Red Hat Enterprise Linux 6

Logical Volume Manager Administration

LVM Administrator Guide

Red Hat Enterprise Linux 6 Logical Volume Manager Administration

LVM Administrator Guide

Steven Levine
Red Hat Customer Content Services
slevine@redhat.com

Legal Notice

Abstract

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

Table of Contents

CHAPTER 1. INTRODUCTION

This book describes the Logical Volume Manager (LVM), including information on running LVM in a clustered environment.

1.1. AUDIENCE

This book is intended to be used by system administrators managing systems running the Linux operating system. It requires familiarity with Red Hat Enterprise Linux 6.

1.2. SOFTWARE VERSIONS

Table 1.1. Software Versions

Software	Description
Red Hat Enterprise Linux 6	refers to Red Hat Enterprise Linux 6 and higher
GFS2	refers to GFS2 for Red Hat Enterprise Linux 6 and higher

1.3. RELATED DOCUMENTATION

For more information about using Red Hat Enterprise Linux, see the following resources:

- *Installation Guide* — Documents relevant information regarding the installation of Red Hat Enterprise Linux 6.

- *Deployment Guide* — Documents relevant information regarding the deployment, configuration and administration of Red Hat Enterprise Linux 6.

- *Storage Administration Guide* — Provides instructions on how to effectively manage storage devices and file systems on Red Hat Enterprise Linux 6.

For more information about the High Availability Add-On and the Resilient Storage Add-On for Red Hat Enterprise Linux 6, see the following resources:

- *High Availability Add-On Overview* — Provides a high-level overview of the Red Hat High Availability Add-On.

- *Cluster Administration* — Provides information about installing, configuring and managing the Red Hat High Availability Add-On,

- *Global File System 2: Configuration and Administration* — Provides information about installing, configuring, and maintaining Red Hat GFS2 (Red Hat Global File System 2), which is included in the Resilient Storage Add-On.

- *DM Multipath* — Provides information about using the Device-Mapper Multipath feature of Red Hat Enterprise Linux 6.

- *Load Balancer Administration* — Provides information on configuring high-performance systems and services with the Load Balancer Add-On, a set of integrated software components that provide Linux Virtual Servers (LVS) for balancing IP load across a set of real servers.

- *Release Notes* — Provides information about the current release of Red Hat products.

Red Hat documents are available in HTML, PDF, and RPM versions on the Red Hat Enterprise Linux Documentation CD and online at https://access.redhat.com/site/documentation/.

1.4. WE NEED FEEDBACK!

If you find a typographical error in this manual, or if you have thought of a way to make this manual better, we would love to hear from you! Please submit a report in Bugzilla: http://bugzilla.redhat.com/ against the product **Red Hat Enterprise Linux 6**and the component **doc-Logical_Volume_Manager**. When submitting a bug report, be sure to mention the manual's identifier:

> Logical_Volume_Manager_Administration(EN)-6 (2017-3-8-15:20)

If you have a suggestion for improving the documentation, try to be as specific as possible when describing it. If you have found an error, include the section number and some of the surrounding text so we can find it easily.

CHAPTER 2. THE LVM LOGICAL VOLUME MANAGER

This chapter provides a summary of the features of the LVM logical volume manager that are new for the initial and subsequent releases of Red Hat Enterprise Linux 6. Following that, this chapter provides a high-level overview of the components of the Logical Volume Manager (LVM).

2.1. NEW AND CHANGED FEATURES

This section lists new and changed features of the LVM logical volume manager that are included with the initial and subsequent releases of Red Hat Enterprise Linux 6.

2.1.1. New and Changed Features for Red Hat Enterprise Linux 6.0

Red Hat Enterprise Linux 6.0 includes the following documentation and feature updates and changes.

- 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 this parameter is set to **remove**, the system attempts to remove the faulty device and run without it. When this parameter is 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. For information on the LVM mirror failure policies, see Section 5.4.3.1, "Mirrored Logical Volume Failure Policy" .

- For the Red Hat Enterprise Linux 6 release, the Linux I/O stack has been enhanced to process vendor-provided I/O limit information. This allows storage management tools, including LVM, to optimize data placement and access. This support can be disabled by changing the default values of **data_alignment_detection** and **data_alignment_offset_detection** in the **lvm.conf** file, although disabling this support is not recommended.

  For information on data alignment in LVM as well as information on changing the default values of **data_alignment_detection** and **data_alignment_offset_detection**, see the inline documentation for the **/etc/lvm/lvm.conf** file, which is also documented in Appendix B, *The LVM Configuration Files*. For general information on support for the I/O Stack and I/O limits in Red Hat Enterprise Linux 6, see the *Storage Administration Guide*.

- In Red Hat Enterprise Linux 6, 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. For information on Device Mapper support for the **udev** device manager, see Section A.3, "Device Mapper Support for the udev Device Manager" .

- For the Red Hat Enterprise Linux 6 release, you can use the **lvconvert --repair** command to repair a mirror after disk failure. This brings the mirror back into a consistent state. For information on the **lvconvert --repair** command, see Section 5.4.3.3, "Repairing a Mirrored Logical Device".

- As of the Red Hat Enterprise Linux 6 release, you can use the **--merge** option of the **lvconvert** command to merge a snapshot into its origin volume. For information on merging snapshots, see Section 5.4.8, "Merging Snapshot Volumes" .

- As of the Red Hat Enterprise Linux 6 release, you can use the **--splitmirrors** argument of the **lvconvert** command to split off a redundant image of a mirrored logical volume to form a new logical volume. For information on using this option, see Section 5.4.3.2, "Splitting Off a Redundant Image of a Mirrored Logical Volume".

- You can now create a mirror log for a mirrored logical device that is itself mirrored by using the **--mirrorlog mirrored** argument of the **lvcreate** command when creating a mirrored logical device. For information on using this option, see Section 5.4.3, "Creating Mirrored Volumes".

2.1.2. New and Changed Features for Red Hat Enterprise Linux 6.1

Red Hat Enterprise Linux 6.1 includes the following documentation and feature updates and changes.

- The Red Hat Enterprise Linux 6.1 release supports the creation of snapshot logical volumes of mirrored logical volumes. You create a snapshot of a mirrored volume just as you would create a snapshot of a linear or striped logical volume. For information on creating snapshot volumes, see Section 5.4.5, "Creating Snapshot Volumes" .

- When extending an LVM volume, you can now 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 information on extending LVM mirrored volumes with the **--alloc cling** option of the **lvextend** command, see Section 5.4.14.3, "Extending a Logical Volume with the **cling** Allocation Policy".

- You can now specify multiple **--addtag** and **--deltag** arguments within a single **pvchange**, **vgchange**, or **lvchange** command. For information on adding and removing object tags, see Section D.1, "Adding and Removing Object Tags".

- The list of allowed characters in LVM object tags has been extended, and tags can contain the "/", "=", "!", ":", "#", and "&" characters. For information on LVM object tags, see Appendix D, *LVM Object Tags*.

- You can now combine RAID0 (striping) and RAID1 (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. For information on creating mirrored logical volumes, see Section 5.4.3, "Creating Mirrored Volumes".

- As of the Red Hat Enterprise Linux 6.1 release, 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 5.7, "Activating Logical Volumes on Individual Nodes in a Cluster" .

2.1.3. New and Changed Features for Red Hat Enterprise Linux 6.2

Red Hat Enterprise Linux 6.2 includes the following documentation and feature updates and changes.

- The Red Hat Enterprise Linux 6.2 release supports the **issue_discards** parameter in the **lvm.conf** configuration file. When this parameter is set, LVM will issue discards to a logical volume's underlying physical volumes when the logical volume is no longer using the space on the physical volumes. For information on this parameter, see the inline documentation for the **/etc/lvm/lvm.conf** file, which is also documented in Appendix B, *The LVM Configuration Files* .

2.1.4. New and Changed Features for Red Hat Enterprise Linux 6.3

Red Hat Enterprise Linux 6.3 includes the following documentation and feature updates and changes.

- As of the Red Hat Enterprise Linux 6.3 release, LVM supports RAID4/5/6 and a new implementation of mirroring. For information on RAID logical volumes, see Section 5.4.16, "RAID Logical Volumes".

- 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. For information on creating mirrored volumes, see Section 5.4.3, "Creating Mirrored Volumes".

- This manual now documents the snapshot **autoextend** feature. For information on creating snapshot volumes, see Section 5.4.5, "Creating Snapshot Volumes" .

2.1.5. New and Changed Features for Red Hat Enterprise Linux 6.4

Red Hat Enterprise Linux 6.4 includes the following documentation and feature updates and changes.

- Logical volumes can now 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, to 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.

  For general information on thinly-provisioned logical volumes, see Section 3.3.5, "Thinly-Provisioned Logical Volumes (Thin Volumes)". For information on creating thin volumes, see Section 5.4.4, "Creating Thinly-Provisioned Logical Volumes" .

- The Red Hat Enterprise Linux release 6.4 version of LVM 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.

  For general information on thinly-provisioned snapshot volumes, see Section 3.3.7, "Thinly-Provisioned Snapshot Volumes". For information on creating thin snapshot volumes, see Section 5.4.6, "Creating Thinly-Provisioned Snapshot Volumes" .

- This document includes a new section detailing LVM allocation policy, Section 5.3.2, "LVM Allocation".

- LVM now provides support for **raid10** logical volumes. For information on RAID logical volumes, see Section 5.4.16, "RAID Logical Volumes" .

- The LVM metadata daemon, **lvmetad**, is supported in Red Hat Enterprise Linux release 6.4. Enabling this daemon reduces the amount of scanning on systems with many block devices. The **lvmetad** daemon is not currently supported across the nodes of a cluster, and requires that the locking type be local file-based locking.

  For information on the metadata daemon, see Section 4.6, "The Metadata Daemon (lvmetad)" .

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

2.1.6. New and Changed Features for Red Hat Enterprise Linux 6.5

Red Hat Enterprise Linux 6.5 includes the following documentation and feature updates and changes.

- You can control I/O operations on a RAID1 logical volume with the **--writemostly** and **--writebehind** parameters of the **lvchange** command. For information on these parameters, see Section 5.4.16.11, "Controlling I/O Operations on a RAID1 Logical Volume" .

- The **lvchange** command now supports a **--refresh** parameter that allows you to restore a transiently failed device without having to reactivate the device. This feature is described in Section 5.4.16.8.1, "The allocate RAID Fault Policy" .

- LVM provides scrubbing support for RAID logical volumes. For information on this feature, see Section 5.4.16.10, "Scrubbing a RAID Logical Volume" .

- The fields that the **lvs** command supports have been updated. For information on the **lvs** command, see Table 5.4, "lvs Display Fields" .

- The **lvchange** command supports the new **--maxrecoveryrate** and **--minrecoveryrate** parameters, which allow you to control the rate at which **sync** operations are performed. For information on these parameters, see Section 5.4.16.10, "Scrubbing a RAID Logical Volume" .

- 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, as described in Section 5.4.16.1, "Creating a RAID Logical Volume".

- You can now create a thinly-provisioned snapshot of a non-thinly-provisioned logical volume. For information on creating these volumes, known as external volumes, see Section 3.3.7, "Thinly-Provisioned Snapshot Volumes".

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

2.1.7. New and Changed Features for Red Hat Enterprise Linux 6.6

Red Hat Enterprise Linux 6.6 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 3.3.5, "Thinly-Provisioned Logical Volumes (Thin Volumes)". For information on thinly-provisioned snapshot volumes, see Section 3.3.7, "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 6.8 release, this command was renamed **lvmconf**, 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 4.7, "Displaying LVM Information with the **lvm** Command".

- In the Red Hat Enterprise Linux 6.6 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 5.4.17, "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, see Section 5.3.15, "Moving a Volume Group to Another System".

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

2.1.8. New and Changed Features for Red Hat Enterprise Linux 6.7

Red Hat Enterprise Linux 6.7 includes the following documentation and feature updates and changes.

- As of Red Hat Enterprise Linux release 6.7, many LVM processing commands 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 5.4.7, "Creating LVM Cache Logical Volumes".

- The troubleshooting chapter of this document includes a new section, Section 7.8, "Duplicate PV Warnings for Multipathed Devices".

2.1.9. New and Changed Features for Red Hat Enterprise Linux 6.8

Red Hat Enterprise Linux 6.8 includes the following documentation and feature updates and changes.

- When defining selection criteria for LVM commands, you can now specify time values as selection criteria for fields with a field type of **time**. For information on specifying time values as selection criteria, see Section C.3.1, "Specifying Time Values".

2.2. 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 groups, which you can name according to your convenience.

- 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.

2.3. LVM ARCHITECTURE OVERVIEW

For the Red Hat Enterprise Linux 4 release of the Linux operating system, the original LVM1 logical volume manager was replaced by LVM2, which has a more generic kernel framework than LVM1. LVM2 provides the following improvements over LVM1:

- flexible capacity

- more efficient metadata storage

- better recovery format

- new ASCII metadata format

- atomic changes to metadata

- redundant copies of metadata

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 2.1, "LVM Logical Volume Components" shows the components of a simple LVM logical volume:
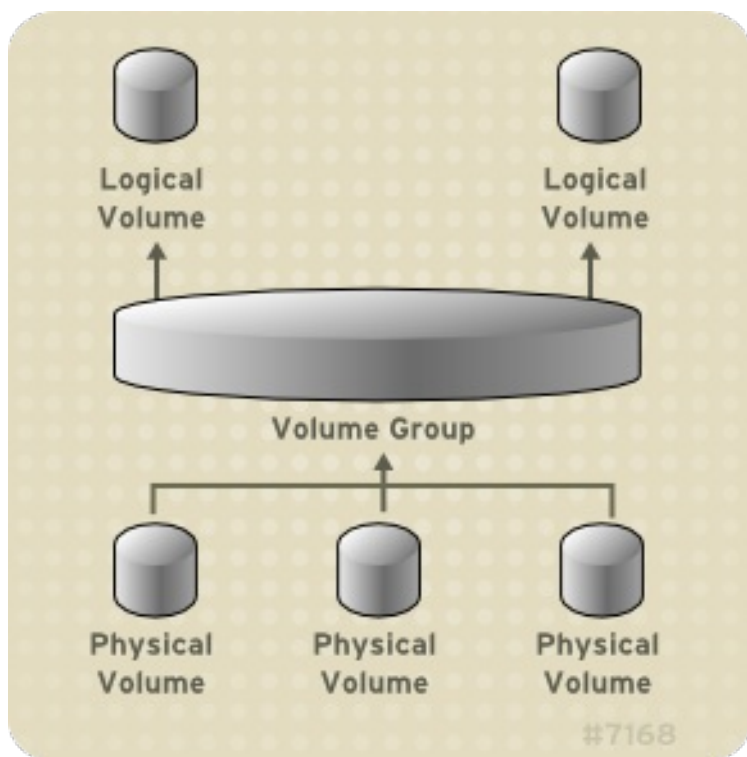
Figure 2.1. LVM Logical Volume Components

For detailed information on the components of an LVM logical volume, see Chapter 3, *LVM Components*.

2.4. THE CLUSTERED LOGICAL VOLUME MANAGER (CLVM)

The Clustered Logical Volume Manager (CLVM) is a set of clustering extensions to LVM. These extensions allow a cluster of computers to manage shared storage (for example, on a SAN) using LVM. CLVM is part of the Resilient Storage Add-On.

Whether you should use CLVM depends on your system requirements:

- If only one node of your system requires access to the storage you are configuring as logical volumes, then you can use LVM without the CLVM extensions and the logical volumes created with that node are all local to the node.

- If you are using a clustered system for failover where only a single node that accesses the storage is active at any one time, you should use High Availability Logical Volume Management agents (HA-LVM).

- If more than one node of your cluster will require access to your storage which is then shared among the active nodes, then you must use CLVM. CLVM allows a user to configure logical volumes on shared storage by locking access to physical storage while a logical volume is being configured, and uses clustered locking services to manage the shared storage.

In order to use CLVM, the High Availability Add-On and Resilient Storage Add-On software, including the **clvmd** daemon, must be running. The **clvmd** daemon is the key clustering extension to LVM. The **clvmd** daemon runs in each cluster computer and distributes LVM metadata updates in a cluster, presenting each cluster computer with the same view of the logical volumes. For information on installing and administering the High Availability Add-On see *Cluster Administration*.

To ensure that **clvmd** is started at boot time, you can execute a **chkconfig ... on** command on the **clvmd** service, as follows:

```
# chkconfig clvmd on
```

If the **clvmd** daemon has not been started, you can execute a **service ... start** command on the **clvmd** service, as follows:

```
# service clvmd start
```

Creating LVM logical volumes in a cluster environment is identical to creating LVM logical volumes on a single node. There is no difference in the LVM commands themselves, or in the LVM graphical user interface, as described in Chapter 5, *LVM Administration with CLI Commands* and Chapter 8, *LVM Administration with the LVM GUI*. In order to enable the LVM volumes you are creating in a cluster, the cluster infrastructure must be running and the cluster must be quorate.

By default, logical volumes created with CLVM on shared storage are visible to all systems that have access to the shared storage. It is possible to create volume groups in which all of the storage devices are visible to only one node in the cluster. It is also possible to change the status of a volume group from a local volume group to a clustered volume group. For information, see Section 5.3.3, "Creating Volume Groups in a Cluster" and Section 5.3.8, "Changing the Parameters of a Volume Group".

> **WARNING**
>
> When you create volume groups with CLVM on shared storage, you must ensure that all nodes in the cluster have access to the physical volumes that constitute the volume group. Asymmetric cluster configurations in which some nodes have access to the storage and others do not are not supported.

Figure 2.2, "CLVM Overview" shows a CLVM overview in a cluster.
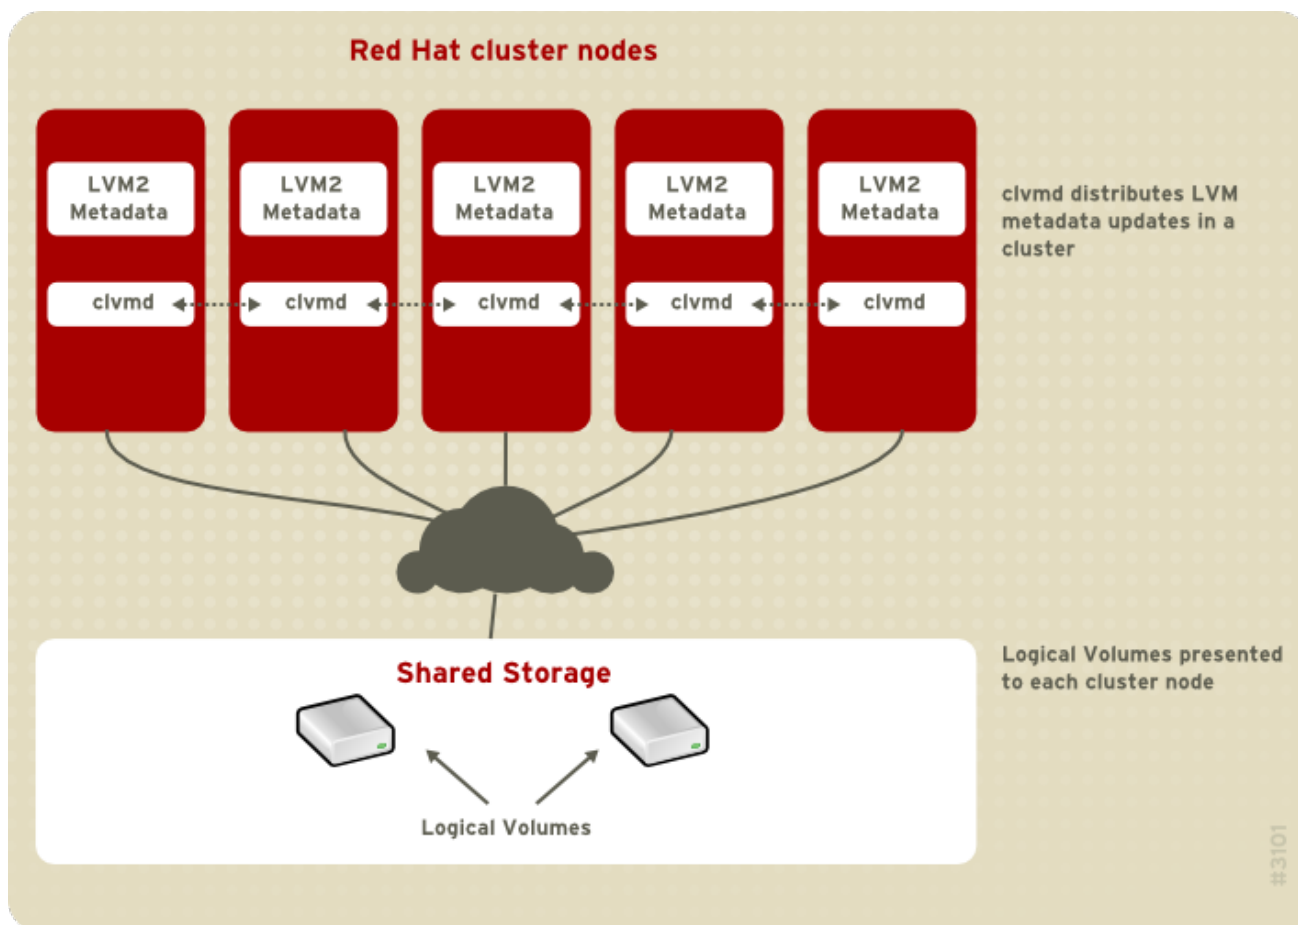
Figure 2.2. CLVM Overview

> **NOTE**
>
> CLVM requires changes to the **lvm.conf** file for cluster-wide locking. Information on configuring the **lvm.conf** file to support clustered locking is provided within the **lvm.conf** file itself. For information about the **lvm.conf** file, see Appendix B, *The LVM Configuration Files*.

2.5. DOCUMENT OVERVIEW

This remainder of this document includes the following chapters:

- Chapter 3, *LVM Components* describes the components that make up an LVM logical volume.

- Chapter 4, *LVM Administration Overview* provides an overview of the basic steps you perform to configure LVM logical volumes, whether you are using the LVM Command Line Interface (CLI) commands or the LVM Graphical User Interface (GUI).

- Chapter 5, *LVM Administration with CLI Commands* summarizes the individual administrative tasks you can perform with the LVM CLI commands to create and maintain logical volumes.

- Chapter 6, *LVM Configuration Examples* provides a variety of LVM configuration examples.

- Chapter 7, *LVM Troubleshooting* provides instructions for troubleshooting a variety of LVM issues.

- Chapter 8, *LVM Administration with the LVM GUI* summarizes the operating of the LVM GUI.

- Appendix A, *The Device Mapper* describes the Device Mapper that LVM uses to map logical and physical volumes.

- Appendix B, *The LVM Configuration Files* describes the LVM configuration files.

- Appendix D, *LVM Object Tags* describes LVM object tags and host tags.

- Appendix E, *LVM Volume Group Metadata* describes LVM volume group metadata, and includes a sample copy of metadata for an LVM volume group.

CHAPTER 3. LVM COMPONENTS

This chapter describes the components of an LVM Logical volume.

3.1. 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. 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* .

3.1.1. LVM Physical Volume Layout

Figure 3.1, "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 3.1. Physical Volume layout

3.1.2. Multiple Partitions on a Disk

LVM allows you to create physical volumes out of disk partitions. It is generally recommended 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-up.

- 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.

3.2. 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.

3.3. LVM LOGICAL VOLUMES

In LVM, a volume group is divided up into logical volumes. There are three types of LVM logical volumes: *linear* volumes, *striped* volumes, and *mirrored* volumes. These are described in the following sections.

3.3.1. 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 3.2, "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.
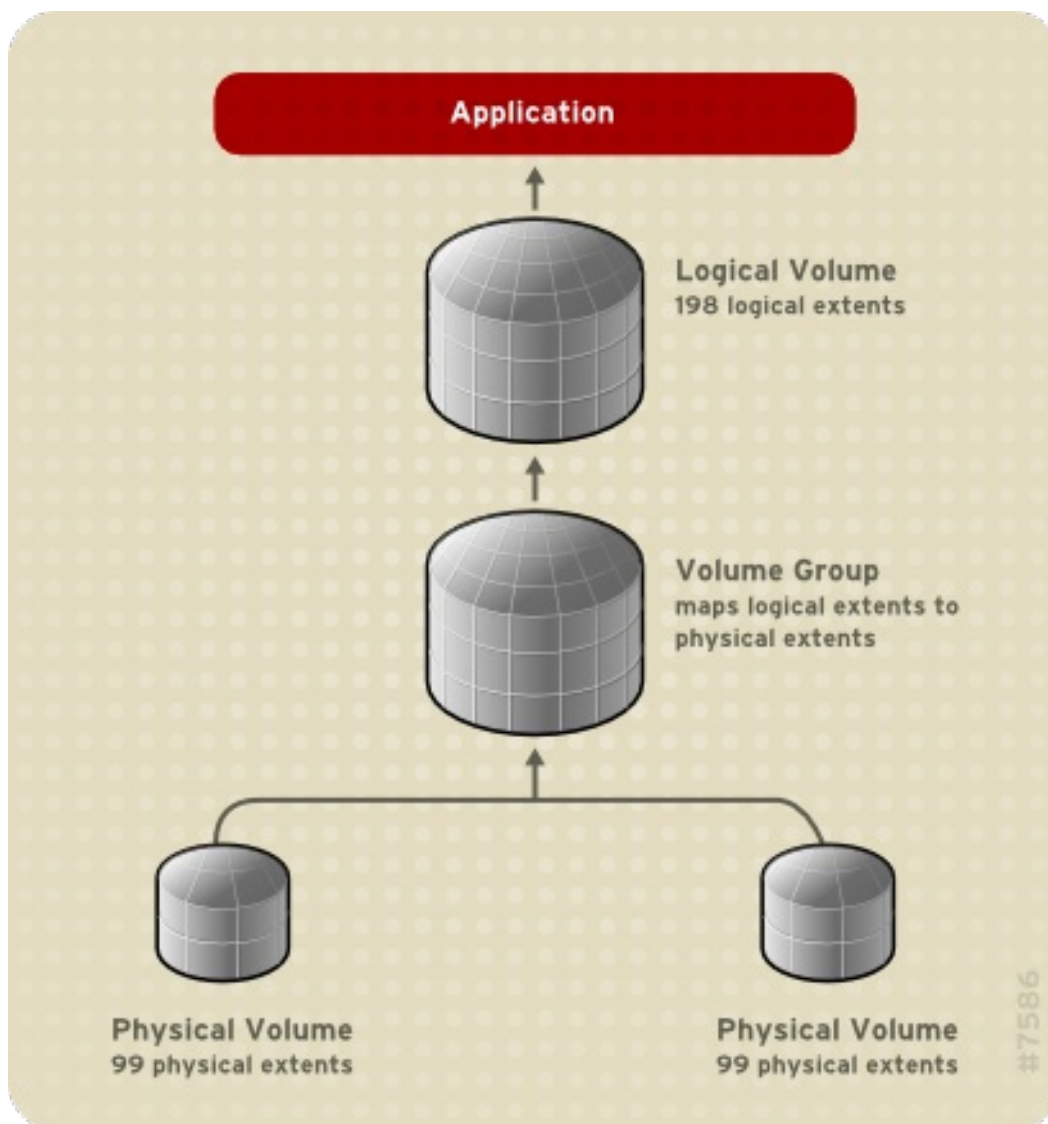
Figure 3.2. Extent Mapping

The physical volumes that make up a logical volume do not have to be the same size. Figure 3.3, "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.
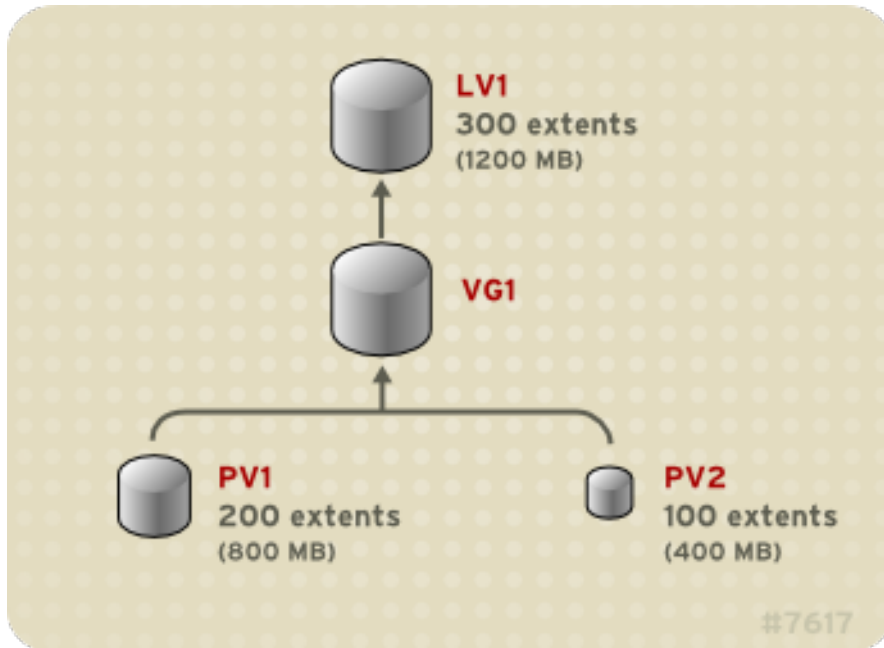
Figure 3.3. Linear Volume with Unequal Physical Volumes

You can configure more than one linear logical volume of whatever size you require from the pool of physical extents. Figure 3.4, "Multiple Logical Volumes" shows the same volume group as in Figure 3.3, "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 3.4. Multiple Logical Volumes

3.3.2. 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 PV1

- the second stripe of data is written to PV2

- the third stripe of data is written to PV3

- the fourth stripe of data is written to PV1

In a striped logical volume, the size of the stripe cannot exceed the size of an extent.
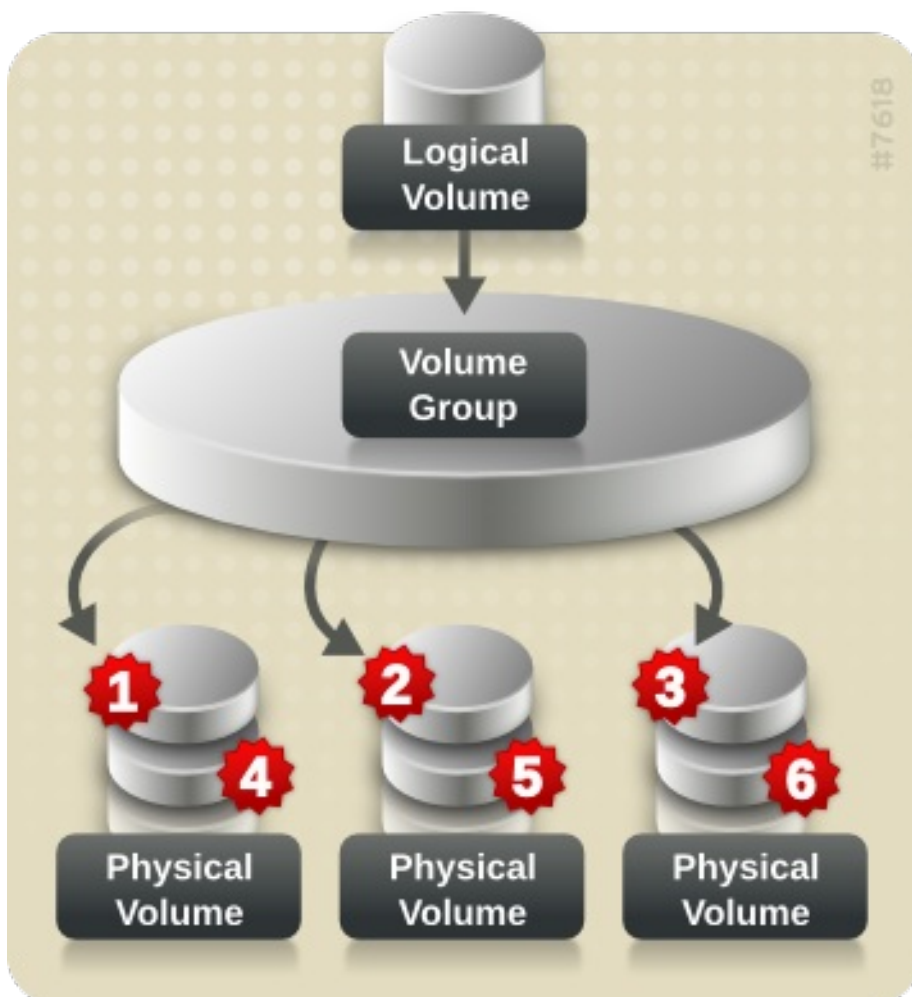
Figure 3.5. Striping Data Across Three PVs

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 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 5.4.14.1, "Extending a Striped Volume".

3.3.3. Mirrored Logical Volumes

A mirror maintains identical copies of data on different devices. When data is written to one device, it is written to a second device as well, mirroring the data. This provides protection for device failures. When one leg of a mirror fails, the logical volume becomes a linear volume and can still be accessed.

LVM supports mirrored volumes. When you create a mirrored logical volume, LVM ensures that data written to an underlying physical volume is mirrored onto a separate physical volume. With LVM, you can create mirrored logical volumes with multiple mirrors.

An LVM mirror divides the device being copied into regions that are typically 512KB in size. LVM maintains a small log which it uses to keep track of which regions are in sync with the mirror or mirrors. This log can be kept on disk, which will keep it persistent across reboots, or it can be maintained in memory.

Figure 3.6, "Mirrored Logical Volume" shows a mirrored logical volume with one mirror. In this configuration, the log is maintained on disk.

Figure 3.6. Mirrored Logical Volume

For information on creating and modifying mirrors, see Section 5.4.3, "Creating Mirrored Volumes".

3.3.4. RAID Logical Volumes

As of the Red Hat Enterprise Linux 6.3 release, LVM supports RAID logical volumes. For information on the RAID implementations that LVM supports, see Section 5.4.16, "RAID Logical Volumes".

3.3.5. Thinly-Provisioned Logical Volumes (Thin Volumes)

As of the Red Hat Enterprise Linux 6.4 release, 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 over-commit 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 5.4.4, "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 3.3.7, "Thinly-Provisioned Snapshot Volumes".

3.3.6. 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**
>
> As of the Red Hat Enterprise Linux 6.4 release, LVM supports thinly-provisioned snapshots. For information on thinly provisioned snapshot volumes, see Section 3.3.7, "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 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 resizeable, 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 5.4.5, "Creating Snapshot Volumes" .

As of the Red Hat Enterprise Linux 6 release, 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 5.4.8, "Merging Snapshot Volumes" .

3.3.7. Thinly-Provisioned Snapshot Volumes

The Red Hat Enterprise Linux release 6.4 version of LVM 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 5.4.6, "Creating Thinly-Provisioned Snapshot Volumes".

3.3.8. Cache Volumes

As of the Red Hat Enterprise Linux 6.7 release, LVM supports the use of fast block devices (such as SSD drives) as write-back or write-though 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 5.4.7, "Creating LVM Cache Logical Volumes".

CHAPTER 4. 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 6, *LVM Configuration Examples*.

For descriptions of the CLI commands you can use to perform LVM administration, see Chapter 5, *LVM Administration with CLI Commands*. Alternately, you can use the LVM GUI, which is described in Chapter 8, *LVM Administration with the LVM GUI* .

4.1. CREATING LVM VOLUMES IN A CLUSTER

To create logical volumes in a cluster environment, you use the Clustered Logical Volume Manager (CLVM), which is a set of clustering extensions to LVM. These extensions allow a cluster of computers to manage shared storage (for example, on a SAN) using LVM. In order to use CLVM, the High Availability Add-On and Resilient Storage Add-On software, including the **clvmd** daemon, must be started at boot time, as described in Section 2.4, "The Clustered Logical Volume Manager (CLVM)" .

Creating LVM logical volumes in a cluster environment is identical to creating LVM logical volumes on a single node. There is no difference in the LVM commands themselves, or in the LVM GUI interface. In order to enable the LVM volumes you are creating in a cluster, the cluster infrastructure must be running and the cluster must be quorate.

CLVM requires changes to the **lvm.conf** file for cluster-wide locking. Information on configuring the **lvm.conf** file to support clustered locking is provided within the   **lvm.conf** file itself. For information about the **lvm.conf** file, see Appendix B, *The LVM Configuration Files* .

By default, logical volumes created with CLVM on shared storage are visible to all systems that have access to the shared storage. It is possible to create volume groups in which all of the storage devices are visible to only one node in the cluster. It is also possible to change the status of a volume group from a local volume group to a clustered volume group. For information, see Section 5.3.3, "Creating Volume Groups in a Cluster" and Section 5.3.8, "Changing the Parameters of a Volume Group"

> ⚠️ **WARNING**
>
> When you create volume groups with CLVM on shared storage, you must ensure that all nodes in the cluster have access to the physical volumes that constitute the volume group. Asymmetric cluster configurations in which some nodes have access to the storage and others do not are not supported.

For information on how to install the High Availability Add-On and set up the cluster infrastructure, see *Cluster Administration*.

For an example of creating a mirrored logical volume in a cluster, see Section 6.5, "Creating a Mirrored LVM Logical Volume in a Cluster".

4.2. 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.

> **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 6 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).

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.

Alternately, you can create and mount the GFS2 file system with the LVM GUI.

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.

4.3. GROWING A FILE SYSTEM ON A LOGICAL VOLUME

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

1. Make a new physical volume.

2. Extend the volume group that contains the logical volume with the file system you are growing to include the new physical volume.

3. Extend the logical volume to include the new physical volume.

4. Grow the file system.

If you have sufficient unallocated space in the volume group, you can use that space to extend the logical volume instead of performing steps 1 and 2.

4.4. LOGICAL VOLUME BACKUP

Metadata backups and archives are automatically created on every volume group and logical volume configuration change 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/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 5.3.13, "Backing Up Volume Group Metadata" .

4.5. 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*.

4.6. THE METADATA DAEMON (LVMETAD)

LVM can optionally use a central metadata cache, implemented through a daemon (**lvmetad**) 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.

> **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.

To take advantage of the daemon, you must do the following:

- Start the daemon through the **lvm2-lvmetad** service. To start the daemon automatically at boot time, use the **chkconfig lvm2-lvmetad on** command. To start the daemon manually, use the **service lvm2-lvmetad start** command.

- Configure LVM to make use of the daemon by setting the **global/use_lvmetad** variable to 1 in the **lvm.conf** configuration file. For information on the **lvm.conf** configuration file, see Appendix B, *The LVM Configuration Files* .

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, by means of **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. For information on the **udev** device manager and **udev** rules, see Section A.3, "Device Mapper Support for the udev Device Manager" .

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 and logical volumes that should be automatically activated. Without the **lvmetad** daemon, a manual activation is necessary. By default, this list is not defined, which means that all volumes are autoactivated once all of the physical volumes are in place. The autoactivation works recursively for LVM stacked on top of other devices, as it is event-based.

> **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.

4.7. 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 built-in block device types (Red Hat Enterprise Linux release 6.6 and later).

- **lvm formats**

  Displays recognizes 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 5. 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.

> **NOTE**
>
> If you are creating or modifying an LVM volume for a clustered environment, you must ensure that you are running the **clvmd** daemon. For information, see Section 4.1, "Creating LVM Volumes in a Cluster".

5.1. 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 **lvol0** in a volume group called **vg0** can be specified as **vg0/lvol0**. 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 "lvol0" 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 lvol0
    Creating volume group backup "/etc/lvm/backup/new_vg" (seqno 5).
    Found volume group "new_vg"
    Creating new_vg-lvol0
    Loading new_vg-lvol0 table
    Resuming new_vg-lvol0 (253:2)
    Clearing start of logical volume "lvol0"
    Creating volume group backup "/etc/lvm/backup/new_vg" (seqno 5).
  Logical volume "lvol0" 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 lvm** 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 7.4, "Recovering Physical Volume Metadata".

5.2. PHYSICAL VOLUME ADMINISTRATION

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

5.2.1. Creating Physical Volumes

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

5.2.1.1. 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
```

5.2.1.2. 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
```

5.2.1.3. 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

5.2.2. 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 5.8, "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-ldwM-01S9-XO8M-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          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 **/etc/lvm/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 5.5, "Controlling LVM Device Scans with Filters".

5.2.3. 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.

5.2.4. 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.

5.2.5. 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 5.3.7, "Removing Physical Volumes from a Volume Group".

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

5.3. VOLUME GROUP ADMINISTRATION

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

5.3.1. 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 5.3.2, "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**, six device special files are created:

> /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.

5.3.2. 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 to be added to an existing logical volume 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 5.4.14.3, "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.

5.3.3. Creating Volume Groups in a Cluster

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

By default, volume groups created with CLVM 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 **-c n** 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 5.3.8, "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 5.3.5, "Displaying Volume Groups" Section 5.8, "Customized Reporting for LVM", and the **vgs** man page.

5.3.4. 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
```

5.3.5. 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 5.3.6, "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 5.8, "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 a **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-OpMO-0118-nlwO-wwqd-fD5D32
```

5.3.6. 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 a **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 **lvm.conf** file to restrict the scan to avoid specific devices. For information on using filters to control which devices are scanned, see Section 5.5, "Controlling LVM Device Scans with Filters".

The following example shows the output of a **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
```

5.3.7. 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-MOkn-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.

5.3.8. Changing the Parameters of a Volume Group

The **vgchange** command is used to deactivate and activate volume groups, as described in Section 5.3.9, "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.

5.3.9. 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 5.4.10, "Changing the Parameters of a Logical Volume Group" , For information on activating logical volumes on individual nodes in a cluster, see Section 5.7, "Activating Logical Volumes on Individual Nodes in a Cluster".

5.3.10. Removing Volume Groups

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

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

5.3.11. 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 off the new volume group **smallvg** from the original volume group **bigvg**.

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

5.3.12. 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
```

5.3.13. Backing Up Volume Group Metadata

Metadata backups and archives are automatically created on every volume group and logical volume configuration change 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 **vgcfrestore** 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 7.4, "Recovering Physical Volume Metadata" .

5.3.14. 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
```

5.3.15. 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**
>
> As of Red Hat Enterprise Linux 6.5, 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 form 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.

5.3.16. 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 removes unused ones.

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

5.4. LOGICAL VOLUME ADMINISTRATION

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

5.4.1. 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.

As of the Red Hat Enterprise Linux 6.3 release, you can use LVM to create, display, rename, use, and remove RAID logical volumes. For information on RAID logical volumes, see Section 5.4.16, "RAID Logical 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 MB 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 volume group to use for the logical volume. 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
```

You can also use the **-l** argument of the **lvcreate** command to specify the percentage of the remaining free space in a volume group as the size of the logical volume. 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 testvg -n mylv
```

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 5.3.7, "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 -ntestlv 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 5.3.1, "Creating Volume Groups".

5.4.2. 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 3.3.2, "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 64kB. The logical volume is 50 gigabytes in size, is named **gfslv**, and is carved out of volume group **vg0**.

```
# lvcreate -L 50G -i2 -I64 -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 -i2 -nstripelv testvg /dev/sda1:0-49 /dev/sdb1:50-99
  Using default stripesize 64.00 KB
  Logical volume "stripelv" created
```

5.4.3. Creating Mirrored Volumes

NOTE

As of the Red Hat Enterprise Linux 6.3 release, LVM supports RAID4/5/6 and a new implementation of mirroring. For information on this new implementation, see Section 5.4.16, "RAID Logical Volumes".

NOTE

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. 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 6.5, "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 -L 50G -m1 -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 -m1 -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. As of the Red Hat Enterprise Linux 6.3 release, 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 resynced 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 -L 12MB -m1 --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 -L 500M -m1 -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 -L 12MB -m1 --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 necessary 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 -L 500M -m1 -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 -L 500M -m1 -n mirrorlv vg0 /dev/sda1:0-499 /dev/sdb1:0-499 /dev/sdc1:0
```

NOTE

As of the Red Hat Enterprise Linux 6.1 release, 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.

5.4.3.1. 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 this parameter is 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 and the mirror will not remember its sync status across crashes and reboots and the entire mirror will be resynced.

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 7.3, "Recovering from LVM Mirror Failure".

5.4.3.2. 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, you 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
```

5.4.3.3. 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.

5.4.3.4. 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, you can use the **lvconvert** command to restore the mirror. This procedure is provided in Section 7.3, "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
  lvol1           100.00 lvol1_mimage_0(0),lvol1_mimage_1(0)
  [lvol1_mimage_0]      /dev/sda1(0)
  [lvol1_mimage_1]      /dev/sdb1(0)
  [lvol1_mlog]          /dev/sdd1(0)
# lvconvert -m 2 vg00/lvol1
```

```
  vg00/lvol1: Converted: 13.0%
  vg00/lvol1: Converted: 100.0%
  Logical volume lvol1 converted.
# lvs -a -o name,copy_percent,devices vg00
  LV              Copy%  Devices
  lvol1              100.00 lvol1_mimage_0(0),lvol1_mimage_1(0),lvol1_mimage_2(0)
  [lvol1_mimage_0]      /dev/sda1(0)
  [lvol1_mimage_1]      /dev/sdb1(0)
  [lvol1_mimage_2]      /dev/sdc1(0)
  [lvol1_mlog]          /dev/sdd1(0)
```

5.4.4. Creating Thinly-Provisioned Logical Volumes

As of the Red Hat Enterprise Linux 6.4 release, 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 **lvmthin**(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, you 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** that is 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 -V1G -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 -V1G -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** that is 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 **pool** 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 -L100M -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 (Pool_LV_size / Pool_LV_chunk_size * 64). You cannot currently resize the metadata volume, however, so if you expect significant growth of the size of thin pool at a later time you should increase this value 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.
```

5.4.5. Creating Snapshot Volumes

NOTE

As of the Red Hat Enterprise Linux 6.4 release, LVM supports thinly-provisioned snapshots. For information on creating thinly provisioned snapshot volumes, see Section 5.4.6, "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. As of the Red Hat Enterprise Linux 6.1 release, 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 5.7, "Activating Logical Volumes on Individual Nodes in a Cluster".

> **NOTE**
>
> As of the Red Hat Enterprise Linux 6.1 release, LVM snapshots are supported for mirrored logical volumes.
>
> As of the Red Hat Enterprise Linux 6.3 release, snapshots are supported for RAID logical volumes. For information on RAID logical volumes, see Section 5.4.16, "RAID Logical Volumes".

As of the Red Hat Enterprise Linux 6.5 release, 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 5.4.17, "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/lvol1**. 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/lvol1
```

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-OjsI-LT03-nHLC-y8XW-EhCl78
  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 for each snapshot volume. 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.

As of the Red Hat Enterprise Linux 6.2 release, there are two new features related to snapshots. First, 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. Second, 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.

As of the Red Hat Enterprise Linux 6.5 release, 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*.

5.4.6. Creating Thinly-Provisioned Snapshot Volumes

The Red Hat Enterprise Linux release 6.4 version of LVM provides support for thinly-provisioned snapshot volumes. For information on the benefits and limitations of thin snapshot volumes, see Section 3.3.7, "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. As of the Red Hat Enterprise Linux 6.5 release, thin snapshots can also be created 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. It is recommended that you use this option when creating a logical volume so that you can more easily see the volume you have created when you display logical volumes with the **lvs** command.

The following command creates a thinly-provisioned snapshot volume of the thinly-provisioned logical volume **vg001/thinvolume** that is named **mysnapshot1**.

```
# lvcreate -s --name mysnapshot1 vg001/thinvolume
  Logical volume "mysnapshot1" created
# lvs
  LV          VG      Attr    LSize  Pool       Origin     Data%  Move Log Copy%  Convert
  mysnapshot1 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 5.4.17, "Controlling Logical Volume Activation".

As of the Red Hat Enterprise Linux 6.5 release, you can create a thinly-provisioned snapshot of a non-thinly-provisioned logical volume. Since the non-thinly-provisioned logical volume is not contained within a thinpool, 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 shapshot 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
```

5.4.7. Creating LVM Cache Logical Volumes

As of the Red Hat Enterprise Linux 6.7 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 4G in size and that consists of **/dev/sde1**, the slow physical volume.

   ```
   # lvcreate -L 4G -n lv VG /dev/sde1
   ```

3. Create the cache data logical volume. This logical volume will hold data blocks from the origin volume. The size of this logical volume is the size of the cache and will be reported as the size of the cache pool logical volume. This example creates the cache data volume named **lv_cache**. It is 2G in size and is contained on the fast device **/dev/sdf1**, which is part of the volume group **VG**.

   ```
   # lvcreate -L 2G -n lv_cache VG /dev/sdf1
   ```

4. Create the cache metadata logical volume. This logical volume will hold cache pool metadata. The ratio of the size of the cache data logical volume to the size of the cache metadata logical volume should be about 1000:1, with a minimum size of 8MiB for the cache metadata logical volume. This example creates the cache metadata volume named **lv_cache_meta**. It is 12M in size and is also contained on the fast device **/dev/sdf1**, which is part of the volume group **VG**.

```
# lvcreate -L 12M -n lv_cache_meta VG /dev/sdf1
```

5. Create the cache pool logical volume by combining the cache data and the cache metadata logical volumes into a logical volume of type **cache-pool**. You can set the behavior of the cache pool in this step; in this example the **cachemode** argument is set to **writethrough**, which indicates that a write is considered complete only when it has been stored in both the cache pool logical volume and on the origin logical volume.

When you execute this command, the cache data logical volume is renamed with **_cdata** appended to the original name of the cache data logical volume, and the cache metadata logical volume is renamed with **_cmeta** appended to the original name of the cache data logical volume; both of these volumes become hidden.

```
# lvconvert --type cache-pool --cachemode writethrough --poolmetadata VG/lv_cache_meta
VG/lv_cache
  WARNING: Converting logical volume VG/lv_cache and VG/lv_cache_meta to pool's data
and metadata volumes.
  THIS WILL DESTROY CONTENT OF LOGICAL VOLUME (filesystem etc.)
  Converted VG/lv_cache to cache pool.
# lvs -a -o +devices
  LV              VG Attr     LSize   Pool Origin Data%  Meta% Cpy%Sync Devices
  lv              VG -wi-a-----   4.00g                           /dev/sde1(0)
  lv_cache        VG Cwi---C---   2.00g                           lv_cache_cdata(0)
  [lv_cache_cdata] VG  Cwi-------   2.00g                           /dev/sdf1(0)
  [lv_cache_cmeta] VG ewi-------  12.00m                            /dev/sdf1(512)
  [lvol0_pmspare]  VG ewi-------  12.00m                            /dev/sde1(1024)
```

6. Create the cache logical volume by combining the cache pool logical volume with the origin logical volume. The 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. You can execute this command when the origin logical volume is in use.

```
# lvconvert --type cache --cachepool VG/lv_cache VG/lv
  Logical volume VG/lv is now cached.
# lvs -a -o +devices
  LV              VG Attr     LSize  Pool      Origin       Data%  Meta% Cpy%Sync Devices
  lv              VG Cwi-a-C---   4.00g [lv_cache] [lv_corig] 0.02   2.31  0.00         lv_corig(0)
  [lv_corig]       VG  owi-aoC---   4.00g                                          /dev/sde1(0)
  [lv_cache]       VG Cwi---C---   2.00g                     0.02   2.31  0.00
lv_cache_cdata(0)
  [lv_cache_cdata] VG  Cwi-ao----   2.00g                                          /dev/sdf1(0)
  [lv_cache_cmeta] VG ewi-ao----  12.00m                                           /dev/sdf1(512)
  [lvol0_pmspare]   VG ewi-------  12.00m                                           /dev/sde1(1024)
```

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

5.4.8. Merging Snapshot Volumes

As of the Red Hat Enterprise Linux 6 release, 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.

5.4.9. Persistent Device Numbers

Major and minor device numbers are allocated dynamically at module load. Some applications work best if the block device always is 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.

5.4.10. 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 5.3.8, "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
```

5.4.11. 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
```

Renaming the root logical volume requires additional reconfiguration. For information on renaming a root volume, see How to rename root volume group or logical volume in Red Hat Enterprise Linux .

For more information on activating logical volumes on individual nodes in a cluster, see Section 5.7, "Activating Logical Volumes on Individual Nodes in a Cluster".

5.4.12. 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.

5.4.13. 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 5.8, "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 **lvol2** 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/lvol2
```

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
```

5.4.14. 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.

5.4.14.1. 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 135G 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
```

5.4.14.2. Extending a Mirrored Volume

As of the Red Hat Enterprise Linux 6.3 release, it is possible to grow mirrored logical volumes with the **lvextend** command without performing a synchronization of the new mirror regions.

If you specify the **--nosync** option when you create a mirrored logical volume with the **lvcreate** command, the mirror regions are not synchronized when the mirror is created, as described in Section 5.4.3, "Creating Mirrored Volumes". If you later extend a mirror that you have created with the **--nosync** option, the mirror 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 have an attribute bit 1 of "M" if it is a mirrored volume that was created without an initial synchronization, and it will have an attribute bit 1 of "m" if it was created with initial synchronization.

The following command displays the attributes of a mirrored logical volume named **lv** that was created without initial synchronization, showing attribute bit 1 as "M". Attribute bit 7 is "m", indicating a target type of **mirror**. For information on the meaning of the attribute bits, see Table 5.4, "lvs Display Fields".

```
# lvs vg
  LV   VG   Attr      LSize Pool Origin Snap%  Move Log     Copy%  Convert
  lv   vg   Mwi-a-m- 5.00g                     lv_mlog 100.00
```

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

If you created a mirrored 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 logical volume that was created without the **--nosync** option, indicated that the mirror was synchronized when it was created. This example, however, specifies that the mirror not be synchronized when the volume is extended. Note that the volume has an attribute of "m", but after executing the **lvextend** commmand with the **--nosync** option the volume has an attribute of "M".

```
# lvs vg
  LV   VG   Attr      LSize  Pool Origin Snap%  Move Log     Copy%  Convert
  lv   vg   mwi-a-m- 20.00m                     lv_mlog 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     Copy%  Convert
  lv  vg   Mwi-a-m- 5.02g                  lv_mlog 100.00
```

If a mirror is inactive, it will not automatically skip synchronization when you extend the mirror, even if you create the mirror 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 mirror is performing recovery, you cannot extend the mirrored 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 mirror while it is recovering.

5.4.14.3. 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 and 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]1
  PV        VG   Fmt  Attr PSize   PFree   PV Tags
  /dev/sdb1  taft lvm2 a-   135.66g 135.66g A
  /dev/sdc1  taft lvm2 a-   135.66g 135.66g B
  /dev/sdd1  taft lvm2 a-   135.66g 135.66g B
  /dev/sde1  taft lvm2 a-   135.66g 135.66g C
  /dev/sdf1  taft lvm2 a-   135.66g 135.66g C
  /dev/sdg1  taft lvm2 a-   135.66g 135.66g A
  /dev/sdh1  taft lvm2 a-   135.66g 135.66g A
```

The following command creates a 100GB mirrored volume from the volume group **taft**.

```
# lvcreate -m 1 -n mirror --nosync -L 100G taft
```

The following command shows which devices are used for the mirror legs and mirror log.

```
# lvs -a -o +devices
  LV               VG      Attr  LSize  Log        Copy%  Devices
  mirror           taft    Mwi-a- 100.00g mirror_mlog 100.00
mirror_mimage_0(0),mirror_mimage_1(0)
  [mirror_mimage_0] taft   iwi-ao 100.00g                 /dev/sdb1(0)
  [mirror_mimage_1] taft   iwi-ao 100.00g                 /dev/sdc1(0)
  [mirror_mlog]    taft    lwi-ao  4.00m               /dev/sdh1(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 +100G taft/mirror
  Extending 2 mirror images.
  Extending logical volume mirror to 200.00 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        Copy%  Devices
  mirror           taft    Mwi-a- 200.00g mirror_mlog 50.16
mirror_mimage_0(0),mirror_mimage_1(0)
  [mirror_mimage_0] taft   lwi-ao 200.00g                 /dev/sdb1(0)
  [mirror_mimage_0] taft   lwi-ao 200.00g                 /dev/sdg1(0)
  [mirror_mimage_1] taft   lwi-ao 200.00g                 /dev/sdc1(0)
  [mirror_mimage_1] taft   lwi-ao 200.00g                 /dev/sdd1(0)
  [mirror_mlog]    taft    lwi-ao  4.00m               /dev/sdh1(0)
```

5.4.15. 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 voume. If shrinking the file sytem 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 **lvol1** in volume group **vg00** to be 64 megabytes. In
this example, **lvol1** 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/lvol1
fsck from util-linux 2.23.2
/dev/mapper/vg00-lvol1: clean, 11/25688 files, 8896/102400 blocks
resize2fs 1.42.9 (28-Dec-2013)
Resizing the filesystem on /dev/mapper/vg00-lvol1 to 65536 (1k) blocks.
The filesystem on /dev/mapper/vg00-lvol1 is now 65536 blocks long.

  Size of logical volume vg00/lvol1 changed from 100.00 MiB (25 extents) to 64.00 MiB (16 extents).
  Logical volume vg00/lvol1 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/lvol1
```

5.4.16. RAID Logical Volumes

As of the Red Hat Enterprise Linux 6.3 release, LVM supports RAID4/5/6 and a new implementation of
mirroring. The latest implementation of mirroring differs from the previous implementation of mirroring
(documented in Section 5.4.3, "Creating Mirrored Volumes") in the following ways:

- The segment type for the new implementation of mirroring is **raid1**. For the earlier
  implementation, the segment type is **mirror**.

- The new implementation of mirroring leverages MD software RAID, just as for the RAID 4/5/6
  implementations.

- The new implementation of mirroring maintains a fully redundant bitmap area for each mirror
  image, which increases its fault handling capabilities. This means that there is no **--mirrorlog**
  option or **--corelog** option for mirrors created with this segment type.

- The new implementation of mirroring can handle transient failures.

- Mirror images can be temporarily split from the array and merged back into the array later.

- The new implementation of mirroring supports snapshots (as do the higher-level RAID implementations).

- The new RAID implementations are not cluster-aware. You cannot create an LVM RAID logical volume in a clustered volume group.

For information on how failures are handled by the RAID logical volumes, see Section 5.4.16.8, "Setting a RAID fault policy".

The remainder of this section describes the following administrative tasks you can perform on LVM RAID devices:

- Section 5.4.16.1, "Creating a RAID Logical Volume"

- Section 5.4.16.2, "Converting a Linear Device to a RAID Device"

- Section 5.4.16.3, "Converting an LVM RAID1 Logical Volume to an LVM Linear Logical Volume"

- Section 5.4.16.4, "Converting a Mirrored LVM Device to a RAID1 Device"

- Section 5.4.16.5, "Changing the Number of Images in an Existing RAID1 Device"

- Section 5.4.16.6, "Splitting off a RAID Image as a Separate Logical Volume"

- Section 5.4.16.7, "Splitting and Merging a RAID Image"

- Section 5.4.16.8, "Setting a RAID fault policy"

- Section 5.4.16.9, "Replacing a RAID device"

- Section 5.4.16.10, "Scrubbing a RAID Logical Volume"

- Section 5.4.16.11, "Controlling I/O Operations on a RAID1 Logical Volume"

5.4.16.1. Creating a RAID Logical Volume

To create a RAID logical volume, you specify a raid type as the **--type** argument of the **lvcreate** command. Usually when you create a logical volume with the **lvcreate** command, the **--type** argument is implicit. For example, when you specify the **-i** *stripes* argument, the **lvcreate** command assumes the **--type stripe** option. When you specify the **-m** *mirrors* argument, the **lvcreate** command assumes the **--type mirror** option. When you create a RAID logical volume, however, you must explicitly specify the segment type you desire. The possible RAID segment types are described in Table 5.1, "RAID Segment Types".

Table 5.1. RAID Segment Types

Segment type	Description
raid1	RAID1 mirroring
raid4	RAID4 dedicated parity disk
raid5	Same as **raid5_ls**

Segment type	Description
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
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 (Red Hat Enterprise Linux 6.4 and later	Striped mirrors Striping of mirror sets

For most users, specifying one of the five available primary types (**raid1**, **raid4**, **raid5**, **raid6**, **raid10**) should be sufficient. For more information on the different algorithms used by RAID 5/6, see chapter four of the *Common RAID Disk Data Format Specification* at

http://www.snia.org/sites/default/files/SNIA_DDF_Technical_Position_v2.0.pdf.

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 1G 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. Although the **-m** argument is the same argument used to specify the number of copies for the previous mirror implementation, in this case you override the default segment type **mirror** by explicitly setting the segment type as **raid1**. Similarly, you specify the number of stripes for a RAID 4/5/6 logical volume with the familiar **-i argument**, overriding the default segment type with the desired RAID type. You can also specify the stripe size with the **-I** argument.

> **NOTE**
>
> You can set the default mirror segment type to **raid1** by changing **mirror_segtype_default** in the **lvm.conf** file.

The following command creates a RAID5 array (3 stripes + 1 implicit parity drive) named **my_lv** in the volume group **my_vg** that is 1G 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 1G 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.

As of Red Hat Enterprise Linux 6.5, 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 10G is 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 5.4.16.10, "Scrubbing a RAID Logical Volume" .

5.4.16.2. 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.

5.4.16.3. 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         /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        /dev/sdb1(1)
```

5.4.16.4. Converting a Mirrored LVM Device to a RAID1 Device

You can convert an existing mirrored LVM device 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)
```

5.4.16.5. 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, by using 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 5.4.3.4, "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 display shows 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)
```

5.4.16.6. 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 5.4.3.2, "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 images from an existing RAID1 logical volume (as described in Section 5.4.16.5, "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)
```

5.4.16.7. 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 -m2 -L1G -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
```

5.4.16.8. 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.

5.4.16.8.1. 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-H20O-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-H20O-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, as described in Section 5.4.16.8.2, "The warn RAID Fault Policy". Alternately, you can replace the failed device, as described in Section 5.4.16.9, "Replacing a RAID device".

5.4.16.8.2. 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 fbI0YO-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 fbI0YO-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, as of Red Hat Enterprise Linux release 6.5 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
```

5.4.16.9. 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 -m2 -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 -m1 -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.

5.4.16.10. Scrubbing a RAID Logical Volume

As of the Red Hat Enterprise Linux 6.5 release, 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** display 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 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 5.8.2, "Object Selection".

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.

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

As of the Red Hat Enterprise Linux release 6.5, 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. The default behavior is to set the **write-mostly** attribute 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.

5.4.17. 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.

5.5. 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 that are 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 best 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.

5.6. 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 move as a percentage at five second intervals.

> # pvmove -i5 /dev/sdd1

5.7. 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* .

5.8. CUSTOMIZED REPORTING FOR LVM

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.

The following sections provide:

- A summary of command arguments you can use to control the format of the generated report.

- A list of the fields you can select for each LVM object.

- A summary of command arguments you can use to sort the generated report.

- Instructions for specifying the units of the report output.

5.8.1. 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-XO8M-
  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 7.2, "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
```

5.8.2. Object Selection

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 5.2, "pvs 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 5.2. pvs 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
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-XO8M-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 5.3, "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 5.3. 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)
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 **vgs** 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 **vgs** 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-OpMO-0118-nlwO-wwqd-fD5D32
```

The lvs Command

Table 5.4, "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.

Table 5.4. 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_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

Argument	Header	Description
		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.
		Bit 4: fixed (m)inor
		Bit 5: State: (a)ctive, (s)uspended, (I)nvalid 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)ritemostly. (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)ritemostly 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)

Argument	Header	Description
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)
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

Argument	Header	Description
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-OjsI-LT03-
nHLC-y8XW-EhCl78
  newvgsnap1 new_vg    1 swi-a-  8.00M  -1  -1 253  5    lvol0    0.20                          1ye1OU-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
```

5.8.3. 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
```

5.8.4. 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
```

CHAPTER 6. LVM CONFIGURATION EXAMPLES

This chapter provides some basic LVM configuration examples.

6.1. CREATING AN LVM LOGICAL VOLUME ON THREE DISKS

This example creates an LVM logical volume called **new_logical_volume** that consists of the disks at **/dev/sda1**, **/dev/sdb1**, and **/dev/sdc1**.

6.1.1. Creating the Physical Volumes

To use disks in a volume group, you label them as LVM physical volumes.

> **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
```

6.1.2. Creating the Volume Group

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
```

6.1.3. Creating the Logical Volume

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 2GB of the volume group.

```
# lvcreate -L2G -n new_logical_volume new_vol_group
  Logical volume "new_logical_volume" created
```

6.1.4. Creating the File System

The following command creates a GFS2 file system on the logical volume.

```
# mkfs.gfs2 -plock_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
[root@tng3-1 ~]# df
Filesystem          1K-blocks      Used Available Use% Mounted on
/dev/new_vol_group/new_logical_volume
                    1965840        20  1965820   1% /mnt
```

6.2. CREATING A STRIPED LOGICAL VOLUME

This example creates an LVM striped logical volume called **striped_logical_volume** that stripes data across the disks at **/dev/sda1**, **/dev/sdb1**, and **/dev/sdc1**.

6.2.1. Creating the Physical Volumes

Label the disks you will use in the volume groups as LVM physical volumes.

> **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
```

6.2.2. Creating the Volume Group

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
```

6.2.3. Creating the Logical Volume

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 -i3 -I4 -L2G -nstriped_logical_volume volgroup01
  Rounding size (512 extents) up to stripe boundary size (513 extents)
  Logical volume "striped_logical_volume" created
```

6.2.4. Creating the File System

The following command creates a GFS2 file system on the logical volume.

```
# mkfs.gfs2 -plock_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
[root@tng3-1 ~]# 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
```

6.3. SPLITTING A VOLUME GROUP

In this example, 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**.

6.3.1. Determining Free Space

You can 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   ]
```

6.3.2. Moving the Data

You can 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   ]
```

6.3.3. Splitting the Volume Group

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.

You can 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
```

6.3.4. Creating the New Logical Volume

After creating the new volume group, you can create the new logical volume **yourlv**.

```
# lvcreate -L5G -n yourlv yourvg
  Logical volume "yourlv" created
```

6.3.5. Making a File System and Mounting the New Logical Volume

You can make a file system on the new logical volume and mount it.

```
#  mkfs.gfs2 -plock_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

[root@tng3-1 ~]# mount /dev/yourvg/yourlv /mnt
```

6.3.6. Activating and Mounting the Original Logical Volume

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

[root@tng3-1 ~]# mount /dev/myvg/mylv /mnt
[root@tng3-1 ~]# 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
```

6.4. REMOVING A DISK FROM A LOGICAL VOLUME

This example shows 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.

6.4.1. 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
```

We want to move the extents off of **/dev/sdb1** so that we can remove it from the volume group.

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
```

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"
[root@tng3-1 ~]# 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.

6.4.2. 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
```

We want to move the extents of **/dev/sdb1** to a new device, **/dev/sdd1**.

6.4.2.1. Creating the New Physical Volume

Create a new physical volume from **/dev/sdd1**.

```
# pvcreate /dev/sdd1
  Physical volume "/dev/sdd1" successfully created
```

6.4.2.2. Adding the New Physical Volume to the Volume Group

Add **/dev/sdd1** to the existing volume group **myvg**.

```
# vgextend myvg /dev/sdd1
  Volume group "myvg" successfully extended
[root@tng3-1]# 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
```

6.4.2.3. Moving the Data

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%

[root@tng3-1]# 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
```

6.4.2.4. Removing the Old Physical Volume from the Volume Group

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.

6.5. 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. 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, either directly or by means of the **lvmconf** command as described in Section 4.1, "Creating LVM Volumes in a Cluster".

The following procedure creates a mirrored LVM volume in a cluster. First the procedure checks to see whether the cluster services are installed and running, then the procedure creates the mirrored volume.

1. 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
   ```

   > **NOTE**
   >
   > As of Red Hat Enterprise Linux 6.7, the **lvmconf** command provides a **--services** option that will also enable the services required for LVM in a cluster, a **--mirrorservice** option that enables the **cmirrord** service, and a **--startstopservices** option that immediately starts or stops the services that have been enabled. For information on the **lvmconf** command, see the **lvmconf** man page.

2. To create a clustered logical volume, the cluster infrastructure must be up and running on every node in the cluster. The following example verifies that the **clvmd** daemon is running on the node from which it was issued:

   ```
    ps auxw | grep clvmd
   root    17642  0.0  0.1 32164 1072 ?        Ssl  Apr06   0:00 clvmd -T20 -t 90
   ```

   The following command shows the local view of the cluster status:

   ```
   # cman_tool services
   fence domain
   member count  3
   victim count  0
   victim now    0
   master nodeid 2
   wait state    none
   members       1 2 3

   dlm lockspaces
   name          clvmd
   ```

```
id        0x4104eefa
flags     0x00000000
change    member 3 joined 1 remove 0 failed 0 seq 1,1
members   1 2 3
```

3. Ensure that the **cmirror** package is installed.

4. Start the **cmirrord** service.

```
# service cmirrord start
Starting cmirrord:                              [ OK ]
```

5. 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
[root@doc-07 ~]# pvcreate /dev/sdc1
  Physical volume "/dev/sdc1" successfully created
[root@doc-07 ~]# pvcreate /dev/sdd1
  Physical volume "/dev/sdd1" successfully created
```

6. 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
```

7. 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 -l 1000 -m1 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
[root@doc-07 log]# lvs vg001/mirrorlv
LV       VG      Attr   LSize Origin Snap%  Move Log         Copy%  Convert
```

```
  mirrorlv vg001    mwi-a- 3.91G                      vg001_mlog    91.00
[root@doc-07 ~]#  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
```

8. You can use the **lvs** 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 7.3, "Recovering from LVM Mirror Failure" .

CHAPTER 7. LVM TROUBLESHOOTING

This chapter provide instructions for troubleshooting a variety of LVM issues.

7.1. 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.

7.2. 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)
```

7.3. 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**, requiring that you manually rebuild the mirror. For information on setting the **mirror_log_fault_policy** parameter, see Section 5.4.3.1, "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/sdg1" successfully created
Physical volume "/dev/sdg2" 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
[root@link-08 ~]# 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)

[root@link-08 ~]# 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
  /dev/sda2: 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

[root@link-08 ~]# 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          lvm2 [603.94 GB]
  PV /dev/sdi2          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)
```

7.4. 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 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
[root@link-07 backup]# 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.

7.5. 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 7.4, "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 **pvcreate** 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.

7.6. 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.

It is recommended that you run the **vgreduce** command with the **--test** argument to verify what you will be destroying.

Like most LVM operations, the **vgreduce** command is reversible in a sense 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.

7.7. 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 some 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 -l8780 -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 5.4.1, "Creating Linear Logical Volumes".

7.8. 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 GDjTZf7Y03GJHjteqOwrye2dcSCjdaUi: using /dev/dm-5 not /dev/sdd
> Found duplicate PV GDjTZf7Y03GJHjteqOwrye2dcSCjdaUi: using /dev/emcpowerb not /dev/sde
> Found duplicate PV GDjTZf7Y03GJHjteqOwrye2dcSCjdaUi: using /dev/sddlmab 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

7.8.1. 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/sddlmab** 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.

7.8.2. 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 GDjTZf7Y03GJHjteqOwrye2dcSCjdaUi: 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 5.5, "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.

7.8.3. 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 GDjTZf7Y03GJHjteqOwrye2dcSCjdaUi: using **/dev/mapper/mpatha** not **/dev/mapper/mpathc**

> Found duplicate PV GDjTZf7Y03GJHjteqOwrye2dcSCjdaUi: 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.

CHAPTER 8. LVM ADMINISTRATION WITH THE LVM GUI

In addition to the Command Line Interface (CLI), LVM provides a Graphical User Interface (GUI) which you can use to configure LVM logical volumes. You can open this utility by typing **system-config-lvm**. The LVM chapter of the *Storage Administration Guide* provides step-by-step instructions for configuring an LVM logical volume using this utility.

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. So, 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 these mapping formats:

- linear

- striped

- mirror

- snapshot and snapshot-origin

- error

- zero

- multipath

- crypt

- device-mapper RAID

- thin

- thin-pool

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 16384000 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 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.

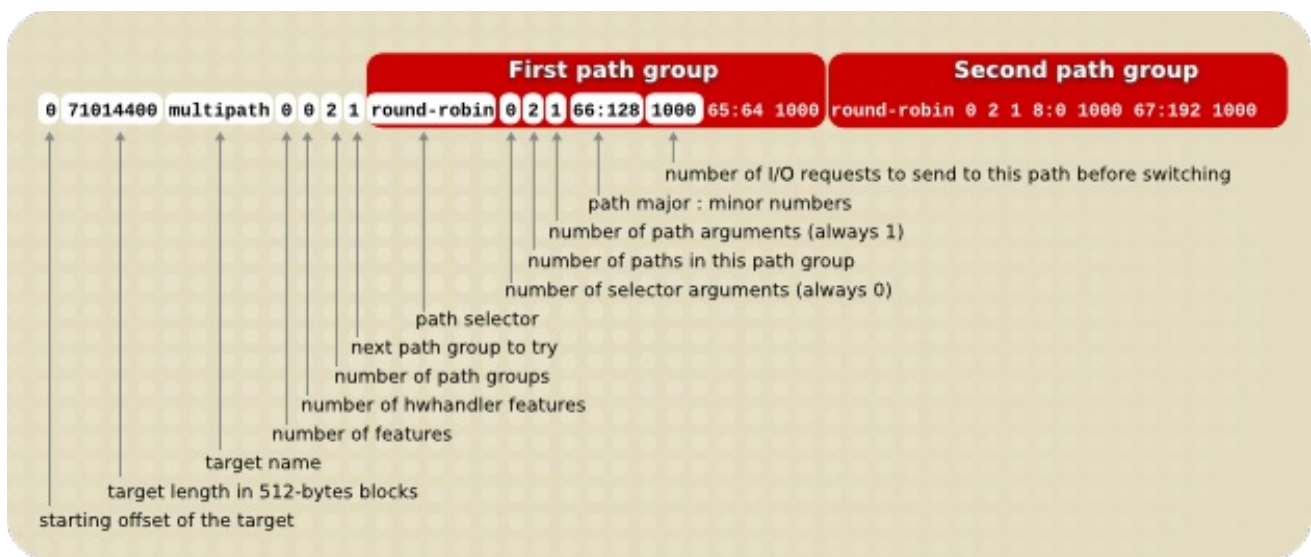Figure A.1, "Multipath Mapping Target" shows the format of a multipath target with two path groups.

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 out 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 _Using Device Mapper Multipath_ document.

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.1.9. The device-mapper RAID Mapping Target

The device-mapper RAID (dm-raid) target provides a bridge from DM to MD. It allows the MD RAID drivers to be accessed using a device-mapper interface. The format of the dm-raid target is as follows

> *start length* raid *raid_type #raid_params raid_params #raid_devs metadata_dev0 dev0* [.. *metadata_devN devN*]

start

starting block in virtual device

length

length of this segment

raid_type

The RAID type can be one of the following

raid1

RAID1 mirroring

raid4

RAID4 dedicated parity disk

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_zr

RAID6 zero restart

— rotating parity 0 (left to right) with data restart

raid6_nr

RAID6 N restart

— rotating parity N (right to left) with data restart

raid6_nc

RAID6 N continue

— rotating parity N (right to left) with data continuation

raid10

Various RAID10-inspired algorithms selected by further optional arguments

— RAID 10: Striped mirrors (striping on top of mirrors)

— RAID 1E: Integrated adjacent striped mirroring

— RAID 1E: Integrated offset striped mirroring

— Other similar RAID10 variants

#raid_params

The number of parameters that follow

raid_params

Mandatory parameters:

chunk_size

Chunk size in sectors. This parameter is often known as "stripe size". It is the only mandatory parameter and is placed first.

Followed by optional parameters (in any order):

[sync|nosync]

Force or prevent RAID initialization.

rebuild *idx*

Rebuild drive number *idx* (first drive is 0).

daemon_sleep *ms*

Interval between runs of the bitmap daemon that clear bits. A longer interval means less bitmap I/O but resyncing after a failure is likely to take longer.

min_recovery_rate *KB/sec/disk*

Throttle RAID initialization

max_recovery_rate *KB/sec/disk*

Throttle RAID initialization

write_mostly *idx*

Mark drive index *idx* write-mostly.

max_write_behind *sectors*

See the description of **--write-behind** in the **mdadm** man page.

stripe_cache *sectors*

> Stripe cache size (RAID 4/5/6 only)

region_size *sectors*

> The **region_size** multiplied by the number of regions is the logical size of the array. The bitmap records the device synchronization state for each region.

raid10_copies *#copies*

> The number of RAID10 copies. This parameter is used in conjunction with the **raid10_format** parameter to alter the default layout of a RAID10 configuration. The default value is 2.

raid10_format near|far|offset

> This parameter is used in conjunction with the **raid10_copies** parameter to alter the default layout of a RAID10 configuration. The default value is **near**, which specifies a standard mirroring layout.
>
> If the **raid10_copies** and **raid10_format** are left unspecified, or **raid10_copies 2** and/or **raid10_format near** is specified, then the layouts for 2, 3 and 4 devices are as follows:

```
2 drives    3 drives      4 drives
--------   ----------   --------------
A1  A1     A1  A1  A2    A1  A1  A2  A2
A2  A2     A2  A3  A3    A3  A3  A4  A4
A3  A3     A4  A4  A5    A5  A5  A6  A6
A4  A4     A5  A6  A6    A7  A7  A8  A8
..  ..     ..  ..  ..    ..  ..  ..  ..
```

> The 2-device layout is equivalent to 2-way RAID1. The 4-device layout is what a traditional RAID10 would look like. The 3-device layout is what might be called a 'RAID1E - Integrated Adjacent Stripe Mirroring'.
>
> If **raid10_copies 2** and **raid10_format far** are specified, then the layouts for 2, 3 and 4 devices are as follows:

```
2 drives    3 drives       4 drives
--------   -----------    ------------------
A1  A2     A1  A2  A3     A1  A2  A3  A4
A3  A4     A4  A5  A6     A5  A6  A7  A8
A5  A6     A7  A8  A9     A9  A10 A11 A12
..  ..     ..  ..  ..     ..  ..  ..  ..
A2  A1     A3  A1  A2     A2  A1  A4  A3
A4  A3     A6  A4  A5     A6  A5  A8  A7
A6  A5     A9  A7  A8     A10 A9  A12 A11
..  ..     ..  ..  ..     ..  ..  ..  ..
```

> If **raid10_copies 2** and **raid10_format offset** are specified, then the layouts for 2, 3 and 4 devices are as follows:

```
2 drives    3 drives       4 drives
--------   --------      ------------------
A1  A2     A1  A2  A3     A1  A2  A3  A4
A2  A1     A3  A1  A2     A2  A1  A4  A3
A3  A4     A4  A5  A6     A5  A6  A7  A8
```

```
A4  A3     A6  A4  A5     A6   A5   A8   A7
A5  A6     A7  A8  A9     A9   A10  A11  A12
A6  A5     A9  A7  A8     A10  A9   A12  A11
..  ..     ..  ..  ..     ..   ..   ..   ..
```

These layouts closely resemble the layouts fo RAID1E - Integrated Offset Stripe Mirroring'

#raid_devs

The number of devices composing the array

Each device consists of two entries. The first is the device containing the metadata (if any); the second is the one containing the data.

If a drive has failed or is missing at creation time, a '-' can be given for both the metadata and data drives for a given position.

The following example shows a RAID4 target with a starting block of 0 and a segment length of 1960893648. There are 4 data drives, 1 parity, with no metadata devices specified to hold superblock/bitmap info and a chunk size of 1MiB

```
0 1960893648 raid raid4 1 2048 5 - 8:17 - 8:33 - 8:49 - 8:65 - 8:81
```

The following example shows a RAID4 target with a starting block of 0 and a segment length of 1960893648. there are 4 data drives, 1 parity, with metadata devices, a chunk size of 1MiB, force RAID initialization, and a **min_recovery** rate of 20 kiB/sec/disks.

```
0 1960893648 raid raid4 4 2048 sync min_recovery_rate 20 5 8:17 8:18 8:33 8:34 8:49 8:50 8:65 8:66
8:81 8:82
```

A.1.10. The thin and thin-pool Mapping Targets

The format of a thin-pool target is as follows:

start length thin-pool *metadata_dev data_dev data_block_size low_water_mark* [*#feature_args* [*arg**]]

start

starting block in virtual device

length

length of this segment

metadata_dev

The metadata device

data_dev

The data device

data_block_size

The data block size (in sectors). The data block size gives the smallest unit of disk space that can be allocated at a time expressed in units of 512-byte sectors. Data block size must be between 64KB (128 sectors) and 1GB (2097152 sectors) inclusive and it must be a mutlipole of 128 (64KB).

low_water_mark

The low water mark, expressed in blocks of size ***data_block_size***. If free space on the data device drops below this level then a device-mapper event will be triggered which a user-space daemon should catch allowing it to extend the pool device. Only one such event will be sent. Resuming a device with a new table itself triggers an event so the user-space daemon can use this to detect a situation where a new table already exceeds the threshold.

A low water mark for the metadata device is maintained in the kernel and will trigger a device-mapper event if free space on the metadata device drops below it.

#feature_args

The number of feature arguments

arg

The thin pool feature argument are as follows:

skip_block_zeroing

Skip the zeroing of newly-provisioned blocks.

ignore_discard

Disable discard support.

no_discard_passdown

Do not pass discards down to the underlying data device, but just remove the mapping.

read_only

Do not allow any changes to be made to the pool metadata.

error_if_no_space

Error IOs, instead of queuing, if no space.

The following example shows a thin-pool target with a starting block in the virtual device of 0, a segment length of 1638400. **/dev/sdc1** is a small metadata device and **/dev/sdc2** is a larger data device. The chunksize is 64k, the low_water_mark is 0, and there are no features.

```
0 16384000 thin-pool /dev/sdc1 /dev/sdc2 128 0 0
```

The format of a thin target is as follows:

```
start length thin pool_dev dev_id [external_origin_dev]
```

start

starting block in virtual device

length

length of this segment

pool_dev

The thin-pool device, for example **/dev/mapper/my_pool** or 253:0

dev_id

The internal device identifier of the device to be activated.

external_origin_dev

An optional block device outside the pool to be treated as a read-only snapshot origin. Reads to unprovisioned areas of the thin target will be mapped to this device.

The following example shows a 1 GiB thinLV that uses **/dev/mapper/pool** as its backing store (thin-pool). The target has a starting block in the virtual device of 0 and a segment length of 2097152.

```
0 2097152 thin /dev/mapper/pool 1
```

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 the user 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 fragments 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-K528WUGQgPadNXYcFrrf9LnPIUMswgkCkpgPIgYzSvigM7SfeWCypddNSWtNzc2N

...
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-tOcS1kqFV9drb0X1Vr8sxeYP0tqcrpdegyqj5lZxe45JMGlmvtqLmbLpBcenh2L3
```

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**, 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 devices 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-testgfslv3: 0 312352768 linear
testgfsvg-testgfslv2: 0 312352768 linear
testgfsvg-testgfslv1: 0 312352768 linear
testgfsvg-testgfslv1: 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-testgfslv3: 1 dependencies   : (8, 16)
testgfsvg-testgfslv2: 1 dependencies   : (8, 16)
testgfsvg-testgfslv1: 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 the user can name. This provides users 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, the users 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. A user 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

In Red Hat Enterprise Linux 6, 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.

The Red Hat Enterprise Linux 6 release 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 basic/general Device Mapper rules and creates the symlinks in **/dev/mapper** with a **/dev/dm-*N*** target where N is a number assigned dynamically to a device by the kernel (**/dev/dm-*N*** is a node)
	NOTE: **/dev/dm-*N*** nodes should $_{never}$ 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 **/dev/mapper** 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 **/dev/*vgname*** directory with a **/dev/dm-*N*** 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.

Filename	Description
13-dm-disk.rules	Contains rules to be applied for all Device Mapper devices in general and creates symlinks in the **/dev/disk/by-id**, **/dev/disk/by-uuid** and the **/dev/disk/by-uuid** 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 **lvmeetad** 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
dmsetup udevcomplete	Used to notify that udev has completed processing the rules and unlocks waiting process (called from within **udev** rules in **95-dm-notify.rules**).
dmsetup udevcomplete_all	Used for debugging purposes to manually unlock all waiting processes.
dmsetup udevcookies	Used for debugging purposes, to show all existing cookies (system-wide semaphores).
dmsetup udevcreatecookie	Used to create a cookie (semaphore) manually. This is useful to run more processes under one synchronization resource.
dmsetup udevreleasecookie	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** 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".

LVM profiles

An LVM profile is a set of selected customizable configuration settings that can be implemented for specific environments. The settings in an LVM profile can be used to override existing configuration. For information on LVM profiles see Section B.3, "LVM Profiles".

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. The **lvmconfig** command displays the LVM configuration information after loading the **/etc/lvm/lvm.conf** file and any other configuration files.

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 profilable configuration settings, 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 profilable settings 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.
#
# 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.

# This section allows you to set the way the configuration settings are handled.
```

```
config {

    # If enabled, any LVM2 configuration mismatch is reported.
    # This implies checking that the configuration key is understood
    # by LVM2 and that the value of the key is of a proper type.
    # If disabled, any configuration mismatch is ignored and default
    # value is used instead without any warning (a message about the
    # configuration key not being found is issued in verbose mode only).
    checks = 1

    # If enabled, any configuration mismatch aborts the LVM2 process.
    abort_on_errors = 0

    # Directory where LVM looks for configuration profiles.
    profile_dir = "/etc/lvm/profile"
}

# This section allows you to configure which block devices should
# be used by the LVM system.
devices {

    # Where do you want your volume groups to appear ?
    dir = "/dev"

    # An array of directories that contain the device nodes you wish
    # to use with LVM2.
    scan = [ "/dev" ]

    # If set, the cache of block device nodes with all associated symlinks
    # will be constructed out of the existing udev database content.
    # This avoids using and opening any inapplicable non-block devices or
    # subdirectories found in the device directory. This setting is applied
    # to udev-managed device directory only, other directories will be scanned
    # fully. LVM2 needs to be compiled with udev support for this setting to
    # take effect. N.B. Any device node or symlink not managed by udev in
    # udev directory will be ignored with this setting on.
    obtain_device_list_from_udev = 1

    # If several entries in the scanned directories correspond to the
    # same block device and the tools need to display a name for device,
    # all the pathnames are matched against each item in the following
    # list of regular expressions in turn and the first match is used.
    # preferred_names = [ ]

    # Try to avoid using undescriptive /dev/dm-N names, if present.
    preferred_names = [ "^/dev/mpath/", "^/dev/mapper/mpath", "^/dev/[hs]d" ]

    # A filter that tells LVM2 to only use a restricted set of devices.
    # The filter consists of an array of regular expressions.  These
    # expressions can be delimited by a character of your choice, and
    # prefixed with either an 'a' (for accept) or 'r' (for reject).
    # The first expression found to match a device name determines if
    # the device will be accepted or rejected (ignored).  Devices that
    # don't match any patterns are accepted.

    # Be careful if there there are symbolic links or multiple filesystem
```

```
# entries for the same device as each name is checked separately against
# the list of patterns.  The effect is that if the first pattern in the
# list to match a name is an 'a' pattern for any of the names, the device
# is accepted; otherwise if the first pattern in the list to match a name
# is an 'r' pattern for any of the names it is rejected; otherwise it is
# accepted.

# Don't have more than one filter line active at once: only one gets used.

# Run vgscan after you change this parameter to ensure that
# the cache file gets regenerated (see below).
# If it doesn't do what you expect, check the output of 'vgscan -vvvv'.


# By default we accept every block device:
filter = [ "a/.*/" ]

# Exclude the cdrom drive
# filter = [ "r|/dev/cdrom|" ]

# When testing I like to work with just loopback devices:
# filter = [ "a/loop/", "r/.*/" ]

# Or maybe all loops and ide drives except hdc:
# filter =[ "a|loop|", "r|/dev/hdc|", "a|/dev/ide|", "r|.*|" ]

# Use anchors if you want to be really specific
# filter = [ "a|^/dev/hda8$|", "r/.*/" ]

# Since "filter" is often overridden from command line, it is not suitable
# for system-wide device filtering (udev rules, lvmetad). To hide devices
# from LVM-specific udev processing and/or from lvmetad, you need to set
# global_filter. The syntax is the same as for normal "filter"
# above. Devices that fail the global_filter are not even opened by LVM.

# global_filter = []

# The results of the filtering are cached on disk to avoid
# rescanning dud devices (which can take a very long time).
# By default this cache is stored in the /etc/lvm/cache directory
# in a file called '.cache'.
# It is safe to delete the contents: the tools regenerate it.
# (The old setting 'cache' is still respected if neither of
# these new ones is present.)
# N.B. If obtain_device_list_from_udev is set to 1 the list of
# devices is instead obtained from udev and any existing .cache
# file is removed.
cache_dir = "/etc/lvm/cache"
cache_file_prefix = ""

# You can turn off writing this cache file by setting this to 0.
write_cache_state = 1

# Advanced settings.

# List of pairs of additional acceptable block device types found
```

```
# in /proc/devices with maximum (non-zero) number of partitions.
# types = [ "fd", 16 ]

# If sysfs is mounted (2.6 kernels) restrict device scanning to
# the block devices it believes are valid.
# 1 enables; 0 disables.
sysfs_scan = 1

# By default, LVM2 will ignore devices used as component paths
# of device-mapper multipath devices.
# 1 enables; 0 disables.
multipath_component_detection = 1

# By default, LVM2 will ignore devices used as components of
# software RAID (md) devices by looking for md superblocks.
# 1 enables; 0 disables.
md_component_detection = 1

# By default, if a PV is placed directly upon an md device, LVM2
# will align its data blocks with the md device's stripe-width.
# 1 enables; 0 disables.
md_chunk_alignment = 1

# Default alignment of the start of a data area in MB.  If set to 0,
# a value of 64KB will be used.  Set to 1 for 1MiB, 2 for 2MiB, etc.
# default_data_alignment = 1

# By default, the start of a PV's data area will be a multiple of
# the 'minimum_io_size' or 'optimal_io_size' exposed in sysfs.
# - minimum_io_size - the smallest request the device can perform
#   w/o incurring a read-modify-write penalty (e.g. MD's chunk size)
# - optimal_io_size - the device's preferred unit of receiving I/O
#   (e.g. MD's 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.
# 1 enables; 0 disables.
data_alignment_detection = 1

# Alignment (in KB) of start of data area when creating a new PV.
# md_chunk_alignment and data_alignment_detection are disabled if set.
# Set to 0 for the default alignment (see: data_alignment_default)
# or page size, if larger.
data_alignment = 0

# By default, the start of the PV's aligned data area will be shifted by
# the 'alignment_offset' exposed in sysfs.  This offset is often 0 but
# may be non-zero; e.g.: certain 4KB sector drives that compensate for
# windows partitioning will have an alignment_offset of 3584 bytes
# (sector 7 is the lowest aligned logical block, the 4KB sectors start
# at LBA -1, and consequently sector 63 is aligned on a 4KB boundary).
# But note that pvcreate --dataalignmentoffset will skip this detection.
# 1 enables; 0 disables.
data_alignment_offset_detection = 1

# If, while scanning the system for PVs, LVM2 encounters a device-mapper
```

```
  # device that has its I/O suspended, it waits for it to become accessible.
  # Set this to 1 to skip such devices.  This should only be needed
  # in recovery situations.
  ignore_suspended_devices = 0

  # During each LVM operation errors received from each device are counted.
  # If the counter of a particular device exceeds the limit set here, no
  # further I/O is sent to that device for the remainder of the respective
  # operation. Setting the parameter to 0 disables the counters altogether.
  disable_after_error_count = 0

  # Allow use of pvcreate --uuid without requiring --restorefile.
  require_restorefile_with_uuid = 1

  # Minimum size (in KB) of block devices which can be used as PVs.
  # In a clustered environment all nodes must use the same value.
  # Any value smaller than 512KB is ignored.

  # Ignore devices smaller than 2MB such as floppy drives.
  pv_min_size = 2048

  # The original built-in setting was 512 up to and including version 2.02.84.
  # pv_min_size = 512

  # Issue discards to a logical volumes's underlying physical volume(s) when
  # the logical volume is no longer using the physical volumes' space (e.g.
  # lvremove, lvreduce, etc).  Discards inform the storage that a region is
  # no longer in use.  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 set
  # to 1, discards will only be issued if both the storage and kernel provide
  # support.
  # 1 enables; 0 disables.
  issue_discards = 0
}

# This section allows you to configure the way in which LVM selects
# free space for its Logical Volumes.
allocation {

  # 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.
  # Use the special tag "@*" as a wildcard to match any PV tag.

  # Example: LVs are mirrored between two sites within a single VG.
  # PVs are tagged with either @site1 or @site2 to indicate where
  # they are situated.

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

```
        # 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.  Set this to 0 to revert to the previous
        # algorithm.
        maximise_cling = 1

        # Set to 1 to guarantee that mirror logs will always be placed on
        # different PVs from the mirror images.  This was the default
        # until version 2.02.85.
        mirror_logs_require_separate_pvs = 0

        # Set to 1 to guarantee that thin pool metadata will always
        # be placed on different PVs from the pool data.
        thin_pool_metadata_require_separate_pvs = 0

        # Specify the minimal chunk size (in KB) for thin pool volumes.
        # Use of the larger chunk size may improve perfomance 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 64KB
        # Supported values are in range from 64 to 1048576.
        # thin_pool_chunk_size = 64

        # Specify discards behavior of the thin pool volume.
        # Select one of  "ignore", "nopassdown", "passdown"
        # thin_pool_discards = "passdown"

        # Set to 0, to disable zeroing of thin pool data chunks before their
        # first use.
        # N.B. zeroing larger thin pool chunk size degrades performance.
        # thin_pool_zero = 1
}

# This section that allows you to configure the nature of the
# information that LVM2 reports.
log {

        # Controls the messages sent to stdout or stderr.
        # There are three levels of verbosity, 3 being the most verbose.
        verbose = 0

        # Set to 1 to suppress all non-essential messages from stdout.
        # This has the same effect as -qq.
        # When this is set, 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

        # Should we send log messages through syslog?
        # 1 is yes; 0 is no.
        syslog = 1
```

```
# Should we log error and debug messages to a file?
# By default there is no log file.
#file = "/var/log/lvm2.log"

# Should we overwrite the log file each time the program is run?
# By default we append.
overwrite = 0

# What level of log messages should we send to the log file and/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

# Format of output messages
# Whether or not (1 or 0) to indent messages according to their severity
indent = 1

# Whether or not (1 or 0) to display the command name on each line output
command_names = 0

# A prefix to use before the message text (but after the command name,
# if selected).  Default is two spaces, so you can see/grep the severity
# of each message.
prefix = "  "

# To make the messages look similar to the original LVM tools use:
#   indent = 0
#   command_names = 1
#   prefix = " -- "

# Set this if you want log messages during activation.
# Don't use this in low memory situations (can deadlock).
# activation = 0

# 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
# Use "all" to see everything.
debug_classes = [ "memory", "devices", "activation", "allocation",
    "lvmetad", "metadata", "cache", "locking" ]
}

# Configuration of metadata backups and archiving.  In LVM2 when we
# talk about a 'backup' we mean making a copy of the metadata for the
# *current* system.  The 'archive' contains old metadata configurations.
# Backups are stored in a human readeable text format.
backup {

# Should we maintain a backup of the current metadata configuration ?
# Use 1 for Yes; 0 for No.
# Think very hard before turning this off!
backup = 1

# Where shall we keep it ?
```

```
    # Remember to back up this directory regularly!
    backup_dir = "/etc/lvm/backup"

    # Should we maintain an archive of old metadata configurations.
    # Use 1 for Yes; 0 for No.
    # On by default.  Think very hard before turning this off.
    archive = 1

    # Where should archived files go ?
    # Remember to back up this directory regularly!
    archive_dir = "/etc/lvm/archive"

    # What is the minimum number of archive files you wish to keep ?
    retain_min = 10

    # What is the minimum time you wish to keep an archive file for ?
    retain_days = 30
}

# Settings for the running LVM2 in shell (readline) mode.
shell {

    # Number of lines of history to store in ~/.lvm_history
    history_size = 100
}


# Miscellaneous global LVM2 settings
global {
    # The file creation mask for any files and directories created.
    # Interpreted as octal if the first digit is zero.
    umask = 077

    # Allow other users to read the files
    #umask = 022

    # Enabling test mode means that no changes to the on disk metadata
    # will be made.  Equivalent to having the -t option on every
    # command.  Defaults to off.
    test = 0

    # Default value for --units argument
    units = "h"

    # Since version 2.02.54, the tools distinguish between powers of
    # 1024 bytes (e.g. KiB, MiB, GiB) and powers of 1000 bytes (e.g.
    # KB, MB, GB).
    # If you have scripts that depend on the old behaviour, set this to 0
    # temporarily until you update them.
    si_unit_consistency = 1

    # Whether or not to communicate with the kernel device-mapper.
    # Set to 0 if you want to use the tools to manipulate LVM metadata
    # without activating any logical volumes.
    # If the device-mapper kernel driver is not present in your kernel
    # setting this to 0 should suppress the error messages.
```

```
activation = 1

# If we can't communicate with device-mapper, should we try running
# the LVM1 tools?
# This option only applies to 2.4 kernels and is provided to help you
# switch between device-mapper kernels and LVM1 kernels.
# The LVM1 tools need to be installed with .lvm1 suffices
# e.g. vgscan.lvm1 and they will stop working after you start using
# the new lvm2 on-disk metadata format.
# The default value is set when the tools are built.
# fallback_to_lvm1 = 0

# The default metadata format that commands should use - "lvm1" or "lvm2".
# The command line override is -M1 or -M2.
# Defaults to "lvm2".
# format = "lvm2"

# Location of proc filesystem
proc = "/proc"

# Type of locking to use. Defaults to local file-based locking (1).
# Turn locking off by setting to 0 (dangerous: risks metadata corruption
# if LVM2 commands get run concurrently).
# Type 2 uses the external shared library locking_library.
# Type 3 uses built-in clustered locking.
# Type 4 uses read-only locking which forbids any operations that might
# change metadata.
locking_type = 1

# Set to 0 to fail when a lock request cannot be satisfied immediately.
wait_for_locks = 1

# If using external locking (type 2) and initialisation fails,
# with this set to 1 an attempt will be made to use the built-in
# clustered locking.
# If you are using a customised locking_library you should set this to 0.
fallback_to_clustered_locking = 1

# If an attempt to initialise type 2 or type 3 locking failed, perhaps
# because cluster components such as clvmd are not running, with this set
# to 1 an attempt will be made to use local file-based locking (type 1).
# If this succeeds, only commands against local volume groups will proceed.
# Volume Groups marked as clustered will be ignored.
fallback_to_local_locking = 1

# 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 = "/var/lock/lvm"

# Whenever 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.
# NB. This option only affects locking_type = 1 viz. local file-based
# locking.
```

```
prioritise_write_locks = 1

# Other entries can go here to allow you to load shared libraries
# e.g. if support for LVM1 metadata was compiled as a shared library use
#   format_libraries = "liblvm2format1.so"
# Full pathnames can be given.

# Search this directory first for shared libraries.
#   library_dir = "/lib"

# The external locking library to load if locking_type is set to 2.
#   locking_library = "liblvm2clusterlock.so"

# 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

# Check whether CRC is matching when parsed VG is used multiple times.
# This is useful to catch unexpected internal cached volume group
# structure modification. Please only enable for debugging.
detect_internal_vg_cache_corruption = 0

# If set to 1, no operations that change on-disk metadata will be 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

# 'mirror_segtype_default' defines which segtype will be used when the
# shorthand '-m' option is used for mirroring.  The possible options are:
#
# "mirror" - The original RAID1 implementation provided by LVM2/DM.  It is
#         characterized by a flexible log solution (core, disk, mirrored)
#   and by the necessity to block I/O while reconfiguring in the
#   event of 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.
#     Ref: https://bugzilla.redhat.com/show_bug.cgi?id=817130#c10
#
# "raid1"  - This implementation leverages MD's 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 - no separate devices are
#   required.)  This mirror implementation does not require I/O
#   to be blocked in the kernel in the event of a failure.
#   This mirror implementation is not cluster-aware and cannot be
#   used in a shared (active/active) fashion in a cluster.
#
# Specify the '--type <mirror|raid1>' option to override this default
# setting.
mirror_segtype_default = "mirror"

# 'raid10_segtype_default' determines the segment types used by default
```

when the '--stripes/-i' and '--mirrors/-m' arguments are both specified
during the creation of a logical volume.
Possible settings include:
#
"raid10" - This implementation leverages MD's RAID10 personality through
device-mapper.
#
# "mirror" - LVM will layer the 'mirror' and 'stripe' segment types. It
#   will do this by creating a mirror on top of striped sub-LVs;
#   effectively creating a RAID 0+1 array. This is suboptimal
#   in terms of providing redunancy and performance. Changing to
#   this setting is not advised.
# Specify the '--type <raid10|mirror>' option to override this default
# setting.
raid10_segtype_default = "mirror"

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.
Set to 1 to reinstate the previous format.
#
lvdisplay_shows_full_device_path = 0

Whether to use (trust) a running instance of lvmetad. If this is set to
0, all commands fall back to the usual scanning mechanisms. When set to 1
and when lvmetad is running (it is not auto-started), the volume group
metadata and PV state flags are obtained from the lvmetad instance and no
scanning is done by the individual commands. In a setup with lvmetad,
lvmetad udev rules *must* be set up for LVM to work correctly. Without
proper udev rules, all changes in block device configuration will be
ignored until a manual 'pvscan --cache' is performed.
#
# If lvmetad has been running while use_lvmetad was 0, it MUST be stopped
# before changing use_lvmetad to 1 and started again afterwards.
use_lvmetad = 0

Full path of the utility called to check that a thin metadata device
is in a state that allows it to be used.
Each time a thin pool needs to be activated or after it is deactivated
this utility is executed. The activation will only proceed if the utility
has an exit status of 0.
Set to "" to skip this check. (Not recommended.)
The thin tools are available as part of the device-mapper-persistent-data
package from https://github.com/jthornber/thin-provisioning-tools.
#
thin_check_executable = "/usr/sbin/thin_check"

Array of string options passed with thin_check command. By default,
option "-q" is for quiet output.
With thin_check version 2.1 or newer you can add "--ignore-non-fatal-errors"
to let it pass through ignoreable errors and fix them later.
#
thin_check_options = ["-q"]

Full path of the utility called to repair a thin metadata device

```
    # is in a state that allows it to be used.
    # Each time a thin pool needs repair this utility is executed.
    # See thin_check_executable how to obtain binaries.
    #
    # thin_repair_executable = "/usr/sbin/thin_repair"

    # Array of extra string options passed with thin_repair command.
    # thin_repair_options = [ "" ]

    # Full path of the utility called to dump thin metadata content.
    # See thin_check_executable how to obtain binaries.
    #
    # thin_dump_executable = "/usr/sbin/thin_dump"

    # If set, given features are not used by thin driver.
    # This can be helpful not just for testing, but i.e. allows to avoid
    # using problematic implementation of some thin feature.
    # Features:
    #   block_size
    #   discards
    #   discards_non_power_2
    #   external_origin
    #   metadata_resize
    #
    # thin_disabled_features = [ "discards", "block_size" ]
}

activation {
    # Set to 1 to perform internal checks on the operations issued to
    # libdevmapper.  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

    # Set to 0 to disable udev synchronisation (if compiled into the binaries).
    # Processes will not wait for notification from udev.
    # They will continue irrespective of any possible udev processing
    # in the background.  You should only use this if udev is not running
    # or has rules that ignore the devices LVM2 creates.
    # The command line argument --nodevsync takes precedence over this setting.
    # If set to 1 when udev is not running, and there are LVM2 processes
    # waiting for udev, run 'dmsetup udevcomplete_all' manually to wake them up.
    udev_sync = 1

    # Set to 0 to disable the udev rules installed by LVM2 (if built with
    # --enable-udev_rules). LVM2 will then manage the /dev nodes and symlinks
    # for active logical volumes directly itself.
    # N.B. Manual intervention may be required if this setting is changed
    # while any logical volumes are active.
    udev_rules = 1

    # Set to 1 for LVM2 to verify operations performed by udev. This turns on
    # additional checks (and if necessary, repairs) on entries in the device
    # directory after udev has completed processing its events.
    # Useful for diagnosing problems with LVM2/udev interactions.
    verify_udev_operations = 0
```

# If set to 1 and if deactivation of an LV fails, perhaps because
# a process run from a quick udev rule temporarily opened the device,
# retry the operation for a few seconds before failing.
retry_deactivation = 1

# How to fill in missing stripes if activating an incomplete volume.
# 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 to in place of missing stripes.
# But note that using anything other than "error" with mirrored
# or snapshotted volumes is likely to result in data corruption.
missing_stripe_filler = "error"

# The linear target is an optimised version of the striped target
# that only handles a single stripe. Set this to 0 to disable this
# optimisation and always use the striped target.
use_linear_target = 1

# How much stack (in KB) to reserve for use while devices suspended
# Prior to version 2.02.89 this used to be set to 256KB
reserved_stack = 64

# How much memory (in KB) to reserve for use while devices suspended
reserved_memory = 8192

# Nice value used while devices suspended
process_priority = -18

If volume_list is defined, each LV is only activated if there is a
match against the list.
#
"vgname" and "vgname/lvname" are matched exactly.
"@tag" matches any tag set in the LV or VG.
"@*" matches if any tag defined on the host is also set in the LV or VG
#
If any host tags exist but volume_list is not defined, a default
single-entry list containing "@*" is assumed.
#
volume_list = ["vg1", "vg2/lvol1", "@tag1", "@*"]

If auto_activation_volume_list is defined, each LV that is to be
activated with the autoactivation option (--activate ay/-a ay) is
first checked against the list. There are two scenarios in which
the autoactivation option is used:
#
- automatic activation of volumes based on incoming PVs. If all the
PVs making up a VG are present in the system, the autoactivation
is triggered. This requires lvmetad (global/use_lvmetad=1) and udev
to be running. In this case, "pvscan --cache -aay" is called
automatically without any user intervention while processing
udev events. Please, make sure you define auto_activation_volume_list
properly so only the volumes you want and expect are autoactivated.
#
- direct activation on command line with the autoactivation option.
In this case, the user calls "vgchange --activate ay/-a ay" or

```
#     "lvchange --activate ay/-a ay" directly.
#
# By default, the auto_activation_volume_list is not defined and all
# volumes will be activated either automatically or by using --activate ay/-a ay.
#
# N.B. The "activation/volume_list" is still honoured in all cases so even
# if the VG/LV passes the auto_activation_volume_list, it still needs to
# pass the volume_list for it to be activated in the end.

# If auto_activation_volume_list is defined but empty, no volumes will be
# activated automatically and --activate ay/-a ay will do nothing.
#
# auto_activation_volume_list = []

# If auto_activation_volume_list is defined and it's not empty, only matching
# volumes will be activated either automatically or by using --activate ay/-a ay.
#
#   "vgname" and "vgname/lvname" are matched exactly.
#   "@tag" matches any tag set in the LV or VG.
#   "@*" matches if any tag defined on the host is also set in the LV or VG
#
# auto_activation_volume_list = [ "vg1", "vg2/lvol1", "@tag1", "@*" ]

# If read_only_volume_list is defined, each LV that is to be activated
# is checked against the list, and if it matches, it as activated
# in read-only mode.  (This overrides '--permission rw' stored in the
# metadata.)
#
#   "vgname" and "vgname/lvname" are matched exactly.
#   "@tag" matches any tag set in the LV or VG.
#   "@*" matches if any tag defined on the host is also set in the LV or VG
#
# read_only_volume_list = [ "vg1", "vg2/lvol1", "@tag1", "@*" ]

# Each LV can have an 'activation skip' flag stored persistently against it.
# During activation, this flag is used to decide whether such an LV is skipped.
# The 'activation skip' flag can be set during LV creation and by default it
# is automatically set for thin snapshot LVs. The 'auto_set_activation_skip'
# enables or disables this automatic setting of the flag while LVs are created.
# auto_set_activation_skip = 1

# For RAID or 'mirror' segment types, 'raid_region_size' is the
# size (in kiB) of each:
# - synchronization operation when initializing
# - each copy operation when performing a 'pvmove' (using 'mirror' segtype)
# This setting has replaced 'mirror_region_size' since version 2.02.99
raid_region_size = 512

# Setting to use when there is no readahead value stored in the metadata.
#
# "none" - Disable readahead.
# "auto" - Use default value chosen by kernel.
readahead = "auto"

# 'raid_fault_policy' defines how a device failure in a RAID logical
# volume is handled.  This includes logical volumes that have the following
```

```
# segment types: raid1, raid4, raid5*, and raid6*.
#
# In the event of a failure, the following policies will determine what
# actions are performed during the automated response to failures (when
# dmeventd is monitoring the RAID logical volume) and when 'lvconvert' is
# called manually with the options '--repair' and '--use-policies'.
#
# "warn" - Use the system log to warn the user that a device in the RAID
#    logical volume 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 logical volume (1 device for
#    raid4/5, 2 for raid6, etc) the logical volume will remain
#    usable.
#
# "allocate" - Attempt to use any extra physical volumes in the volume
#    group as spares and replace faulty devices.
#
raid_fault_policy = "warn"


# 'mirror_image_fault_policy' and 'mirror_log_fault_policy' define
# how a device failure affecting a mirror (of "mirror" segment type) is
# handled.  A mirror is composed of mirror images (copies) and a log.
# A disk log ensures that a mirror does not need to be re-synced
# (all copies made the same) every time a machine reboots or crashes.
#
# In the event of a failure, the specified policy will be used to determine
# what happens. This applies to automatic repairs (when the mirror is being
# monitored by dmeventd) and to manual lvconvert --repair when
# --use-policies is given.
#
# "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 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_log_fault_policy = "allocate"
mirror_image_fault_policy = "remove"

# 'snapshot_autoextend_threshold' and 'snapshot_autoextend_percent' define
# how to handle automatic snapshot extension. The former defines when the
# snapshot should be extended: when its space usage exceeds this many
# percent. The latter defines how much extra space should be allocated for
# the snapshot, in percent of its current size.
#
# For example, if you set snapshot_autoextend_threshold to 70 and
# snapshot_autoextend_percent to 20, whenever a snapshot exceeds 70% usage,
# it will be extended by another 20%. For a 1G snapshot, using up 700M will
# trigger a resize to 1.2G. When the usage exceeds 840M, the snapshot will
# be extended to 1.44G, and so on.
#
# Setting snapshot_autoextend_threshold to 100 disables automatic
# extensions. The minimum value is 50 (A setting below 50 will be treated
# as 50).

snapshot_autoextend_threshold = 100
snapshot_autoextend_percent = 20

# 'thin_pool_autoextend_threshold' and 'thin_pool_autoextend_percent' define
# how to handle automatic pool extension. The former defines when the
# pool should be extended: when its space usage exceeds this many
# percent. The latter defines how much extra space should be allocated for
# the pool, in percent of its current size.
#
# For example, if you set thin_pool_autoextend_threshold to 70 and
# thin_pool_autoextend_percent to 20, whenever a pool exceeds 70% usage,
# it will be extended by another 20%. For a 1G pool, using up 700M will
# trigger a resize to 1.2G. When the usage exceeds 840M, the pool will
# be extended to 1.44G, and so on.
#
# Setting thin_pool_autoextend_threshold to 100 disables automatic
# extensions. The minimum value is 50 (A setting below 50 will be treated
# as 50).

thin_pool_autoextend_threshold = 100
thin_pool_autoextend_percent = 20

# While activating devices, I/O to devices being (re)configured is
# suspended, and as a precaution against deadlocks, LVM2 needs to pin
# any memory it is using so it is not paged out.  Groups of pages that
# are known not to be accessed during activation need not be pinned
# into memory.  Each string listed in this setting is compared against
# each line in /proc/self/maps, and the pages corresponding to any
# 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.
# mlock_filter = [ "locale/locale-archive", "gconv/gconv-modules.cache" ]

# Set to 1 to revert to the default behaviour prior to version 2.02.62
# which used mlockall() to pin the whole process's memory while activating
# devices.
use_mlockall = 0
```

```
# Monitoring is enabled by default when activating logical volumes.
# Set to 0 to disable monitoring or use the --ignoremonitoring option.
monitoring = 1

# 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.  The default is 15 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 the
# operation is complete.
polling_interval = 15
}


####################
# Advanced section #
###################

# Metadata settings
#
# metadata {
    # Default number of copies of metadata to hold on each PV.  0, 1 or 2.
    # You might want to override it from the command line with 0
    # when running pvcreate on new PVs which are to be added to large VGs.

    # pvmetadatacopies = 1

    # Default number of copies of metadata to maintain for each VG.
    # 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 default value of 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'.

    # vgmetadatacopies = 0

    # Approximate default size of on-disk metadata areas in sectors.
    # You should increase this if you have large volume groups or
    # you want to retain a large on-disk history of your metadata changes.

    # pvmetadatasize = 255

    # List of directories holding live copies of text format metadata.
    # These directories must not be on logical volumes!
    # It's possible to use LVM2 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
    # you are absolutely sure you know what you are doing! Use
    # the supplied toolset to make changes (e.g. vgcfgrestore).

    # dirs = [ "/etc/lvm/metadata", "/mnt/disk2/lvm/metadata2" ]
#}

# Event daemon
#
dmeventd {
    # mirror_library is the library used 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"

    # snapshot_library is the library used when monitoring a snapshot device.
    #
    # "libdevmapper-event-lvm2snapshot.so" monitors the filling of
    # snapshots and emits a warning through syslog when the use of
    # the snapshot exceeds 80%. The warning is repeated when 85%, 90% and
    # 95% of the snapshot is filled.

    snapshot_library = "libdevmapper-event-lvm2snapshot.so"

    # thin_library is the library used when monitoring a thin device.
    #
    # "libdevmapper-event-lvm2thin.so" monitors the filling of
    # pool and emits a warning through syslog when the use of
    # the pool exceeds 80%. The warning is repeated when 85%, 90% and
    # 95% of the pool is filled.

    thin_library = "libdevmapper-event-lvm2thin.so"

    # Full path of the dmeventd binary.
    #
    # executable = "/sbin/dmeventd"
}
```

APPENDIX C. LVM SELECTION CRITERIA

As of Red Hat Enterpise Linux release 6.6, many LVM reporting commmands accept the **-S** or **--select** option to define selection criteria for those commands. As of Red Hat Enterprise Linux release 6.7, 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 6.8 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.3.1, "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**, and **size**.

```
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]
```

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
!=	Not equal to	number, size, percent, string, string list
>=	Greater than or equal to	number, size, percent
>	Greater than	number, size, percent
<=	Less than or equal to	number, size, percent
<	Less than	number, size, percent

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
(	Left parenthesis (grouping operator)
)	Right parenthesis (grouping operator)
[	List start (grouping operator)
]	List end (grouping operator)
{	List subset start (grouping operator)

Logical and Grouping Operator	Description
}	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
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

Logical Volume Field	Description	Field Type
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
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

Logical Volume Field	Description	Field Type
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
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	string

Logical Volume Field	Description	Field Type
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
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. Pysical 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

Physical Volume Field	Description	Field Type
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
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

Volume Group Field	Description	Field Type
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
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

Volume Group Field	Description	Field Type
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

Logical Volume Segment Field	Description	Field Type
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
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, "Pysical Volume Segment Fields" describes the physical volume segment fields and their field types.

Table C.13. Pysical 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
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	writeable	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

Field	Field Value	Synonyms
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
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

Field	Field Value	Synonyms
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
zero	""	0
lv_permissions	writeable	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.3.1. 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.3.1.1, "Standard time selection format" and Section C.3.1.2, "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.3.1.1. 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.3.1.2. 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 the show the freeform date/time specificaiton 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.4. 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                      linear      public
  swap            f1 -wi-ao---- 512.00m                     linear      public
  [lvol0_pmspare] vg ewi-------  4.00m                       linear      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                       linear      private, \
                                                      thin,pool, \
                                                      data
```

```
  [pool_tmeta]    vg ewi-ao----   4.00m                            linear       private,  \
                                                                   thin,pool, \
                                                                   metadata
```

The following command displays all logical volumes with "lvol[13]" in their name, using a regular expression to specify this.

```
# lvs -a -o+layout,role -S 'lv_name=~lvol[13]'
  LV    VG  Attr      LSize Pool Origin Data%  Layout      Role
  lvol1 vg   Vwi-a-tz-- 1.00g pool        0.00   thin,sparse public
  lvol3 vg   Vwi---tz-k 1.00g pool lvol2       thin,sparse public,snapshot,thinsnapshot
```

The following command displays all logical volumes greater than 500m 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
  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 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
  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 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.5. 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 **lvol4** 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 volmes 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                    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                 lvol2  thin,sparse public,snapshot,thinsnapshot
lvol4 active skip activation      linear     public
lvol5 active                thin,sparse public,origin,thinorigin
lvol6                 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
lvol6                 lvol5  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 **lvol1**, 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
lvol1 vg
lvol2 vg
lvol3 vg
lvol4 vg
lvol5 vg
lvol6 vg
pool  vg
```

```
# vgchange -ay -S 'lv_name=lvol1'
  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 lvol[456] or the logical volume size is more than 5g.

```
# lvchange --addtag mytag -S '(role=origin && lv_name=~lvol[456]) || lv_size > 5g'
  Logical volume "root" changed.
  Logical volume "lvol5" 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.

As of the Red Hat Enterprise Linux 6.1 release, LVM tags are strings of up to 1024 characters (for earlier releases the upper length limit was 128 characters). LVM tags cannot start with a hyphen.

A valid tag can consist of a limited range of characters only. For the Red Hat Enterprise Linux 6.0 release, the allowed characters are [A-Za-z0-9_+.-]. As of the Red Hat Enterprise Linux 6.1 release, the list of allowed characters has been 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.

As of the Red Hat Enterprise Linux 6.1 release, 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 allow 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 hostnames 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: Read/Write? Resizeable?

- 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-jO6s-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-Ur9OF9"
                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 10.0-3 Version for 6.9 GA Publication	**Wed Mar 8 2017**	**Steven Levine**
Revision 10.0-1 Version for 6.9 Beta Publication	**Fri Dec 16 2016**	**Steven Levine**
Revision 9.0-5 Updating document for 6.8.	**Tue Nov 29 2016**	**Steven Levine**
Revision 9.0-4 Preparing document for 6.8 GA publication.	**Wed Apr 27 2016**	**Steven Levine**
Revision 9.0-3 Initial revision for Red Hat Enterprise Linux 6.8 Beta release	**Wed Mar 9 2016**	**Steven Levine**
Revision 8.0-16 Republish version for Red Hat Enterprise Linux 6.7	**Mon Jul 13 2015**	**Steven Levine**
Revision 8.0-14 Initial revision for Red Hat Enterprise Linux 6.7	**Wed Jul 8 2015**	**Steven Levine**
Revision 8.0-13 Initial revision for Red Hat Enterprise Linux 6.7 Beta release	**Tue Apr 14 2015**	**Steven Levine**
Revision 7.0-9 Initial revision for Red Hat Enterprise Linux 6.6	**Wed Oct 8 2014**	**Steven Levine**
Revision 7.0-8 Initial revision for Red Hat Enterprise Linux 6.6 Beta release	**Thu Aug 7 2014**	**Steven Levine**
Revision 6.0-14 Initial revision for Red Hat Enterprise Linux 6.5	**Wed Nov 13 2013**	**Steven Levine**
Revision 6.0-10 Initial revision for Red Hat Enterprise Linux 6.5 Beta release	**Fri Sep 27 2013**	**Steven Levine**
Revision 5.0-19 Initial revision for Red Hat Enterprise Linux 6.4	**Mon Feb 18 2013**	**John Ha**
Revision 5.0-12 Initial revision for Red Hat Enterprise Linux 6.4 Beta release	**Mon Nov 27 2012**	**Steven Levine**
Revision 4.0-2 Initial revision for Red Hat Enterprise Linux 6.3	**Fri Jun 15 2012**	**Steven Levine**
Revision 3.0-1 Initial revision for Red Hat Enterprise Linux 6.2 Beta release	**Mon Sep 19 2011**	**Steven Levine**
Revision 2.0-1 Initial revision for Red Hat Enterprise Linux 6.1	**Thu May 19 2011**	**Steven Levine**
Revision 1.0-1 Initial release for Red Hat Enterprise Linux 6	**Wed Nov 10 2010**	**Steven Levine**

INDEX

Symbols

A

B

C