This is displayed when the device for a PV is not known.
# This configuration option has an automatic default value.
# two_word_unknown_device = 0
# }

# Configuration section dmeventd.
# Settings for the LVM event daemon.
dmeventd {

# Configuration option dmeventd/mirror_library.
# The library dmeventd uses when monitoring a mirror device.
# libdevmapper-event-lvm2mirror.so attempts to recover from
# failures. It removes failed devices from a volume group and
# reconfigures a mirror as necessary. If no mirror library is
# provided, mirrors are not monitored through dmeventd.
mirror_library = "libdevmapper-event-lvm2mirror.so"

# Configuration option dmeventd/raid_library.
# This configuration option has an automatic default value.
# raid_library = "libdevmapper-event-lvm2raid.so"

# Configuration option dmeventd/snapshot_library.
# The library dmeventd uses when monitoring a snapshot device.
# libdevmapper-event-lvm2snapshot.so monitors the filling of snapshots
# and emits a warning through syslog when the usage exceeds 80%. The
# warning is repeated when 85%, 90% and 95% of the snapshot is filled.
snapshot_library = "libdevmapper-event-lvm2snapshot.so"

# Configuration option dmeventd/thin_library.
# The library dmeventd uses when monitoring a thin device.
# libdevmapper-event-lvm2thin.so monitors the filling of a pool
# and emits a warning through syslog when the usage exceeds 80%. The
# warning is repeated when 85%, 90% and 95% of the pool is filled.
thin_library = "libdevmapper-event-lvm2thin.so"

# Configuration option dmeventd/executable.
# The full path to the dmeventd binary.
# This configuration option has an automatic default value.
# executable = "/usr/sbin/dmeventd"
}

# Configuration section tags.
# Host tag settings.
# This configuration section has an automatic default value.
# tags {

# Configuration option tags/hosttags.
# Create a host tag using the machine name.
# The machine name is nodename returned by uname(2).
# This configuration option has an automatic default value.

```

```
# hosttags = 0

# Configuration section tags/<tag>.
# Replace this subsection name with a custom tag name.
# Multiple subsections like this can be created. The '@' prefix for
# tags is optional. This subsection can contain host_list, which is a
# list of machine names. If the name of the local machine is found in
# host_list, then the name of this subsection is used as a tag and is
# applied to the local machine as a 'host tag'. If this subsection is
# empty (has no host_list), then the subsection name is always applied
# as a 'host tag'.
#
# Example
# The host tag foo is given to all hosts, and the host tag
# bar is given to the hosts named machine1 and machine2.
# tags { foo { } bar { host_list = [ "machine1", "machine2" ] } }
#
# This configuration section has variable name.
# This configuration section has an automatic default value.
# tag {

# Configuration option tags/<tag>/host_list.
# A list of machine names.
# These machine names are compared to the nodename returned
# by uname(2). If the local machine name matches an entry in
# this list, the name of the subsection is applied to the
# machine as a 'host tag'.
# This configuration option does not have a default value defined.
# }
# }
```

APPENDIX C. LVM SELECTION CRITERIA

As of Red Hat Enterprise Linux release 7.1, many LVM reporting commands accept the **-S** or **--select** option to define selection criteria for those commands. As of Red Hat Enterprise Linux release 7.2, many processing commands support selection criteria as well. These two categories of commands for which you can define selection criteria are defined as follows:

- Reporting commands — Display only the lines that satisfy the selection criteria. Examples of reporting commands for which you can define selection criteria include **pvs**, **vgs**, **lvs**, **pvdisplay**, **vgdisplay**, **lvdisplay**, **lvm devtypes**, and **dmsetup info -c**.

Specifying the **-o selected** option in addition to the **-S** option displays all rows and adds a "selected" column that shows 1 if the row matches the selection criteria and 0 if it does not.

- Processing commands — Process only the items that satisfy the selection criteria. Examples of processing commands for which you can define selection criteria include **pvchange**, **vgchange**, **lvchange**, **vgimport**, **vgexport**, **vgremove**, and **lvremove**.

Selection criteria are a set of statements that use comparison operators to define the valid values for particular fields to display or process. The selected fields are, in turn, combined by logical and grouping operators.

When specifying which fields to display using selection criteria, there is no requirement for the field which is in the selection criteria to be displayed. The selection criteria can contain one set of fields while the output can contain a different set of fields.

- For a listing of available fields for the various LVM components, see [Section C.3, “Selection Criteria Fields”](#).
- For a listing of allowed operators, see [Section C.2, “Selection Criteria Operators”](#). The operators are also provided on the `lvm(8)` man page.
- You can also see full sets of fields and possible operators by specifying the **help** (or **?**) keyword for the **-S/--select** option of a reporting commands. For example, the following command displays the fields and possible operators for the **lvs** command.

```
# lvs -S help
```

For the Red Hat Enterprise Linux 7.2 release, you can specify time values as selection criteria for fields with a field type of **time**. For information on specifying time values, see [Section C.4, “Specifying Time Values”](#).

C.1. SELECTION CRITERIA FIELD TYPES

The fields you specify for selection criteria are of a particular type. The help output for each field display the field type enclosed in brackets. The following help output examples show the output indicating the field types **string**, **string_list**, **number**, **percent**, **size** and **time**.

```
lv_name          - Name. LVs created for internal use are enclosed in
brackets.[string]
lv_role          - LV role. [string list]
raid_mismatch_count - For RAID, number of mismatches found or repaired.
[number]
copy_percent     - For RAID, mirrors and pvmove, current percentage in-
```

```

sync. [percent]
lv_size          - Size of LV in current units. [size]
lv_time          - Creation time of the LV, if known [time]

```

Table C.1, “Selection Criteria Field Types” describes the selection criteria field types

Table C.1. Selection Criteria Field Types

Field Type	Description
number	Non-negative integer value.
size	Floating point value with units, 'm' unit used by default if not specified.
percent	Non-negative integer with or without % suffix.
string	Characters quoted by ' or " or unquoted.
string list	Strings enclosed by [] or { } and elements delimited by either "all items must match" or "at least one item must match" operator.

The values you specify for a field can be the following:

- Concrete values of the field type
- Regular expressions that include any fields of the **string** field type, such as "+~" operator.
- Reserved values; for example -1, unknown, undefined, undef are all keywords to denote an undefined numeric value.
- Defined synonyms for the field values, which can be used in selection criteria for values just as for their original values. For a listing of defined synonyms for field values, see [Table C.14](#), “Selection Criteria Synonyms”.

C.2. SELECTION CRITERIA OPERATORS

Table C.2, “Selection Criteria Grouping Operators” describes the selection criteria grouping operators.

Table C.2. Selection Criteria Grouping Operators

Grouping Operator	Description
()	Used for grouping statements
[]	Used to group strings into a string list (exact match)
{ }	Used to group strings into a string list (subset match)

Table C.3, “Selection Criteria Comparison Operators” describes the selection criteria comparison operators and the field types with which they can be used.

Table C.3. Selection Criteria Comparison Operators

Comparison Operator	Description	Field Type
=~	Matching regular expression	regex
!~	Not matching regular expression.	regex
=	Equal to	number, size, percent, string, string list, time
!=	Not equal to	number, size, percent, string, string list, time
>=	Greater than or equal to	number, size, percent, time
>	Greater than	number, size, percent, time
<=	Less than or equal to	number, size, percent, time
<	Less than	number, size, percent, time
since	Since specified time (same as >=)	time
after	After specified time (same as >)	time
until	Until specified time (same as <=)	time
before	Before specified time (same as <)	time

Table C.4, “[Selection Criteria Logical and Grouping Operators](#)” describes the selection criteria logical and grouping operators.

Table C.4. Selection Criteria Logical and Grouping Operators

Logical and Grouping Operator	Description
&&	All fields must match
,	All fields must match (same as &&)
	At least one field must match
#	At least one field must match (same as)
!	Logical negation

Logical and Grouping Operator	Description
(	Left parenthesis (grouping operator)
)	Right parenthesis (grouping operator)
[	List start (grouping operator)
]	List end (grouping operator)
{	List subset start (grouping operator)
}	List subset end (grouping operator)

C.3. SELECTION CRITERIA FIELDS

This section describes the logical and physical volume selection criteria fields you can specify.

[Table C.5, “Logical Volume Fields”](#) describes the logical volume fields and their field types.

Table C.5. Logical Volume Fields

Logical Volume Field	Description	Field Type
lv_uuid	Unique identifier	string
lv_name	Name (logical volumes created for internal use are enclosed in brackets)	string
lv_full_name	Full name of logical volume including its volume group, namely VG/LV	string
lv_path	Full pathname for logical volume (blank for internal logical volumes)	string
lv_dm_path	Internal device mapper pathname for logical volume (in /dev/mapper directory)	string
lv_parent	For logical volumes that are components of another logical volume, the parent logical volume	string
lv_layout	logical volume layout	string list
lv_role	logical volume role	string list
lv_initial_image_sync	Set if mirror/RAID images underwent initial resynchronization	number

Logical Volume Field	Description	Field Type
lv_image_synced	Set if mirror/RAID image is synchronized	number
lv_merging	Set if snapshot logical volume is being merged to origin	number
lv_converting	Set if logical volume is being converted	number
lv_allocation_policy	logical volume allocation policy	string
lv_allocation_locked	Set if logical volume is locked against allocation changes	number
lv_fixed_minor	Set if logical volume has fixed minor number assigned	number
lv_merge_failed	Set if snapshot merge failed	number
lv_snapshot_invalid	Set if snapshot logical volume is invalid	number
lv_skip_activation	Set if logical volume is skipped on activation	number
lv_when_full	For thin pools, behavior when full	string
lv_active	Active state of the logical volume	string
lv_active_locally	Set if the logical volume is active locally	number
lv_active_remotely	Set if the logical volume is active remotely	number
lv_active_exclusively	Set if the logical volume is active exclusively	number
lv_major	Persistent major number or -1 if not persistent	number
lv_minor	Persistent minor number or -1 if not persistent	number
lv_read_ahead	Read ahead setting in current units	size
lv_size	Size of logical volume in current units	size

Logical Volume Field	Description	Field Type
lv_metadata_size	For thin and cache pools, the size of the logical volume that holds the metadata	size
seg_count	Number of segments in logical volume	number
origin	For snapshots, the origin device of this logical volume	string
origin_size	For snapshots, the size of the origin device of this logical volume	size
data_percent	For snapshot and thin pools and volumes, the percentage full if logical volume is active	percent
snap_percent	For snapshots, the percentage full if logical volume is active	percent
metadata_percent	For thin pools, the percentage of metadata full if logical volume is active	percent
copy_percent	For RAID, mirrors and pvmove, current percentage in-sync	percent
sync_percent	For RAID, mirrors and pvmove, current percentage in-sync	percent
raid_mismatch_count	For RAID, number of mismatches found or repaired	number
raid_sync_action	For RAID, the current synchronization action being performed	string
raid_write_behind	For RAID1, the number of outstanding writes allowed to writemostly devices	number
raid_min_recovery_rate	For RAID1, the minimum recovery I/O load in kiB/sec/disk	number
raid_max_recovery_rate	For RAID1, the maximum recovery I/O load in kiB/sec/disk	number
move_pv	For pvmove, source physical volume of temporary logical volume created by pvmove	string
convert_lv	For lvconvert, name of temporary logical volume created by lvconvert	string

Logical Volume Field	Description	Field Type
mirror_log	For mirrors, the logical volume holding the synchronization log	string
data_lv	For thin and cache pools, the logical volume holding the associated data	string
metadata_lv	For thin and cache pools, the logical volume holding the associated metadata	string
pool_lv	For thin volumes, the thin pool logical volume for this volume	string
lv_tags	Tags, if any	string list
lv_profile	Configuration profile attached to this logical volume	string
lv_time	Creation time of the logical volume, if known	time
lv_host	Creation host of the logical volume, if known	string
lv_modules	Kernel device-mapper modules required for this logical volume	string list

Table C.6, “Logical Volume Device Combined Info and Status Fields” describes the logical volume device fields that combine both logical device info and logical device status.

Table C.6. Logical Volume Device Combined Info and Status Fields

Logical Volume Field	Description	Field Type
lv_attr	Selects according to both logical volume device info as well as logical volume status.	string

Table C.7, “Logical Volume Device Info Fields” describes the logical volume device info fields and their field types.

Table C.7. Logical Volume Device Info Fields

Logical Volume Field	Description	Field Type
lv_kernel_major	Currently assigned major number or -1 if logical volume is not active	number
lv_kernel_minor	Currently assigned minor number or -1 if logical volume is not active	number

Logical Volume Field	Description	Field Type
lv_kernel_read_ahead	Currently-in-use read ahead setting in current units	size
lv_permissions	logical volume permissions	string
lv_suspended	Set if logical volume is suspended	number
lv_live_table	Set if logical volume has live table present	number
lv_inactive_table	Set if logical volume has inactive table present	number
lv_device_open	Set if logical volume device is open	number

Table C.8, “Logical Volume Device Status Fields” describes the logical volume device status fields and their field types.

Table C.8. Logical Volume Device Status Fields

Logical Volume Field	Description	Field Type
cache_total_blocks	Total cache blocks	number
cache_used_blocks	Used cache blocks	number
cache_dirty_blocks	Dirty cache blocks	number
cache_read_hits	Cache read hits	number
cache_read_misses	Cache read misses	number
cache_write_hits	Cache write hits	number
cache_write_misses	Cache write misses	number
lv_health_status	logical volume health status	string

Table C.9, “Physical Volume Label Fields” describes the physical volume label fields and their field types.

Table C.9. Physical Volume Label Fields

Physical Volume Field	Description	Field Type
pv_fmt	Type of metadata	string
pv_uuid	Unique identifier	string
dev_size	Size of underlying device in current units	size
pv_name	Name	string
pv_mda_free	Free metadata area space on this device in current units	size
pv_mda_size	Size of smallest metadata area on this device in current units	size

Table C.5, “Logical Volume Fields” describes the physical volume fields and their field types.

Table C.10. Physical Volume Fields

Physical Volume Field	Description	Field Type
pe_start	Offset to the start of data on the underlying device	number
pv_size	Size of physical volume in current units	size
pv_free	Total amount of unallocated space in current units	size
pv_used	Total amount of allocated space in current units	size
pv_attr	Various attributes	string
pv_allocatable	Set if this device can be used for allocation	number
pv_exported	Set if this device is exported	number
pv_missing	Set if this device is missing in system	number
pv_pe_count	Total number of physical extents	number
pv_pe_alloc_count	Total number of allocated physical extents	number
pv_tags	Tags, if any	string list
pv_mda_count	Number of metadata areas on this device	number

Physical Volume Field	Description	Field Type
pv_mda_used_count	Number of metadata areas in use on this device	number
pv_ba_start	Offset to the start of PV Bootloader Area on the underlying device in current units	size
pv_ba_size	Size of PV Bootloader Area in current units	size

Table C.11, “Volume Group Fields” describes the volume group fields and their field types.

Table C.11. Volume Group Fields

Volume Group Field	Description	Field Type
vg_fmt	Type of metadata	string
vg_uuid	Unique identifier	string
vg_name	Name	string
vg_attr	Various attributes	string
vg_permissions	Volume group permissions	string
vg_extendable	Set if volume group is extendable	number
vg_exported	Set if volume group is exported	number
vg_partial	Set if volume group is partial	number
vg_allocation_policy	Volume group allocation policy	string
vg_clustered	Set if volume group is clustered	number
vg_size	Total size of volume group in current units	size
vg_free	Total amount of free space in current units	size
vg_sysid	System ID of the volume group indicating which host owns it	string
vg_systemid	System ID of the volume group indicating which host owns it	string

Volume Group Field	Description	Field Type
vg_extent_size	Size of physical extents in current units	size
vg_extent_count	Total number of physical extents	number
vg_free_count	Total number of unallocated physical extents	number
max_lv	Maximum number of logical volumes allowed in volume group or 0 if unlimited	number
max_pv	Maximum number of physical volumes allowed in volume group or 0 if unlimited	number
pv_count	Number of physical volumes	number
lv_count	Number of logical volumes	number
snap_count	Number of snapshots	number
vg_seqno	Revision number of internal metadata — incremented whenever it changes	number
vg_tags	Tags, if any	string list
vg_profile	Configuration profile attached to this volume group	string
vg_mda_count	Number of metadata areas on this volume group	number
vg_mda_used_count	Number of metadata areas in use on this volume group	number
vg_mda_free	Free metadata area space for this volume group in current units	size
vg_mda_size	Size of smallest metadata area for this volume group in current units	size
vg_mda_copies	Target number of in use metadata areas in the volume group	number

Table C.12, “Logical Volume Segment Fields” describes the logical volume segment fields and their field types.

Table C.12. Logical Volume Segment Fields

Logical Volume Segment Field	Description	Field Type
segtype	Type of logical volume segment	string
stripes	Number of stripes or mirror legs	number
stripesize	For stripes, amount of data placed on one device before switching to the next	size
stripe_size	For stripes, amount of data placed on one device before switching to the next	size
regionsize	For mirrors, the unit of data copied when synchronizing devices	size
region_size	For mirrors, the unit of data copied when synchronizing devices	size
chunksize	For snapshots, the unit of data used when tracking changes	size
chunk_size	For snapshots, the unit of data used when tracking changes	size
thin_count	For thin pools, the number of thin volumes in this pool	number
discards	For thin pools, how discards are handled	string
cachemode	For cache pools, how writes are cached	string
zero	For thin pools, if zeroing is enabled	number
transaction_id	For thin pools, the transaction id	number
thin_id	For thin volumes, the thin device id	number
seg_start	Offset within the logical volume to the start of the segment in current units	size
seg_start_pe	Offset within the logical volume to the start of the segment in physical extents.	number
seg_size	Size of segment in current units	size
seg_size_pe	Size of segment in physical extents	size

Logical Volume Segment Field	Description	Field Type
seg_tags	Tags, if any	string list
seg_pe_ranges	Ranges of physical extents of underlying devices in command line format	string
devices	Underlying devices used with starting extent numbers	string
seg_monitor	dmeventd monitoring status of the segment	string
cache_policy	The cache policy (cached segments only)	string
cache_settings	Cache settings/parameters (cached segments only)	string list

Table C.13, “Physical Volume Segment Fields” describes the physical volume segment fields and their field types.

Table C.13. Physical Volume Segment Fields

Physical Volume Segment Field	Description	Field Type
pvseg_start	Physical extent number of start of segment	number
pvseg_size	Number of extents in segment	number

Table C.14, “Selection Criteria Synonyms” lists the synonyms you can use for field values. These synonyms can be used in selection criteria as well as for values just like their original values. In this table, a field value of "" indicates a blank string, which can be matched by specifying **-S 'field_name=""'**.

In this table, a field indicated by 0 or 1 indicates a binary value. You can specify a **--binary** option for reporting tools which causes binary fields to display 0 or 1 instead of what is indicated in this table as "some text" or "".

Table C.14. Selection Criteria Synonyms

Field	Field Value	Synonyms
pv_allocatable	allocatable	1
pv_allocatable	""	0
pv_exported	exported	1
pv_exported	""	0

Field	Field Value	Synonyms
pv_missing	missing	1
pv_missing	""	0
vg_extendable	extendable	1
vg_extendable	""	0
vg_exported	exported	1
vg_exported	""	0
vg_partial	partial	1
vg_partial	""	0
vg_clustered	clustered	1
vg_clustered	""	0
vg_permissions	writable	rw, read-write
vg_permissions	read-only	r, ro
vg_mda_copies	unmanaged	unknown, undefined, undef, -1
lv_initial_image_sync	initial image sync	sync, 1
lv_initial_image_sync	""	0
lv_image_synced	image synced	synced, 1
lv_image_synce	""	0
lv_merging	merging	1
lv_merging	""	0
lv_converting	converting	1
lv_converting	""	0
lv_allocation_locked	allocation locked	locked, 1
lv_allocation_locked	""	0

Field	Field Value	Synonyms
lv_fixed_minor	fixed minor	fixed, 1
lv_fixed_minor	""	0
lv_active_locally	active locally	active, locally, 1
lv_active_locally	""	0
lv_active_remotely	active remotely	active, remotely, 1
lv_active_remotely	""	0
lv_active_exclusively	active exclusively	active, exclusively, 1
lv_active_exclusively	""	0
lv_merge_failed	merge failed	failed, 1
lv_merge_failed	""	0
lv_snapshot_invalid	snapshot invalid	invalid, 1
lv_snapshot_invalid	""	0
lv_suspended	suspended	1
lv_suspended	""	0
lv_live_table	live table present	live table, live, 1
lv_live_table	""	0
lv_inactive_table	inactive table present	inactive table, inactive, 1
lv_inactive_table	""	0
lv_device_open	open	1
lv_device_open	""	0
lv_skip_activation	skip activation	skip, 1
lv_skip_activation	""	0
zero	zero	1

Field	Field Value	Synonyms
zero	""	0
lv_permissions	writable	rw, read-write
lv_permissions	read-only	r, ro
lv_permissions	read-only-override	ro-override, r-override, R
lv_when_full	error	error when full, error if no space
lv_when_full	queue	queue when full, queue if no space
lv_when_full	""	undefined
cache_policy	""	undefined
seg_monitor	""	undefined
lv_health_status	""	undefined

C.4. SPECIFYING TIME VALUES

When specifying time values for LVM selection, you can use either a standardized time specification format or a more free-form specification, as described in [Section C.4.1, “Standard time selection format”](#) and [Section C.4.2, “Freeform time selection format”](#).

You can specify the way time values are displayed with the report/time format configuration option in the `/etc/lvm/lvm.conf` configuration file. Information on specifying this option is provided in the `lvm.conf` file.

When specifying time values, you can use the comparison operator aliases **since**, **after**, **until**, and **before**, as described in [Table C.3, “Selection Criteria Comparison Operators”](#).

C.4.1. Standard time selection format

You can specify time values for LVM selection in the following format.

```
date time timezone
```

[Table C.15, “Time Specification Formats”](#) summarizes the formats you can use when specifying these time values.

Table C.15. Time Specification Formats

Field	Field Value
date	YYYY-MM-DD
	YYYY-MM, auto DD=1
	YYYY, auto MM=01 and DD=01
time	hh:mm:ss
	hh:mm, auto ss=0
	hh, auto mm=0, auto ss=0
timezone (always with + or - sign)	+hh:mm or -hh:mm
	+hh or -hh

The full date/time specification is YYYY-MM-DD hh:mm:ss. Users are able to leave date/time parts from right to left. Whenever these parts are left out, a range is assumed automatically with second granularity. For example:

- "2015-07-07 9:51" means range of "2015-07-07 9:51:00" - "2015-07-07 9:51:59"
- "2015-07" means range of "2015-07-01 0:00:00" - "2015-07-31 23:59:59"
- "2015" means range of "2015-01-01 0:00:00" - "2015-12-31 23:59:59"

The following examples show the date/time specification as used in selection criteria.

```
lvs -S 'time since "2015-07-07 9:51"'
lvs -S 'time = "2015-07"'
lvs -S 'time = "2015"'
```

C.4.2. Freeform time selection format

You can specify the date/time specification in LVM selection criteria using the following entitles.

- weekday names ("Sunday" - "Saturday" or abbreviated as "Sun" - "Sat")
- labels for points in time ("noon", "midnight")
- labels for a day relative to current day ("today", "yesterday")
- points back in time with relative offset from today (N is a number)
- ("N" "seconds"/"minutes"/"hours"/"days"/"weeks"/"years" "ago")

- ("N" "secs"/"mins"/"hrs" ... "ago")
- ("N" "s"/"m"/"h" ... "ago")
- time specification either in hh:mm:ss format or with AM/PM suffixes
- month names ("January" - "December" or abbreviated as "Jan" - "Dec")

The following examples show the **freeform** date/time specification as used in selection criteria.

```
lvs -S 'time since "yesterday 9AM"'
lvs -S 'time since "Feb 3 years 2 months ago"'
lvs -S 'time = "February 2015"'
lvs -S 'time since "Jan 15 2015" && time until yesterday'
lvs -S 'time since "today 6AM"'
```

C.5. SELECTION CRITERIA DISPLAY EXAMPLES

This section provides a series of examples showing how to use selection criteria for LVM display commands. The examples in this section use a system configured with LVM volumes that yield the following output when selection criteria are not used.

```
# lvs -a -o+layout,role
LV          VG Attr      LSize   Pool Origin Data%  Meta%  Layout
Role
root        f1 -wi-ao----   9.01g
public
swap        f1 -wi-ao---- 512.00m
public
[lvol10_pmspare] vg ewi-----   4.00m
private,    \

pool, spare
lvol1       vg Vwi-a-tz--   1.00g pool      0.00
thin, sparse public
lvol2       vg Vwi-a-tz--   1.00g pool      0.00
thin, sparse public,    \

origin,     \

thinorigin
lvol3       vg Vwi---tz-k   1.00g pool lvol2
thin, sparse public,    \

snapshot,   \

thinsnapshot
pool        vg twi-aotz-- 100.00m      0.00   1.07
thin, pool  private
[pool_tdata] vg Twi-ao---- 100.00m
private,    \

thin, pool, \

data
```

```

    [pool_tmeta]    vg ewi-ao---- 4.00m                linear
private, \

thin,pool, \

metadata

```

The following command displays all logical volumes with "lv[13]" in their name, using a regular expression to specify this.

```

# lvs -a -o+layout,role -S 'lv_name=~lv[13]'
LV      VG      Attr          LSize   Pool Origin Data%   Layout          Role
lv[1]   vg      Vwi-a-tz--    1.00g   pool          0.00   thin,sparse    public
lv[3]   vg      Vwi---tz-k    1.00g   pool lv[2]          thin,sparse
public,snapshot,thinsnapshot

```

The following command displays all logical volumes greater than 500 megabytes in size.

```

# lvs -a -o+layout,role -S 'lv_size>500m'
LV      VG      Attr          LSize   Pool Origin Data%   Layout          Role
root    f1      -wi-ao----    9.01g                linear          public
swap     f1      -wi-ao----   512.00m                linear          public
lv[1]   vg      Vwi-a-tz--    1.00g   pool          0.00   thin,sparse    public
lv[2]   vg      Vwi-a-tz--    1.00g   pool          0.00   thin,sparse
public,origin,thinorigin
lv[3]   vg      Vwi---tz-k    1.00g   pool lv[2]          thin,sparse
public,snapshot, \

thinsnapshot

```

The following command displays all logical volumes that include **thin** as a logical volume role, indicating that the logical volume is used in constructing a thin pool. This example uses braces ({}) to indicate a subset in the display.

```

# lvs -a -o+layout,role -S 'lv_role={thin}'
LV          VG      Attr          LSize   Layout          Role
[pool_tdata] vg      Twi-ao----   100.00m linear          private,thin,pool,data
[pool_tmeta] vg      ewi-ao----    4.00m linear
private,thin,pool,metadata

```

The following command displays all usable top-level logical volumes, which are the logical volumes with a role of "public". If you do not specify braces ({}) in a string list to indicate a subset, it is assumed by default; specifying **lv_role=public** is equivalent to specifying **lv_role={public}**.

```

# lvs -a -o+layout,role -S 'lv_role=public'
LV      VG      Attr          LSize   Pool Origin Data%   Layout          Role
root    f1      -wi-ao----    9.01g                linear          public
swap     f1      -wi-ao----   512.00m                linear          public
lv[1]   vg      Vwi-a-tz--    1.00g   pool          0.00   thin,sparse    public
lv[2]   vg      Vwi-a-tz--    1.00g   pool          0.00   thin,sparse
public,origin,thinorigin
lv[3]   vg      Vwi---tz-k    1.00g   pool lv[2]          thin,sparse
public,snapshot,thinsnapshot

```

The following command displays all logical volumes with a thin layout.

```
# lvs -a -o+layout,role -S 'lv_layout={thin}'
LV      VG Attr          LSize   Pool Origin Data% Meta% Layout      Role
lvol1   vg Vwi-a-tz-- 1.00g   pool          0.00         thin,sparse public
lvol2   vg Vwi-a-tz-- 1.00g   pool          0.00         thin,sparse
public,origin, \
thinorigin
lvol3   vg Vwi---tz-k 1.00g   pool lvol2          thin,sparse
public,snapshot, \
thinsnapshot
pool    vg twi-aotz-- 100.00m          0.00 1.07 thin,pool private
```

The following command displays all logical volumes with a layout field that matches "sparse,thin" exactly. Note that it is not necessary to specify the string list members for the match to be positive.

```
# lvs -a -o+layout,role -S 'lv_layout=[sparse,thin]'
LV      VG Attr          LSize   Pool Origin Data% Layout      Role
lvol1   vg Vwi-a-tz-- 1.00g   pool          0.00 thin,sparse public
lvol2   vg Vwi-a-tz-- 1.00g   pool          0.00 thin,sparse
public,origin,thinorigin
lvol3   vg Vwi---tz-k 1.00g   pool lvol2          thin,sparse
public,snapshot,thinsnapshot
```

The following command displays the logical volume names of the logical volumes that are thin, sparse logical volumes. Note that the list of fields used for selection criteria do not need to be the same as the list of fields to display.

```
# lvs -a -o lv_name -S 'lv_layout=[sparse,thin]'
LV
lvol1
lvol2
lvol3
```

C.6. SELECTION CRITERIA PROCESSING EXAMPLES

This section provides a series of examples showing how to use selection criteria in commands that process LVM logical volumes.

This example shows the initial configuration of a group of logical volumes, including thin snapshots. Thin snapshots have the "skip activation" flag set by default. This example also includes the logical volume **lvol14** which also has the "skip activation" flag set.

```
# lvs -o name,skip_activation,layout,role
LV      SkipAct          Layout      Role
root                    linear      public
swap                    linear      public
lvol1                    thin,sparse public
lvol2                    thin,sparse public,origin,thinorigin
lvol3 skip activation thin,sparse public,snapshot,thinsnapshot
lvol4 skip activation linear      public
pool                    thin,pool   private
```

The following command removes the skip activation flag from all logical volumes that are thin snapshots.

```
# lvchange --setactivationskip n -S 'role=thinsnapshot'
Logical volume "lvol3" changed.
```

The following command shows the configuration of the logical volumes after executing the **lvchange** command. Note that the "skip activation" flag has not been unset from the logical volume that is not a thin snapshot.

```
# lvs -o name,active,skip_activation,layout,role
LV      Active SkipAct      Layout      Role
root    active              linear      public
swap    active              linear      public
lvol1    active              thin,sparse public
lvol2    active              thin,sparse public,origin,thinorigin
lvol3    active              thin,sparse public,snapshot,thinsnapshot
lvol4    active skip activation linear      public
pool     active              thin,pool   private
```

The following command shows the configuration of the logical volumes after an additional thin origin/snapshot volume has been created.

```
# lvs -o name,active,skip_activation,origin,layout,role
LV      Active SkipAct      Origin Layout      Role
root    active              linear      public
swap    active              linear      public
lvol1    active              thin,sparse public
lvol2    active              thin,sparse
public,origin,thinorigin
lvol3    active              lvol2 thin,sparse
public,snapshot,thinsnapshot
lvol4    active skip activation linear      public
lvol5    active              thin,sparse
public,origin,thinorigin
lvol6    active              lvol5 thin,sparse
public,snapshot,thinsnapshot
pool     active              thin,pool   private
```

The following command activates logical volumes that are both thin snapshot volumes and have an origin volume of **lvol2**.

```
# lvchange -ay -S 'lv_role=thinsnapshot && origin=lvol2'

# lvs -o name,active,skip_activation,origin,layout,role
LV      Active SkipAct      Origin Layout      Role
root    active              linear      public
swap    active              linear      public
lvol1    active              thin,sparse public
lvol2    active              thin,sparse
public,origin,thinorigin
lvol3    active              lvol2 thin,sparse
public,snapshot,thinsnapshot
lvol4    active skip activation linear      public
lvol5    active              thin,sparse
public,origin,thinorigin
```

```
lvvol6                                lvvol5  thin, sparse
public, snapshot, thinsnapshot
pool  active                          thin, pool  private
```

If you execute a command on a whole item while specifying selection criteria that match an item from that whole, the entire whole item is processed. For example, if you change a volume group while selecting one or more items from that volume group, the whole volume group is selected. This example selects logical volume **lvvol1**, which is part of volume group **vg**. All of the logical volumes in volume group **vg** are processed.

```
# lvs -o name,vg_name
LV      VG
root    fedora
swap    fedora
lvvol1  vg
lvvol2  vg
lvvol3  vg
lvvol4  vg
lvvol5  vg
lvvol6  vg
pool    vg

# vgchange -ay -S 'lv_name=lvvol1'
7 logical volume(s) in volume group "vg" now active
```

The following example shows a more complex selection criteria statement. In this example, all logical volumes are tagged with **mytag** if they have a role of origin and are also named lvvol[456] or the logical volume size is more than 5 gigabytes.

```
# lvchange --addtag mytag -S '(role=origin && lv_name=~lvvol[456]) ||
lv_size > 5g'
Logical volume "root" changed.
Logical volume "lvvol5" changed.
```

APPENDIX D. LVM OBJECT TAGS

An LVM tag is a word that can be used to group LVM2 objects of the same type together. Tags can be attached to objects such as physical volumes, volume groups, and logical volumes. Tags can be attached to hosts in a cluster configuration.

Tags can be given on the command line in place of PV, VG or LV arguments. Tags should be prefixed with **@** to avoid ambiguity. Each tag is expanded by replacing it with all objects possessing that tag which are of the type expected by its position on the command line.

LVM tags are strings of up to 1024 characters. LVM tags cannot start with a hyphen.

A valid tag can consist of a limited range of characters only. The allowed characters are [A-Za-z0-9_+.-]. As of the Red Hat Enterprise Linux 6.1 release, the list of allowed characters was extended, and tags can contain the /, =, !, :, #, and & characters.

Only objects in a volume group can be tagged. Physical volumes lose their tags if they are removed from a volume group; this is because tags are stored as part of the volume group metadata and that is deleted when a physical volume is removed.

The following command lists all the logical volumes with the **database** tag.

```
# lvs @database
```

The following command lists the currently active host tags.

```
# lvm tags
```

D.1. ADDING AND REMOVING OBJECT TAGS

To add or delete tags from physical volumes, use the **--addtag** or **--deltag** option of the **pvchange** command.

To add or delete tags from volume groups, use the **--addtag** or **--deltag** option of the **vgchange** or **vgcreate** commands.

To add or delete tags from logical volumes, use the **--addtag** or **--deltag** option of the **lvchange** or **lvcreate** commands.

You can specify multiple **--addtag** and **--deltag** arguments within a single **pvchange**, **vgchange**, or **lvchange** command. For example, the following command deletes the tags **T9** and **T10** and adds the tags **T13** and **T14** to the volume group **grant**.

```
# vgchange --deltag T9 --deltag T10 --addtag T13 --addtag T14 grant
```

D.2. HOST TAGS

In a cluster configuration, you can define host tags in the configuration files. If you set **hosttags = 1** in the **tags** section, a host tag is automatically defined using the machine's host name. This allows you to use a common configuration file which can be replicated on all your machines so they hold identical copies of the file, but the behavior can differ between machines according to the host name.

For information on the configuration files, see [Appendix B, The LVM Configuration Files](#).

For each host tag, an extra configuration file is read if it exists: `lvm_hosttag.conf`. If that file defines new tags, then further configuration files will be appended to the list of files to read in.

For example, the following entry in the configuration file always defines **tag1**, and defines **tag2** if the host name is **host1**.

```
tags { tag1 { } tag2 { host_list = ["host1"] } }
```

D.3. CONTROLLING ACTIVATION WITH TAGS

You can specify in the configuration file that only certain logical volumes should be activated on that host. For example, the following entry acts as a filter for activation requests (such as **vgchange -ay**) and only activates **vg1/lvol0** and any logical volumes or volume groups with the **database** tag in the metadata on that host.

```
activation { volume_list = ["vg1/lvol0", "@database" ] }
```

There is a special match `"@"` that causes a match only if any metadata tag matches any host tag on that machine.

As another example, consider a situation where every machine in the cluster has the following entry in the configuration file:

```
tags { hosttags = 1 }
```

If you want to activate **vg1/lvol2** only on host **db2**, do the following:

1. Run **lvchange --addtag @db2 vg1/lvol2** from any host in the cluster.
2. Run **lvchange -ay vg1/lvol2**.

This solution involves storing host names inside the volume group metadata.

APPENDIX E. LVM VOLUME GROUP METADATA

The configuration details of a volume group are referred to as the metadata. By default, an identical copy of the metadata is maintained in every metadata area in every physical volume within the volume group. LVM volume group metadata is stored as ASCII.

If a volume group contains many physical volumes, having many redundant copies of the metadata is inefficient. It is possible to create a physical volume without any metadata copies by using the `--metadatacopies 0` option of the `pvcreate` command. Once you have selected the number of metadata copies the physical volume will contain, you cannot change that at a later point. Selecting 0 copies can result in faster updates on configuration changes. Note, however, that at all times every volume group must contain at least one physical volume with a metadata area (unless you are using the advanced configuration settings that allow you to store volume group metadata in a file system). If you intend to split the volume group in the future, every volume group needs at least one metadata copy.

The core metadata is stored in ASCII. A metadata area is a circular buffer. New metadata is appended to the old metadata and then the pointer to the start of it is updated.

You can specify the size of metadata area with the `--metadatasize` option of the `pvcreate` command. The default size may be too small for volume groups that contain physical volumes and logical volumes that number in the hundreds.

E.1. THE PHYSICAL VOLUME LABEL

By default, the `pvcreate` command places the physical volume label in the 2nd 512-byte sector. This label can optionally be placed in any of the first four sectors, since the LVM tools that scan for a physical volume label check the first 4 sectors. The physical volume label begins with the string **LABELONE**.

The physical volume label contains:

- Physical volume UUID
- Size of block device in bytes
- NULL-terminated list of data area locations
- NULL-terminated lists of metadata area locations

Metadata locations are stored as offset and size (in bytes). There is room in the label for about 15 locations, but the LVM tools currently use 3: a single data area plus up to two metadata areas.

E.2. METADATA CONTENTS

The volume group metadata contains:

- Information about how and when it was created
- Information about the volume group

The volume group information contains:

- Name and unique id
- A version number which is incremented whenever the metadata gets updated

- Any properties, such as: read/write or resizable
- Any administrative limit on the number of physical volumes and logical volumes it may contain
- The extent size (in units of sectors which are defined as 512 bytes)
- An unordered list of physical volumes making up the volume group, each with:
 - Its UUID, used to determine the block device containing it
 - Any properties, such as whether the physical volume is allocatable
 - The offset to the start of the first extent within the physical volume (in sectors)
 - The number of extents
- An unordered list of logical volumes, each consisting of
 - An ordered list of logical volume segments. For each segment the metadata includes a mapping applied to an ordered list of physical volume segments or logical volume segments

E.3. SAMPLE METADATA

The following shows an example of LVM volume group metadata for a volume group called **myvg**.

```
# Generated by LVM2: Tue Jan 30 16:28:15 2007

contents = "Text Format Volume Group"
version = 1

description = "Created *before* executing 'lvextend -L+5G /dev/myvg/mylv
/dev/sdc'"

creation_host = "tng3-1"           # Linux tng3-1 2.6.18-8.el5 #1 SMP Fri Jan
26 14:15:21 EST 2007 i686
creation_time = 1170196095         # Tue Jan 30 16:28:15 2007

myvg {
    id = "0zd3UT-wbYT-lDHq-lMPs-EjoE-0o18-wL28X4"
    seqno = 3
    status = ["RESIZEABLE", "READ", "WRITE"]
    extent_size = 8192              # 4 Megabytes
    max_lv = 0
    max_pv = 0

    physical_volumes {

        pv0 {
            id = "ZBW5qW-dXF2-0bGw-ZCad-2RlV-phwu-1c1RFt"
            device = "/dev/sda"      # Hint only

            status = ["ALLOCATABLE"]
            dev_size = 35964301      # 17.1491 Gigabytes
            pe_start = 384
            pe_count = 4390 # 17.1484 Gigabytes
        }
    }
}
```

```

pv1 {
    id = "ZHEZJW-MR64-D3QM-Rv7V-Hxsa-zU24-wztY19"
    device = "/dev/sdb"      # Hint only

    status = ["ALLOCATABLE"]
    dev_size = 35964301      # 17.1491 Gigabytes
    pe_start = 384
    pe_count = 4390 # 17.1484 Gigabytes
}

pv2 {
    id = "wCoG4p-55Ui-9tbp-VTEA-j06s-RAVx-UREW0G"
    device = "/dev/sdc"      # Hint only

    status = ["ALLOCATABLE"]
    dev_size = 35964301      # 17.1491 Gigabytes
    pe_start = 384
    pe_count = 4390 # 17.1484 Gigabytes
}

pv3 {
    id = "hGlUwi-zsBg-39FF-do88-pHxY-8XA2-9WKIiA"
    device = "/dev/sdd"      # Hint only

    status = ["ALLOCATABLE"]
    dev_size = 35964301      # 17.1491 Gigabytes
    pe_start = 384
    pe_count = 4390 # 17.1484 Gigabytes
}
}
logical_volumes {

    mylv {
        id = "GhUYSF-qVM3-rzQo-a6D2-o0aV-LQet-Ur90F9"
        status = ["READ", "WRITE", "VISIBLE"]
        segment_count = 2

        segment1 {
            start_extent = 0
            extent_count = 1280      # 5 Gigabytes

            type = "striped"
            stripe_count = 1          # linear

            stripes = [
                "pv0", 0
            ]
        }
        segment2 {
            start_extent = 1280
            extent_count = 1280      # 5 Gigabytes

            type = "striped"
            stripe_count = 1          # linear
        }
    }
}

```

```

    stripes = [
        "pv1", 0
    ]
}
}
}
```

APPENDIX F. REVISION HISTORY

Revision 3.0-2 Preparing document for 7.6 GA publication.	Thu Oct 4 2018	Steven Levine
Revision 2.0-2 Preparing document for 7.5 GA publication.	Thu Mar 15 2018	Steven Levine
Revision 2.0-1 Preparing document for 7.5 Beta publication.	Thu Dec 14 2017	Steven Levine
Revision 1.0-11 Document version for 7.4 GA publication.	Wed Jul 19 2017	Steven Levine
Revision 1.0-9 Preparing document for 7.4 Beta publication.	Mon May 15 2017	Steven Levine
Revision 1.0-7 Updated Version for 7.3	Mon Mar 27 2017	Steven Levine
Revision 1.0-5 Version for 7.3 GA publication	Mon Oct 17 2016	Steven Levine
Revision 1.0-4 Preparing document for 7.3 Beta publication.	Wed Aug 17 2016	Steven Levine
Revision 0.3-4 Preparing document for 7.2 GA publication	Mon Nov 9 2015	Steven Levine
Revision 0.3-2 Preparing document for 7.2 Beta publication.	Wed Aug 19 2015	Steven Levine
Revision 0.2-7 Version for 7.1 GA release	Mon Feb 16 2015	Steven Levine
Revision 0.2-6 Version for 7.1 Beta release	Thu Dec 11 2014	Steven Levine
Revision 0.1-22 Version for 7.0 GA release	Mon Jun 2 2014	Steven Levine
Revision 0.1-1 Branched from the Red Hat Enterprise Linux 6 version of the document.	Wed Jan 16 2013	Steven Levine

INDEX

Symbols

[/lib/udev/rules.d directory](#), [udev Integration with the Device Mapper](#)

A

activating logical volumes

individual nodes, [Activating Logical Volumes on Individual Nodes in a Cluster](#)

activating volume groups, [Activating and Deactivating Volume Groups](#)

administrative procedures, [LVM Administration Overview](#)

allocation, [LVM Allocation](#)

policy, [Creating Volume Groups](#)

preventing, [Preventing Allocation on a Physical Volume](#)

archive file, [Logical Volume Backup](#), [Backing Up Volume Group Metadata](#)

B

backup

file, [Logical Volume Backup](#)

metadata, [Logical Volume Backup](#), [Backing Up Volume Group Metadata](#)

backup file, [Backing Up Volume Group Metadata](#)

block device

scanning, [Scanning for Block Devices](#)

C

cache file

building, [Scanning Disks for Volume Groups to Build the Cache File](#)

cache logical volume

creation, [Creating LVM Cache Logical Volumes](#)

cache volumes, [Cache Volumes](#)

cluster environment, [LVM Logical Volumes in a Red Hat High Availability Cluster](#)

CLVM

definition, [LVM Logical Volumes in a Red Hat High Availability Cluster](#)

command line units, [Using CLI Commands](#)

configuration examples, [LVM Configuration Examples](#)

creating

logical volume, [Creating Linear Logical Volumes](#)

logical volume, example, [Creating an LVM Logical Volume on Three Disks](#)

physical volumes, [Creating Physical Volumes](#)

striped logical volume, example, [Creating a Striped Logical Volume](#)
volume group, clustered, [Creating Volume Groups in a Cluster](#)
volume groups, [Creating Volume Groups](#)

creating LVM volumes

overview, [Logical Volume Creation Overview](#)

D

data relocation, online, [Online Data Relocation](#)

deactivating volume groups, [Activating and Deactivating Volume Groups](#)

device numbers

major, [Persistent Device Numbers](#)

minor, [Persistent Device Numbers](#)

persistent, [Persistent Device Numbers](#)

device path names, [Using CLI Commands](#)

device scan filters, [Controlling LVM Device Scans with Filters](#)

device size, maximum, [Creating Volume Groups](#)

device special file directory, [Creating Volume Groups](#)

display

sorting output, [Sorting LVM Reports](#)

displaying

logical volumes, [Displaying Logical Volumes](#), [The lvs Command](#)

physical volumes, [Displaying Physical Volumes](#), [The pvs Command](#)

volume groups, [Displaying Volume Groups](#), [The vgs Command](#)

E

extent

allocation, [Creating Volume Groups](#), [LVM Allocation](#)

definition, [Volume Groups](#), [Creating Volume Groups](#)

F

failed devices

displaying, [Displaying Information on Failed Devices](#)

features, new and changed, [New and Changed Features](#)

file system

growing on a logical volume, [Growing a File System on a Logical Volume](#)

filters, [Controlling LVM Device Scans with Filters](#)

G

growing file system

logical volume, [Growing a File System on a Logical Volume](#)

H

help display, [Using CLI Commands](#)

I

initializing

partitions, [Initializing Physical Volumes](#)

physical volumes, [Initializing Physical Volumes](#)

Insufficient Free Extents message, [Insufficient Free Extents for a Logical Volume](#)

L

linear logical volume

converting to mirrored, [Changing Mirrored Volume Configuration](#)

creation, [Creating Linear Logical Volumes](#)

definition, [Linear Volumes](#)

logging, [Logging](#)

logical volume

activation, [Controlling Logical Volume Activation](#)

administration, general, [Logical Volume Administration](#)

cache, [Creating LVM Cache Logical Volumes](#)

changing parameters, [Changing the Parameters of a Logical Volume Group](#)

creation, [Creating Linear Logical Volumes](#)

creation example, [Creating an LVM Logical Volume on Three Disks](#)

definition, [Logical Volumes](#), [LVM Logical Volumes](#)

displaying, [Displaying Logical Volumes](#), [Customized Reporting for LVM](#), [The lvs Command](#)

exclusive access, [Activating Logical Volumes on Individual Nodes in a Cluster](#)

extending, [Growing Logical Volumes](#)

growing, [Growing Logical Volumes](#)

historical, [Tracking and Displaying Historical Logical Volumes \(Red Hat Enterprise Linux 7.3 and Later\)](#)

linear, [Creating Linear Logical Volumes](#)

local access, [Activating Logical Volumes on Individual Nodes in a Cluster](#)

lvs display arguments, [The lvs Command](#)

mirrored, [Creating Mirrored Volumes](#)

reducing, [Shrinking Logical Volumes](#)

removing, [Removing Logical Volumes](#)

renaming, [Renaming Logical Volumes](#)

snapshot, [Creating Snapshot Volumes](#)

striped, [Creating Striped Volumes](#)

thinly-provisioned, [Creating Thinly-Provisioned Logical Volumes](#)

thinly-provisioned snapshot, [Creating Thinly-Provisioned Snapshot Volumes](#)

lvchange command, [Changing the Parameters of a Logical Volume Group](#)

lvconvert command, [Changing Mirrored Volume Configuration](#)

lvcreate command, [Creating Linear Logical Volumes](#)

lvdisplay command, [Displaying Logical Volumes](#)

lvextend command, [Growing Logical Volumes](#)

LVM

architecture overview, [LVM Architecture Overview](#)

clustered, [LVM Logical Volumes in a Red Hat High Availability Cluster](#)

components, [LVM Architecture Overview](#), [LVM Components](#)

custom report format, [Customized Reporting for LVM](#)

directory structure, [Creating Volume Groups](#)

help, [Using CLI Commands](#)

label, [Physical Volumes](#)

logging, [Logging](#)

logical volume administration, [Logical Volume Administration](#)

physical volume administration, [Physical Volume Administration](#)

physical volume, definition, [Physical Volumes](#)

volume group, definition, [Volume Groups](#)

lvmdiskscan command, [Scanning for Block Devices](#)

lvmetad daemon, [The Metadata Daemon \(lvmetad\)](#)

lvreduce command, [Shrinking Logical Volumes](#)

lvremove command, [Removing Logical Volumes](#)

lvrename command, [Renaming Logical Volumes](#)

lvs command, [Customized Reporting for LVM](#), [The lvs Command](#)

display arguments, [The lvs Command](#)

lvscan command, [Displaying Logical Volumes](#)

M

man page display, [Using CLI Commands](#)

metadata

backup, [Logical Volume Backup](#), [Backing Up Volume Group Metadata](#)

recovery, [Recovering Physical Volume Metadata](#)

metadata daemon, [The Metadata Daemon \(lvmetad\)](#)

mirrored logical volume

clustered, [Creating a Mirrored LVM Logical Volume in a Cluster](#)

converting to linear, [Changing Mirrored Volume Configuration](#)

creation, [Creating Mirrored Volumes](#)

failure policy, [Mirrored Logical Volume Failure Policy](#)
 failure recovery, [Recovering from LVM Mirror Failure](#)
 reconfiguration, [Changing Mirrored Volume Configuration](#)

mirror_image_fault_policy configuration parameter, [Mirrored Logical Volume Failure Policy](#)
 mirror_log_fault_policy configuration parameter, [Mirrored Logical Volume Failure Policy](#)

O

online data relocation, [Online Data Relocation](#)
 overview
 features, new and changed, [New and Changed Features](#)

P

partition type, setting, [Setting the Partition Type](#)
 partitions
 multiple, [Multiple Partitions on a Disk](#)

path names, [Using CLI Commands](#)
 persistent device numbers, [Persistent Device Numbers](#)
 physical extent
 preventing allocation, [Preventing Allocation on a Physical Volume](#)

physical volume

adding to a volume group, [Adding Physical Volumes to a Volume Group](#)
 administration, general, [Physical Volume Administration](#)
 creating, [Creating Physical Volumes](#)
 definition, [Physical Volumes](#)
 display, [The pvs Command](#)
 displaying, [Displaying Physical Volumes](#), [Customized Reporting for LVM](#)
 illustration, [LVM Physical Volume Layout](#)
 initializing, [Initializing Physical Volumes](#)
 layout, [LVM Physical Volume Layout](#)
 pvs display arguments, [The pvs Command](#)
 recovery, [Replacing a Missing Physical Volume](#)
 removing, [Removing Physical Volumes](#)
 removing from volume group, [Removing Physical Volumes from a Volume Group](#)
 removing lost volume, [Removing Lost Physical Volumes from a Volume Group](#)
 resizing, [Resizing a Physical Volume](#)

pvdisplay command, [Displaying Physical Volumes](#)
 pvmove command, [Online Data Relocation](#)
 pvremove command, [Removing Physical Volumes](#)
 pvresize command, [Resizing a Physical Volume](#)

pvs command, [Customized Reporting for LVM](#)

display arguments, [The pvs Command](#)

pvscan command, [Displaying Physical Volumes](#)

R

RAID logical volume, [RAID Logical Volumes](#)

extending, [Extending a RAID Volume](#)

growing, [Extending a RAID Volume](#)

reducing

logical volume, [Shrinking Logical Volumes](#)

removing

disk from a logical volume, [Removing a Disk from a Logical Volume](#)

logical volume, [Removing Logical Volumes](#)

physical volumes, [Removing Physical Volumes](#)

renaming

logical volume, [Renaming Logical Volumes](#)

volume group, [Renaming a Volume Group](#)

report format, LVM devices, [Customized Reporting for LVM](#)

resizing

physical volume, [Resizing a Physical Volume](#)

rules.d directory, [udev Integration with the Device Mapper](#)

S

scanning

block devices, [Scanning for Block Devices](#)

scanning devices, filters, [Controlling LVM Device Scans with Filters](#)

snapshot logical volume

creation, [Creating Snapshot Volumes](#)

snapshot volume

definition, [Snapshot Volumes](#)

striped logical volume

creation, [Creating Striped Volumes](#)

creation example, [Creating a Striped Logical Volume](#)

definition, [Striped Logical Volumes](#)

extending, [Extending a Striped Volume](#)

growing, [Extending a Striped Volume](#)

T

thin snapshot volume, [Thinly-Provisioned Snapshot Volumes](#)

thin volume

creation, [Creating Thinly-Provisioned Logical Volumes](#)

thinly-provisioned logical volume, [Thinly-Provisioned Logical Volumes \(Thin Volumes\)](#)

creation, [Creating Thinly-Provisioned Logical Volumes](#)

thinly-provisioned snapshot logical volume

creation, [Creating Thinly-Provisioned Snapshot Volumes](#)

thinly-provisioned snapshot volume, [Thinly-Provisioned Snapshot Volumes](#)

troubleshooting, [LVM Troubleshooting](#)

U

udev device manager, [Device Mapper Support for the udev Device Manager](#)

udev rules, [udev Integration with the Device Mapper](#)

units, command line, [Using CLI Commands](#)

V

verbose output, [Using CLI Commands](#)

vgcfgbackup command, [Backing Up Volume Group Metadata](#)

vgcfgrestore command, [Backing Up Volume Group Metadata](#)

vgchange command, [Changing the Parameters of a Volume Group](#)

vgcreate command, [Creating Volume Groups](#), [Creating Volume Groups in a Cluster](#)

vgdisplay command, [Displaying Volume Groups](#)

vgexport command, [Moving a Volume Group to Another System](#)

vgextend command, [Adding Physical Volumes to a Volume Group](#)

vgimport command, [Moving a Volume Group to Another System](#)

vgmerge command, [Combining Volume Groups](#)

vgmknodes command, [Recreating a Volume Group Directory](#)

vgreduce command, [Removing Physical Volumes from a Volume Group](#)

vgrename command, [Renaming a Volume Group](#)

vgs command, [Customized Reporting for LVM](#)

display arguments, [The vgs Command](#)

vgscan command, [Scanning Disks for Volume Groups to Build the Cache File](#)

vgsplit command, [Splitting a Volume Group](#)

volume group

activating, [Activating and Deactivating Volume Groups](#)

administration, general, [Volume Group Administration](#)

changing parameters, [Changing the Parameters of a Volume Group](#)

combining, [Combining Volume Groups](#)

- creating, [Creating Volume Groups](#)
- creating in a cluster, [Creating Volume Groups in a Cluster](#)
- deactivating, [Activating and Deactivating Volume Groups](#)
- definition, [Volume Groups](#)
- displaying, [Displaying Volume Groups](#), [Customized Reporting for LVM](#), [The vgs Command](#)
- extending, [Adding Physical Volumes to a Volume Group](#)
- growing, [Adding Physical Volumes to a Volume Group](#)
- merging, [Combining Volume Groups](#)
- moving between systems, [Moving a Volume Group to Another System](#)
- reducing, [Removing Physical Volumes from a Volume Group](#)
- removing, [Removing Volume Groups](#)
- renaming, [Renaming a Volume Group](#)
- shrinking, [Removing Physical Volumes from a Volume Group](#)
- splitting, [Splitting a Volume Group](#)
 - example procedure, [Splitting a Volume Group](#)
- vgs display arguments, [The vgs Command](#)