



# Red Hat Enterprise Linux 7 DM Multipath

---

DM Multipath Configuration and Administration





## Legal Notice

Copyright © 2015 Red Hat, Inc. and others.

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

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

## Abstract

This book provides information on using the Device-Mapper Multipath feature of Red Hat Enterprise Linux 7.

## Table of Contents

<b>Chapter 1. Device Mapper Multipathing</b> .....	<b>2</b>
1.1. New and Changed Features	2
1.2. Overview of DM Multipath	2
1.3. Storage Array Support	5
1.4. DM Multipath Components	5
1.5. DM Multipath Setup Overview	6
<b>Chapter 2. Multipath Devices</b> .....	<b>7</b>
2.1. Multipath Device Identifiers	7
2.2. Consistent Multipath Device Names in a Cluster	7
2.3. Multipath Device Attributes	8
2.4. Multipath Devices in Logical Volumes	8
<b>Chapter 3. Setting Up DM Multipath</b> .....	<b>10</b>
3.1. Setting Up DM Multipath	10
3.2. Ignoring Local Disks when Generating Multipath Devices	11
3.3. Configuring Storage Devices	13
3.4. Setting Up Multipathing in the initramfs File System	14
<b>Chapter 4. The DM Multipath Configuration File</b> .....	<b>15</b>
4.1. Configuration File Overview	15
4.2. Configuration File Blacklist	16
4.3. Configuration File Defaults	18
4.4. Multipaths Device Configuration Attributes	26
4.5. Configuration File Devices	30
<b>Chapter 5. DM Multipath Administration and Troubleshooting</b> .....	<b>36</b>
5.1. Automatic Configuration File Generation with Multipath Helper	36
5.2. Resizing an Online Multipath Device	36
5.3. Moving root File Systems from a Single Path Device to a Multipath Device	36
5.4. Moving swap File Systems from a Single Path Device to a Multipath Device	37
5.5. The Multipath Daemon	38
5.6. Issues with Large Number of LUNs	38
5.7. Issues with queue_if_no_path feature	39
5.8. Multipath Command Output	39
5.9. Multipath Queries with multipath Command	40
5.10. Multipath Command Options	41
5.11. Determining Device Mapper Entries with the dmsetup Command	41
5.12. Troubleshooting with the multipathd Interactive Console	42
<b>Appendix A. Revision History</b> .....	<b>43</b>
<b>Index</b> .....	<b>44</b>

## Chapter 1. Device Mapper Multipathing

Device mapper multipathing (DM Multipath) allows you to configure multiple I/O paths between server nodes and storage arrays into a single device. These I/O paths are physical SAN connections that can include separate cables, switches, and controllers. Multipathing aggregates the I/O paths, creating a new device that consists of the aggregated paths.

This chapter provides a summary of the features of DM-Multipath that were added subsequent to the initial release of Red Hat Enterprise Linux 7. Following that, this chapter provides a high-level overview of DM Multipath and its components, as well as an overview of DM-Multipath setup.

### 1.1. New and Changed Features

This section lists features of the DM Multipath that are new since the initial release of Red Hat Enterprise Linux 7.

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

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

- [Table 5.1, “Useful multipath Command Options”](#). now includes entries for the `-w` and `-W` options of the `multipath` command, which allow you to better manage the `wwids` file.
- Additional options for the `values` argument of the `features` parameter in the `multipath.conf` file are documented in [Chapter 4, The DM Multipath Configuration File](#).
- [Table 4.1, “Multipath Configuration Defaults”](#). includes an entry for the `force_sync` parameter, which prevents path checkers from running in `async` mode when set to `“yes”`.

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

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

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

- This document includes a new section, [Section 5.1, “Automatic Configuration File Generation with Multipath Helper”](#). The Multipath Helper application gives you options to create multipath configurations with custom aliases, device blacklists, and settings for the characteristics of individual multipath devices.
- The `defaults` section of the `multipath.conf` configuration file supports the new `config_dir`, `new_bindings_in_boot`, `ignore_new_boot_devs`, `retrigger_tries`, and `retrigger_delays` parameters. The `defaults` section of the `multipath.conf` file is documented in [Table 4.1, “Multipath Configuration Defaults”](#).
- The `defaults`, `devices`, and `multipaths` sections of the `multipath.conf` configuration file now support the `delay_watch_checks` and `delay_wait_checks` configuration parameters. For information on the configuration parameters, see [Chapter 4, The DM Multipath Configuration File](#).

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

### 1.2. Overview of DM Multipath

DM Multipath can be used to provide:


➤ Redundancy

DM Multipath can provide failover in an active/passive configuration. In an active/passive configuration, only half the paths are used at any time for I/O. If any element of an I/O path (the cable, switch, or controller) fails, DM Multipath switches to an alternate path.

➤ Improved Performance

DM Multipath can be configured in active/active mode, where I/O is spread over the paths in a round-robin fashion. In some configurations, DM Multipath can detect loading on the I/O paths and dynamically re-balance the load.

[Figure 1.1, “Active/Passive Multipath Configuration with One RAID Device”](#) shows an active/passive configuration with two I/O paths from the server to a RAID device. There are 2 HBAs on the server, 2 SAN switches, and 2 RAID controllers.




**Figure 1.1. Active/Passive Multipath Configuration with One RAID Device**

In this configuration, there is one I/O path that goes through hba1, SAN1, and controller 1 and a second I/O path that goes through hba2, SAN2, and controller2. There are many points of possible failure in this configuration:

- HBA failure
- FC cable failure
- SAN switch failure
- Array controller port failure

With DM Multipath configured, a failure at any of these points will cause DM Multipath to switch to the alternate I/O path.

[Figure 1.2, “Active/Passive Multipath Configuration with Two RAID Devices”](#) shows a more complex active/passive configuration with 2 HBAs on the server, 2 SAN switches, and 2 RAID devices with 2 RAID controllers each.



**Figure 1.2. Active/Passive Multipath Configuration with Two RAID Devices**

In the example shown in [Figure 1.2, “Active/Passive Multipath Configuration with Two RAID Devices”](#), there are two I/O paths to each RAID device (just as there are in the example shown in [Figure 1.1, “Active/Passive Multipath Configuration with One RAID Device”](#)). With DM Multipath configured, a failure at any of the points of the I/O path to either of the RAID devices will cause DM Multipath to switch to the alternate I/O path for that device.

[Figure 1.3, “Active/Active Multipath Configuration with One RAID Device”](#) shows an active/active configuration with 2 HBAs on the server, 1 SAN switch, and 2 RAID controllers. There are four I/O paths from the server to a storage device:

- ✦ hba1 to controller1
- ✦ hba1 to controller2
- ✦ hba2 to controller1
- ✦ hba2 to controller2

In this configuration, I/O can be spread among those four paths.





Figure 1.3. Active/Active Multipath Configuration with One RAID Device

### 1.3. Storage Array Support

By default, DM Multipath includes support for the most common storage arrays that themselves support DM Multipath. For information on the default configuration values, including supported devices, run either of the following commands.

```
# multipathd show config
# multipath -t
```

If your storage array supports DM Multipath and is not configured by default, you may need to add it to the DM Multipath configuration file, `multipath.conf`. For information on the DM Multipath configuration file, see [Chapter 4, The DM Multipath Configuration File](#).

Some storage arrays require special handling of I/O errors and path switching. These require separate hardware handler kernel modules.

### 1.4. DM Multipath Components

[Table 1.1, “DM Multipath Components”](#). describes the components of DM Multipath.

Table 1.1. DM Multipath Components

Component	Description
<code>dm_multipath</code> kernel module	Reroutes I/O and supports failover for paths and path groups.
<code>mpathconf</code> utility	Configures and enables device mapper multipathing.

Component	Description
<b>multipath</b> command	Lists and configures multipath devices. Normally started with <code>/etc/rc.sysinit</code> , it can also be started by a <b>udev</b> program whenever a block device is added.
<b>multipathd</b> daemon	Monitors paths; as paths fail and come back, it may initiate path group switches. Allows interactive changes to multipath devices. The daemon must be restarted following any changes to the <code>/etc/multipath.conf</code> file.
<b>kpartx</b> command	Creates device mapper devices for the partitions on a device. It is necessary to use this command for DOS-based partitions with DM Multipath. The <b>kpartx</b> command is provided in its own package, but the <b>device-mapper-multipath</b> package depends on it.

## 1.5. DM Multipath Setup Overview

DM Multipath includes compiled-in default settings that are suitable for common multipath configurations. The basic procedure for configuring your system with DM Multipath is as follows:

1. Install the **device-mapper-multipath** rpm.
2. Create the configuration file and enable multipathing with the **mpathconf** command. You can also start the multipath daemon with this command if you do not need to edit the configuration file.
3. If necessary, edit the **multipath.conf** configuration file to modify default values and save the updated file.
4. Start the multipath daemon.

For detailed setup instructions for multipath configuration see [Chapter 3, Setting Up DM Multipath](#).

## Chapter 2. Multipath Devices

Without DM Multipath, each path from a server node to a storage controller is treated by the system as a separate device, even when the I/O path connects the same server node to the same storage controller. DM Multipath provides a way of organizing the I/O paths logically, by creating a single multipath device on top of the underlying devices.

### 2.1. Multipath Device Identifiers

Each multipath device has a World Wide Identifier (WWID), which is guaranteed to be globally unique and unchanging. By default, the name of a multipath device is set to its WWID. Alternately, you can set the **user\_friendly\_names** option in the multipath configuration file, which sets the alias to a node-unique name of the form **mpathn**.

For example, a node with two HBAs attached to a storage controller with two ports via a single unzoned FC switch sees four devices: **/dev/sda**, **/dev/sdb**, **dev/sdc**, and **/dev/sdd**. DM Multipath creates a single device with a unique WWID that reroutes I/O to those four underlying devices according to the multipath configuration. When the **user\_friendly\_names** configuration option is set to **yes**, the name of the multipath device is set to **mpathn**.

When new devices are brought under the control of DM Multipath, the new devices may be seen in two different places under the **/dev** directory: **/dev/mapper/mpathn** and **/dev/dm-n**.

- The devices in **/dev/mapper** are created early in the boot process. Use these devices to access the multipathed devices, for example when creating logical volumes.
- Any devices of the form **/dev/dm-n** are for internal use only should never be used by the administrator directly.

For information on the multipath configuration defaults, including the **user\_friendly\_names** configuration option, see [Section 4.3, “Configuration File Defaults”](#).

You can also set the name of a multipath device to a name of your choosing by using the **alias** option in the **multipaths** section of the multipath configuration file. For information on the **multipaths** section of the multipath configuration file, see [Section 4.4, “Multipaths Device Configuration Attributes”](#).

### 2.2. Consistent Multipath Device Names in a Cluster

When the **user\_friendly\_names** configuration option is set to **yes**, the name of the multipath device is unique to a node, but it is not guaranteed to be the same on all nodes using the multipath device. Similarly, if you set the **alias** option for a device in the **multipaths** section of the **multipath.conf** configuration file, the name is not automatically consistent across all nodes in the cluster. This should not cause any difficulties if you use LVM to create logical devices from the multipath device, but if you require that your multipath device names be consistent in every node it is recommended that you not set the **user\_friendly\_names** option to **yes** and that you not configure aliases for the devices. By default, if you do not set **user\_friendly\_names** to **yes** or configure an alias for a device, a device name will be the WWID for the device, which is always the same.

If you want the system-defined user-friendly names to be consistent across all nodes in the cluster, however, you can follow this procedure:

1. Set up all of the multipath devices on one machine.
2. Disable all multipath devices on other machines by running the following commands:

```
# service multipathd stop
# multipath -F
```

3. Copy the `/etc/multipath/bindings` file from the first machine to all the other machines in the cluster.
4. Re-enable the `multipathd` daemon on all the other machines in the cluster by running the following command:

```
# service multipathd start
```

If you add a new device, you will need to repeat this process.

Similarly, if you configure an alias for a device that you would like to be consistent across the nodes in the cluster, you should ensure that the `/etc/multipath.conf` file is the same for each node in the cluster by following the same procedure:

1. Configure the aliases for the multipath devices in the `multipath.conf` file on one machine.
2. Disable all multipath devices on other machines by running the following commands:

```
# service multipathd stop
# multipath -F
```

3. Copy the `/etc/multipath.conf` file from the first machine to all the other machines in the cluster.
4. Re-enable the `multipathd` daemon on all the other machines in the cluster by running the following command:

```
# service multipathd start
```

When you add a new device you will need to repeat this process.

## 2.3. Multipath Device Attributes

In addition to the `user_friendly_names` and `alias` options, a multipath device has numerous attributes. You can modify these attributes for a specific multipath device by creating an entry for that device in the `multipaths` section of the multipath configuration file. For information on the `multipaths` section of the multipath configuration file, see [Section 4.4, “Multipaths Device Configuration Attributes”](#).

## 2.4. Multipath Devices in Logical Volumes

After creating multipath devices, you can use the multipath device names just as you would use a physical device name when creating an LVM physical volume. For example, if `/dev/mapper/mpatha` is the name of a multipath device, the following command will mark `/dev/mapper/mpatha` as a physical volume.

```
pvcreate /dev/mapper/mpatha
```

You can use the resulting LVM physical device when you create an LVM volume group just as you would use any other LVM physical device.



## Note

If you attempt to create an LVM physical volume on a whole device on which you have configured partitions, the **pvcreate** command will fail. Note that the Anaconda and Kickstart installation programs create empty partition tables if you do not specify otherwise for every block device. If you wish to use the whole device rather than a partition, you must remove the existing partitions from the device. You can remove existing partitions with the **kpartx -d** and the **fdisk** commands. If your system has block devices that are greater than 2Tb, you can use the **parted** command to remove partitions.

When you create an LVM logical volume that uses active/passive multipath arrays as the underlying physical devices, you should include filters in the **/etc/lvm/lvm.conf** file to exclude the disks that underlie the multipath devices. This is because if the array automatically changes the active path to the passive path when it receives I/O, multipath will failover and failback whenever LVM scans the passive path if these devices are not filtered. For active/passive arrays that require a command to make the passive path active, LVM prints a warning message when this occurs.

To filter all SCSI devices in the LVM configuration file (**lvm.conf**), include the following filter in the **devices** section of the file.

```
filter = [ "r/block/", "r/disk/", "r/sd.*/", "a/.*/" ]
```

## Chapter 3. Setting Up DM Multipath

This chapter provides step-by-step example procedures for configuring DM Multipath. It includes the following procedures:

- ✦ Basic DM Multipath setup
- ✦ Ignoring local disks
- ✦ Adding more devices to the configuration file
- ✦ Starting multipath in the `initramfs` file system

### 3.1. Setting Up DM Multipath

Before setting up DM Multipath on your system, ensure that your system has been updated and includes the `device-mapper-multipath` package.

You set up multipath with the `mpathconf` utility, which creates the multipath configuration file `/etc/multipath.conf`.

- ✦ If the `/etc/multipath.conf` file already exists, the `mpathconf` utility will edit it.
- ✦ If the `/etc/multipath.conf` file does not exist, the `mpathconf` utility will use the `/usr/share/doc/device-mapper-multipath-0.4.9/multipath.conf` file as the starting file.
- ✦ If the `/usr/share/doc/device-mapper-multipath-0.4.9/multipath.conf` file does not exist the `mpathconf` utility will create the `/etc/multipath.conf` file from scratch.

If you do not need to edit the `/etc/multipath.conf` file, you can set up DM Multipath for a basic failover configuration by running the following command. This command enables the multipath configuration file and starts the `multipathd` daemon.

```
# mpathconf --enable --with_multipathd y
```

If you need to edit the `/etc/multipath.conf` file before starting the `multipathd` daemon, use the following procedure to set up DM Multipath for a basic failover configuration.

1. Run the `mpathconf` command with the `--enable` option specified:

```
# mpathconf --enable
```

For information on additional options to the `mpathconf` command you may require, see the `mpathconf` man page or run the `mpathconf` command with the `--help` option specified.

```
# mpathconf --help
usage: /sbin/mpathconf <command>

Commands:
Enable: --enable
Disable: --disable
Set user_friendly_names (Default y): --user_friendly_names <y|n>
```

```
Set find_multipaths (Default y): --find_multipaths <y|n>
Load the dm-multipath modules on enable (Default y): --with_module
<y|n>
start/stop/reload multipathd (Default n): --with_multipathd <y|n>
```

2. Edit the `/etc/multipath.conf` file if necessary. The default settings for DM Multipath are compiled in to the system and do not need to be explicitly set in the `/etc/multipath.conf` file.

The default value of `path_grouping_policy` is set to `failover`, so in this example you do not need to edit the `/etc/multipath.conf` file. For information on changing the values in the configuration file to something other than the defaults, see [Chapter 4, The DM Multipath Configuration File](#).

The initial defaults section of the configuration file configures your system so that the names of the multipath devices are of the form `mpathn`; without this setting, the names of the multipath devices would be aliased to the WWID of the device.

3. Save the configuration file and exit the editor, if necessary.
4. Execute the following command:

```
# service multipathd start
```

Since the value of `user_friendly_name` is set to `yes` in the configuration file, the multipath devices will be created as `/dev/mapper/mpathn`. For information on setting the name of the device to an alias of your choosing, see [Chapter 4, The DM Multipath Configuration File](#).

If you do not want to use user friendly names, you can run the following command:

```
# mpathconf --enable --user_friendly_names n
```



### Note

If you find that you need to edit the multipath configuration file after you have started the multipath daemon, you must execute the `service multipathd reload` command for the changes to take effect.

## 3.2. Ignoring Local Disks when Generating Multipath Devices

Some machines have local SCSI cards for their internal disks. DM Multipath is not recommended for these devices. If you set the `find_multipaths` configuration parameter to `yes`, you should not have to blacklist these devices. For information on the `find_multipaths` configuration parameter, see [Section 4.3, “Configuration File Defaults”](#).

If you do not set the `find_multipaths` configuration parameter to `yes`, can use the following procedure to modify the multipath configuration file to ignore the local disks when configuring multipath.

1. Determine which disks are the internal disks and mark them as the ones to blacklist.

In this example, `/dev/sda` is the internal disk. Note that as originally configured in the default multipath configuration file, executing the `multipath -v2` command shows the local disk, `/dev/sda`, in the multipath map.

For further information on the `multipath` command output, see [Section 5.8, “Multipath Command Output”](#).

```
# multipath -v2
create: SIBM-ESXSST336732LC____F3ET0EP0Q000072428BX1 undef
WINSYS,SF2372
size=33 GB features="0" hwhandler="0" wp=undef
`-+- policy='round-robin 0' prio=1 status=undef
   |- 0:0:0:0 sda 8:0  [-----

device-mapper ioctl cmd 9 failed: Invalid argument
device-mapper ioctl cmd 14 failed: No such device or address
create: 3600a0b80001327d80000006d43621677 undef WINSYS,SF2372
size=12G features='0' hwhandler='0' wp=undef
`-+- policy='round-robin 0' prio=1 status=undef
   |- 2:0:0:0 sdb 8:16  undef ready  running
   ` - 3:0:0:0 sdf 8:80 undef ready  running

create: 3600a0b80001327510000009a436215ec undef WINSYS,SF2372
size=12G features='0' hwhandler='0' wp=undef
`-+- policy='round-robin 0' prio=1 status=undef
   |- 2:0:0:1 sdc 8:32 undef ready  running
   ` - 3:0:0:1 sdg 8:96 undef ready  running

create: 3600a0b80001327d800000070436216b3 undef WINSYS,SF2372
size=12G features='0' hwhandler='0' wp=undef
`-+- policy='round-robin 0' prio=1 status=undef
   |- 2:0:0:2 sdd 8:48 undef ready  running
   ` - 3:0:0:2 sdg 8:112 undef ready  running

create: 3600a0b80001327510000009b4362163e undef WINSYS,SF2372
size=12G features='0' hwhandler='0' wp=undef
`-+- policy='round-robin 0' prio=1 status=undef
   |- 2:0:0:3 sdd 8:64 undef ready  running
   ` - 3:0:0:3 sdg 8:128 undef ready  running
```

- In order to prevent the device mapper from mapping `/dev/sda` in its multipath maps, edit the blacklist section of the `/etc/multipath.conf` file to include this device. Although you could blacklist the `sda` device using a `devnode` type, that would not be a safe procedure since `/dev/sda` is not guaranteed to be the same on reboot. To blacklist individual devices, you can blacklist using the WWID of that device.

Note that in the output to the `multipath -v2` command, the WWID of the `/dev/sda` device is `SIBM-ESXSST336732LC____F3ET0EP0Q000072428BX1`. To blacklist this device, include the following in the `/etc/multipath.conf` file.

```
blacklist {
    wwid SIBM-ESXSST336732LC____F3ET0EP0Q000072428BX1
}
```



- After you have updated the `/etc/multipath.conf` file, you must manually tell the `multipathd` daemon to reload the file. The following command reloads the updated `/etc/multipath.conf` file.

```
# service multipathd reload
```

- Run the following command to remove the multipath device:

```
# multipath -f SIBM-ESXSST336732LC_____F3ET0EP0Q000072428BX1
```

- To check whether the device removal worked, you can run the `multipath -ll` command to display the current multipath configuration. For information on the `multipath -ll` command, see [Section 5.9, “Multipath Queries with multipath Command”](#).

To check that the blacklisted device was not added back, you can run the `multipath` command, as in the following example. The `multipath` command defaults to a verbosity level of `v2` if you do not specify a `-v` option.

```
# multipath

create: 3600a0b80001327d80000006d43621677 undef WINSYS,SF2372
size=12G features='0' hwhandler='0' wp=undef
`-+- policy='round-robin 0' prio=1 status=undef
  |- 2:0:0:0 sdb 8:16 undef ready running
  `- 3:0:0:0 sdf 8:80 undef ready running

create: 3600a0b80001327510000009a436215ec undef WINSYS,SF2372
size=12G features='0' hwhandler='0' wp=undef
`-+- policy='round-robin 0' prio=1 status=undef
  |- 2:0:0:1 sdc 8:32 undef ready running
  `- 3:0:0:1 sdg 8:96 undef ready running

create: 3600a0b80001327d800000070436216b3 undef WINSYS,SF2372
size=12G features='0' hwhandler='0' wp=undef
`-+- policy='round-robin 0' prio=1 status=undef
  |- 2:0:0:2 sdd 8:48 undef ready running
  `- 3:0:0:2 sdg 8:112 undef ready running

create: 3600a0b80001327510000009b4362163e undef WINSYS,SF2372
size=12G features='0' hwhandler='0' wp=undef
`-+- policy='round-robin 0' prio=1 status=undef
  |- 2:0:0:3 sdd 8:64 undef ready running
  `- 3:0:0:3 sdg 8:128 undef ready running
```

### 3.3. Configuring Storage Devices

By default, DM Multipath includes support for the most common storage arrays that themselves support DM Multipath. For information on the default configuration value, including supported devices, run either of the following commands.

```
# multipathd show config
# multipath -t
```

If you need to add a storage device that is not supported by default as a known multipath device, edit

the `/etc/multipath.conf` file and insert the appropriate device information.

For example, to add information about the HP Open-V series the entry looks like this. This example sets the device to queue for a minute (or 12 retries and 5 seconds per retry) after all paths have failed.

```
devices {
    device {
        vendor "HP"
        product "OPEN-V"
        no_path_retry 12
    }
}
```

For more information on the **devices** section of the configuration file, see [Section 4.5, “Configuration File Devices”](#).

### 3.4. Setting Up Multipathing in the `initramfs` File System

You can set up multipathing in the `initramfs` file system. After configuring multipath, you can rebuild the `initramfs` file system with the multipath configuration files by executing the `dracut` command with the following options:

```
# dracut --force --add multipath --include /etc/multipath
```

If you run multipath from the `initramfs` file system and you make any changes to the multipath configuration files, you must rebuild the `initramfs` file system for the changes to take effect.

## Chapter 4. The DM Multipath Configuration File

By default, DM Multipath provides configuration values for the most common uses of multipathing. In addition, DM Multipath includes support for the most common storage arrays that themselves support DM Multipath. For information on the default configuration values, including supported devices, run either of the following commands.

```
# multipathd show config
# multipath -t
```

You can override the default configuration values for DM Multipath by editing the `/etc/multipath.conf` configuration file. If necessary, you can also add a storage array that is not supported by default to the configuration file.



### Note

You can run set up multipathing in the `initramfs` file system. If you run `multipath` from the `initramfs` file system and you make any changes to the multipath configuration files, you must rebuild the `initramfs` file system for the changes to take effect. For information on rebuilding the `initramfs` file system with multipath, refer to [Section 3.4, “Setting Up Multipathing in the initramfs File System”](#).

This chapter provides information on parsing and modifying the `multipath.conf` file. It contains sections on the following topics:

- » Configuration file overview
- » Configuration file blacklist
- » Configuration file defaults
- » Configuration file multipaths
- » Configuration file devices

In the multipath configuration file, you need to specify only the sections that you need for your configuration, or that you wish to change from the default values. If there are sections of the file that are not relevant to your environment or for which you do not need to override the default values, you can leave them commented out, as they are in the initial file.

The configuration file allows regular expression description syntax.

Further information about the configuration file can be found on the `multipath.conf(5)` man page.

### 4.1. Configuration File Overview

The multipath configuration file is divided into the following sections:

#### **blacklist**

Listing of specific devices that will not be considered for multipath.

#### **blacklist\_exceptions**

Listing of multipath candidates that would otherwise be blacklisted according to the parameters of the blacklist section.

### defaults

General default settings for DM Multipath.

### multipaths

Settings for the characteristics of individual multipath devices. These values overwrite what is specified in the **defaults** and **devices** sections of the configuration file.

### devices

Settings for the individual storage controllers. These values overwrite what is specified in the **defaults** section of the configuration file. If you are using a storage array that is not supported by default, you may need to create a **devices** subsection for your array.

When the system determines the attributes of a multipath device, first it checks the multipath settings, then the devices settings, then the multipath system defaults.

## 4.2. Configuration File Blacklist

The **blacklist** section of the multipath configuration file specifies the devices that will not be used when the system configures multipath devices. Devices that are blacklisted will not be grouped into a multipath device.

In older releases of Red Hat Enterprise Linux, multipath always tried to create a multipath device for every path that was not explicitly blacklisted. As of Red Hat Enterprise Linux 6, however, if the **find\_multipaths** configuration parameter is set to **yes**, then multipath will create a device only if one of three conditions are met:

- ✦ There are at least two non-blacklisted paths with the same WWID.
- ✦ The user manually forces the creation of the device by specifying a device with the **multipath** command.
- ✦ A path has the same WWID as a multipath device that was previously created (even if that multipath device does not currently exist). Whenever a multipath device is created, multipath remembers the WWID of the device so that it will automatically create the device again as soon as it sees a path with that WWID. This allows you to have multipath automatically choose the correct paths to make into multipath devices, without have to edit the multipath blacklist.

If you have previously created a multipath device without using the **find\_multipaths** parameter and then you later set the parameter to **yes**, you may need to remove the WWIDs of any device you do not want created as a multipath device from the `/etc/multipath/wwids` file. The following shows a sample `/etc/multipath/wwids` file. The WWIDs are enclosed by slashes (/):

```
# Multipath wwids, Version : 1.0
# NOTE: This file is automatically maintained by multipath and
multipathd.
# You should not need to edit this file in normal circumstances.
#
# Valid WWIDs:
/3600d0230000000000e13955cc3757802/
/3600d0230000000000e13955cc3757801/
```

```
/3600d0230000000000e13955cc3757800/
/3600d02300069c9ce09d41c31f29d4c00/
/SWINSYS SF2372 0E13955CC3757802/
/3600d0230000000000e13955cc3757803/
```

With the `find_multipaths` parameter set to `yes`, you need to blacklist only the devices with multiple paths that you do not want to be multipathed. Because of this, it will generally not be necessary to blacklist devices.

If you do need to blacklist devices, you can do so according to the following criteria:

- ✦ By WWID, as described in [Section 4.2.1, “Blacklisting by WWID”](#)
- ✦ By device name, as described in [Section 4.2.2, “Blacklisting By Device Name”](#)
- ✦ By device type, as described in [Section 4.2.3, “Blacklisting By Device Type”](#)

By default, a variety of device types are blacklisted, even after you comment out the initial blacklist section of the configuration file. For information, see [Section 4.2.2, “Blacklisting By Device Name”](#).

### 4.2.1. Blacklisting by WWID

You can specify individual devices to blacklist by their World-Wide IDentification with a `wwid` entry in the `blacklist` section of the configuration file.

The following example shows the lines in the configuration file that would blacklist a device with a WWID of 26353900f02796769.

```
blacklist {
    wwid 26353900f02796769
}
```

### 4.2.2. Blacklisting By Device Name

You can blacklist device types by device name so that they will not be grouped into a multipath device by specifying a `devnode` entry in the `blacklist` section of the configuration file.

The following example shows the lines in the configuration file that would blacklist all SCSI devices, since it blacklists all `sd*` devices.

```
blacklist {
    devnode "^sd[a-z]"
}
```

You can use a `devnode` entry in the `blacklist` section of the configuration file to specify individual devices to blacklist rather than all devices of a specific type. This is not recommended, however, since unless it is statically mapped by `udev` rules, there is no guarantee that a specific device will have the same name on reboot. For example, a device name could change from `/dev/sda` to `/dev/sdb` on reboot.

By default, the following `devnode` entries are compiled in the default blacklist; the devices that these entries blacklist do not generally support DM Multipath. To enable multipathing on any of these devices, you would need to specify them in the `blacklist_exceptions` section of the configuration file, as described in [Section 4.2.4, “Blacklist Exceptions”](#).

```

blacklist {
    devnode "^(ram|raw|loop|fd|md|dm-|sr|scd|st)[0-9]*"
    devnode "^(td|ha)d[a-z]"
}

```

### 4.2.3. Blacklisting By Device Type

You can specify specific device types in the **blacklist** section of the configuration file with a **device** section. The following example blacklists all IBM DS4200 and HP devices.

```

blacklist {
    device {
        vendor "IBM"
        product "3S42"          #DS4200 Product 10
    }
    device {
        vendor "HP"
        product "*"
    }
}

```

### 4.2.4. Blacklist Exceptions

You can use the **blacklist\_exceptions** section of the configuration file to enable multipathing on devices that have been blacklisted by default.

For example, if you have a large number of devices and want to multipath only one of them (with the WWID of 3600d0230000000000e13955cc3757803), instead of individually blacklisting each of the devices except the one you want, you could instead blacklist all of them, and then allow only the one you want by adding the following lines to the `/etc/multipath.conf` file.

```

blacklist {
    wwid "*"
}

blacklist_exceptions {
    wwid "3600d0230000000000e13955cc3757803"
}

```

When specifying devices in the **blacklist\_exceptions** section of the configuration file, you must specify the exceptions in the same way they were specified in the blacklist. For example, a WWID exception will not apply to devices specified by a **devnode** blacklist entry, even if the blacklisted device is associated with that WWID. Similarly, **devnode** exceptions apply only to **devnode** entries, and **device** exceptions apply only to device entries.

## 4.3. Configuration File Defaults

The `/etc/multipath.conf` configuration file includes a **defaults** section that sets the **user\_friendly\_names** parameter to **yes**, as follows.

```
defaults {
    user_friendly_names yes
}
```

This overwrites the default value of the **user\_friendly\_names** parameter.

The configuration file includes a template of configuration defaults. This section is commented out, as follows.

```
#defaults {
#    polling_interval          10
#    path_selector             "round-robin 0"
#    path_grouping_policy      multibus
#    uid_attribute             ID_SERIAL
#    prio                      alua
#    path_checker              readsector0
#    rr_min_io                 100
#    max_fds                   8192
#    rr_weight                 priorities
#    failback                  immediate
#    no_path_retry             fail
#    user_friendly_names       yes
#}
```

To overwrite the default value for any of the configuration parameters, you can copy the relevant line from this template into the **defaults** section and uncomment it. For example, to overwrite the **path\_grouping\_policy** parameter so that it is **multibus** rather than the default value of **failover**, copy the appropriate line from the template to the initial **defaults** section of the configuration file, and uncomment it, as follows.

```
defaults {
    user_friendly_names    yes
    path_grouping_policy    multibus
}
```

[Table 4.1, “Multipath Configuration Defaults”](#) describes the attributes that are set in the **defaults** section of the **multipath.conf** configuration file. These values are used by DM Multipath unless they are overwritten by the attributes specified in the **devices** and **multipaths** sections of the **multipath.conf** file.

**Table 4.1. Multipath Configuration Defaults**

Attribute	Description
<b>polling_interval</b>	Specifies the interval between two path checks in seconds. For properly functioning paths, the interval between checks will gradually increase to (4 * <b>polling_interval</b> ). The default value is 5.
<b>multipath_dir</b>	The directory where the dynamic shared objects are stored. The default value is system dependent, commonly <b>/lib/multipath</b> .

Attribute	Description
<b>find_multipaths</b>	<p>Defines the mode for setting up multipath devices. If this parameter is set to <b>yes</b>, then multipath will not try to create a device for every non-blacklisted path. Instead multipath will create a device only if one of three conditions are met:</p> <ul style="list-style-type: none"> <li>- There are at least two non-blacklisted paths with the same WWID.</li> <li>- The user manually forces the creation of the device by specifying a device with the <b>multipath</b> command.</li> <li>- A path has the same WWID as a multipath device that was previously created. Whenever a multipath device is created with <b>find_multipaths</b> set, multipath remembers the WWID of the device so that it will automatically create the device again as soon as it sees a path with that WWID. This allows you to have multipath automatically choose the correct paths to make into multipath devices, without having to edit the multipath blacklist. For instructions on the procedure to follow if you have previously created multipath devices when the <b>find_multipaths</b> parameter was not set, see <a href="#">Section 4.2, "Configuration File Blacklist"</a>.</li> </ul> <p>The default value is <b>no</b>. The default <b>multipath.conf</b> file created by <b>mpathconf</b>, however, will enable <b>find_multipaths</b> as of Red Hat Enterprise Linux 7.</p>
<b>verbosity</b>	<p>The default verbosity. Higher values increase the verbosity level. Valid levels are between 0 and 6. The default value is <b>2</b>.</p>
<b>path_selector</b>	<p>Specifies the default algorithm to use in determining what path to use for the next I/O operation. Possible values include:</p> <p><b>round-robin 0</b>: Loop through every path in the path group, sending the same amount of I/O to each.</p> <p><b>queue-length 0</b>: Send the next bunch of I/O down the path with the least number of outstanding I/O requests.</p> <p><b>service-time 0</b>: Send the next bunch of I/O down the path with the shortest estimated service time, which is determined by dividing the total size of the outstanding I/O to each path by its relative throughput.</p> <p>The default value is <b>round-robin 0</b>.</p>



Attribute	Description
<b>path_grouping_policy</b>	<p>Specifies the default path grouping policy to apply to unspecified multipaths. Possible values include:</p> <p><b>failover</b>: 1 path per priority group.</p> <p><b>multibus</b>: all valid paths in 1 priority group.</p> <p><b>group_by_serial</b>: 1 priority group per detected serial number.</p> <p><b>group_by_prio</b>: 1 priority group per path priority value. Priorities are determined by callout programs specified as global, per-controller, or per-multipath options.</p> <p><b>group_by_node_name</b>: 1 priority group per target node name. Target node names are fetched in <code>/sys/class/fc_transport/target*/node_name</code>.</p> <p>The default value is <b>failover</b>.</p>
<b>prio</b>	<p>Specifies the default function to call to obtain a path priority value. For example, the ALUA bits in SPC-3 provide an exploitable <b>prio</b> value. Possible values include:</p> <p><b>const</b>: Set a priority of 1 to all paths.</p> <p><b>emc</b>: Generate the path priority for EMC arrays.</p> <p><b>alua</b>: Generate the path priority based on the SCSI-3 ALUA settings.</p> <p><b>ontap</b>: Generate the path priority for NetApp arrays.</p> <p><b>rdac</b>: Generate the path priority for LSI/Engenio RDAC controller.</p> <p><b>hp_sw</b>: Generate the path priority for Compaq/HP controller in active/standby mode.</p> <p><b>hds</b>: Generate the path priority for Hitachi HDS Modular storage arrays.</p> <p>The default value is <b>const</b>.</p>

Attribute	Description
<b>features</b>	<p>The default extra features of multipath devices, using the format: "<i>number_of_features_plus_arguments feature1 ...</i>".</p> <p>Possible values for <b>features</b> include:</p> <p><b>queue_if_no_path</b>, which is the same as setting <b>no_path_retry</b> to <b>queue</b>. For information on issues that may arise when using this feature, see <a href="#">Section 5.7, "Issues with queue_if_no_path feature"</a>.</p> <p><b>retain_attached_hw_handler</b>: If this parameter is set to <b>yes</b> and the SCSI layer has already attached a hardware handler to the path device, multipath will not force the device to use the <b>hardware_handler</b> specified by the <b>multipath.conf</b> file. If the SCSI layer has not attached a hardware handler, multipath will continue to use its configured hardware handler as usual. The default value is <b>no</b>.</p> <p><b>pg_init_retries</b> <i>n</i>: Retry path group initialization up to <i>n</i> times before failing where <math>1 \leq n \leq 50</math>.</p> <p><b>pg_init_delay_msecs</b> <i>n</i>: Wait <i>n</i> milliseconds between path group initialization retries where <math>0 \leq n \leq 60000</math>.</p>
<b>path_checker</b>	<p>Specifies the default method used to determine the state of the paths. Possible values include:</p> <p><b>readsector0</b>: Read the first sector of the device.</p> <p><b>tur</b>: Issue a <b>TEST UNIT READY</b> command to the device.</p> <p><b>emc_clariion</b>: Query the EMC Clariion specific EVPD page 0xC0 to determine the path.</p> <p><b>hp_sw</b>: Check the path state for HP storage arrays with Active/Standby firmware.</p> <p><b>rdac</b>: Check the path state for LSI/Engenio RDAC storage controller.</p> <p><b>directio</b>: Read the first sector with direct I/O.</p> <p>The default value is <b>directio</b>.</p>

Attribute	Description
<b>failback</b>	<p>Manages path group failback.</p> <p>A value of <b>immediate</b> specifies immediate failback to the highest priority path group that contains active paths.</p> <p>A value of <b>manual</b> specifies that there should not be immediate failback but that failback can happen only with operator intervention.</p> <p>A value of <b>followover</b> specifies that automatic failback should be performed when the first path of a path group becomes active. This keeps a node from automatically failing back when another node requested the failover.</p> <p>A numeric value greater than zero specifies deferred failback, expressed in seconds.</p> <p>The default value is <b>manual</b>.</p>
<b>rr_min_io</b>	<p>Specifies the number of I/O requests to route to a path before switching to the next path in the current path group. This setting is only for systems running kernels older than 2.6.31. Newer systems should use <b>rr_min_io_rq</b>. The default value is 1000.</p>
<b>rr_min_io_rq</b>	<p>Specifies the number of I/O requests to route to a path before switching to the next path in the current path group, using request-based device-mapper-multipath. This setting should be used on systems running current kernels. On systems running kernels older than 2.6.31, use <b>rr_min_io</b>. The default value is 1.</p>
<b>rr_weight</b>	<p>If set to <b>priorities</b>, then instead of sending <b>rr_min_io</b> requests to a path before calling <b>path_selector</b> to choose the next path, the number of requests to send is determined by <b>rr_min_io</b> times the path's priority, as determined by the <b>prio</b> function. If set to <b>uniform</b>, all path weights are equal. The default value is <b>uniform</b>.</p>
<b>no_path_retry</b>	<p>A numeric value for this attribute specifies the number of times the system should attempt to use a failed path before disabling queuing.</p> <p>A value of <b>fail</b> indicates immediate failure, without queuing.</p> <p>A value of <b>queue</b> indicates that queuing should not stop until the path is fixed.</p> <p>The default value is 0.</p>
<b>user_friendly_names</b>	<p>If set to <b>yes</b>, specifies that the system should use the <b>/etc/multipath/bindings</b> file to assign a persistent and unique alias to the multipath, in the form of <b>mpathn</b>. If set to <b>no</b>, specifies that the system should use the WWID as the alias for the multipath. In either case, what is specified here will be overridden by any device-specific aliases you specify in the <b>multipaths</b> section of the configuration file. The default value is <b>no</b>.</p>
<b>queue_without_daemon</b>	<p>If set to <b>no</b>, the <b>multipathd</b> daemon will disable queuing for all devices when it is shut down. The default value is <b>no</b>.</p>

Attribute	Description
<b>flush_on_last_del</b>	If set to <b>yes</b> , the <b>multipathd</b> daemon will disable queuing when the last path to a device has been deleted. The default value is <b>no</b> .
<b>max_fds</b>	Sets the maximum number of open file descriptors that can be opened by multipath and the <b>multipathd</b> daemon. This is equivalent to the <b>ulimit -n</b> command. As of the Red Hat Enterprise Linux 6.3 release, the default value is <b>max</b> , which sets this to the system limit from <b>/proc/sys/fs/nr_open</b> . For earlier releases, if this is not set the maximum number of open file descriptors is taken from the calling process; it is usually 1024. To be safe, this should be set to the maximum number of paths plus 32, if that number is greater than 1024.
<b>checker_timeout</b>	The timeout to use for prioritizers and path checkers that issue SCSI commands with an explicit timeout, in seconds. The default value is taken from <b>sys/block/sdx/device/timeout</b> .
<b>fast_io_fail_tmo</b>	The number of seconds the SCSI layer will wait after a problem has been detected on an FC remote port before failing I/O to devices on that remote port. This value should be smaller than the value of <b>dev_loss_tmo</b> . Setting this to <b>off</b> will disable the timeout. The default value is determined by the OS.
<b>dev_loss_tmo</b>	The number of seconds the SCSI layer will wait after a problem has been detected on an FC remote port before removing it from the system. Setting this to infinity will set this to 2147483647 seconds, or 68 years. The default value is determined by the OS.
<b>hwtable_string_match</b>	<p>Each device configuration in the <b>devices</b> section of the <b>multipath.conf</b> file will either create its own device configuration or it will modify one of the built-in device configurations. If <b>hwtable_string_match</b> is set to <b>yes</b>, then if the vendor, product, and revision strings in a user's device configuration exactly match those strings in a built-in device configuration, the built-in configuration is modified by the options in the user's configuration. Otherwise, the user's device configuration is treated as a new configuration. If <b>hwtable_string_match</b> is set to <b>no</b>, a regular expression match is used instead of a string match.</p> <p>The <b>hwtable_string_match</b> is set to <b>no</b> by default.</p>
<b>retain_attached_hw_handler</b>	If this parameter is set to <b>yes</b> and the SCSI layer has already attached a hardware handler to the path device, multipath will not force the device to use the <b>hardware_handler</b> specified by the <b>multipath.conf</b> file. If the SCSI layer has not attached a hardware handler, multipath will continue to use its configured hardware handler as usual. The default value is <b>no</b> .
<b>detect_prio</b>	If this is set to <b>yes</b> , multipath will first check if the device supports ALUA, and if so it will automatically assign the device the <b>alua</b> prioritizer. If the device doesn't support ALUA, it will determine the prioritizer as it always does. The default value is <b>no</b> .
<b>reload_readwrite</b>	If this is set to <b>yes</b> , the <b>multipathd</b> daemon will listen to path device change events and if a device has become read/write it will reload the multipath device.
<b>uid_attribute</b>	Provides a unique path identifier. The default value is <b>ID_SERIAL</b> .

Attribute	Description
<b>force_sync</b>	(Red Hat Enterprise Linux Release 7.1 and later) If this is set to "yes", it prevents path checkers from running in async mode.
<b>delay_watch_checks</b>	(Red Hat Enterprise Linux Release 7.2 and later) If set to a value greater than 0, the <b>multipathd</b> daemon will watch paths that have recently become valid for the specified number of checks. If they fail again while they are being watched, when they next become valid they will not be used until they have stayed up for the number of consecutive checks specified with <b>delay_wait_checks</b> . This allows you to keep paths that may be unreliable from immediately being put back into use as soon as they come back online. The default value is <b>no</b> .
<b>delay_wait_checks</b>	(Red Hat Enterprise Linux Release 7.2 and later) If set to a value greater than 0, when a device that has recently come back online fails again within the number of checks specified with <b>delay_watch_checks</b> , the next time it comes back online it will be marked and delayed and it will not be used until it has passed the number of checks specified in <b>delay_wait_checks</b> . The default value is <b>no</b> .
<b>ignore_new_boot_devs</b>	(Red Hat Enterprise Linux Release 7.2 and later) If set to <b>yes</b> , when the node is still in the <b>initramfs</b> file system during early boot, multipath will not create any devices whose WWIDs do not already exist in the <b>initramfs copy</b> of the <b>/etc/multipath/wwids</b> . This feature can be used for booting up during installation, when multipath would otherwise attempt to set itself up on devices that it did not claim when they first appeared via the <b>udev</b> rules. This parameter can be set to <b>yes</b> or <b>no</b> . If unset, it defaults to <b>no</b> .
<b>retrigger_tries</b> , <b>retrigger_delay</b>	(Red Hat Enterprise Linux Release 7.2 and later) The <b>retrigger_tries</b> and <b>retrigger_delay</b> parameters are used in conjunction to make <b>multipathd</b> retrigger uevents if <b>udev</b> failed to completely process the original ones, leaving the device unusable by multipath. The <b>retrigger_tries</b> parameter sets the number of times that multipath will try to retrigger a <b>uevent</b> if a device has not been completely set up. The <b>retrigger_delay</b> parameter sets the number of seconds between retries. Both of these options accept numbers greater than or equal to zero. Setting the <b>retrigger_tries</b> parameter to zero disables retries. Setting the <b>retrigger_delay</b> parameter to zero causes the <b>uevent</b> to be reissued on the next loop of the path checker. If the <b>retrigger_tries</b> parameter is unset, it defaults to 3. If the <b>retrigger_delay</b> parameter is unset, it defaults to 10.

Attribute	Description
<b>new_bindings_in_boot</b>	(Red Hat Enterprise Linux Release 7.2 and later) The <b>new_bindings_in_boot</b> parameter is used to keep multipath from giving out a <b>user_friendly_name</b> in the <b>initramfs</b> file system that was already given out by the bindings file in the regular file system, an issue that can arise since the <b>user_friendly_names</b> bindings in the <b>initramfs</b> file system get synced with the bindings in the regular file system only when the <b>initramfs</b> file system is remade. When this parameter is set to <b>no</b> multipath will not create any new bindings in the <b>initramfs</b> file system. If a device does not already have a binding in the <b>initramfs</b> copy of <b>/etc/multipath/bindings</b> , multipath will use its WWID as an alias instead of giving it a <b>user_friendly_name</b> . Later in boot, after the node has mounted the regular filesystem, multipath will give out a <b>user_friendly_name</b> to the device. This parameter can be set to <b>yes</b> or <b>no</b> . If unset, it defaults to <b>no</b> .
<b>deferred_remove</b>	If set to <b>yes</b> , multipathd will do a deferred remove instead of a regular remove when the last path device has been deleted. This ensures that if a multipathed device is in use when a regular remove is performed and the remove fails, the device will automatically be removed when the last user closes the device. The default value is <b>no</b> .
<b>log_checker_err</b>	If set to <b>once</b> , multipathd logs the first path checker error at verbosity level 2. Any later errors are logged at verbosity level 3 until the device is restored. If it is set to <b>always</b> , multipathd always logs the path checker error at verbosity level 2. The default value is <b>always</b> .

## 4.4. Multipaths Device Configuration Attributes

[Table 4.2, “Multipath Attributes”](#) shows the attributes that you can set in the **multipaths** section of the **multipath.conf** configuration file for each specific multipath device. These attributes apply only to the one specified multipath. These defaults are used by DM Multipath and override attributes set in the **defaults** and **devices** sections of the **multipath.conf** file.

**Table 4.2. Multipath Attributes**

Attribute	Description
<b>wwid</b>	Specifies the WWID of the multipath device to which the <b>multipath</b> attributes apply. This parameter is mandatory for this section of the <b>multipath.conf</b> file.
<b>alias</b>	Specifies the symbolic name for the multipath device to which the <b>multipath</b> attributes apply. If you are using <b>user_friendly_names</b> , do not set this value to <b>mpathn</b> ; this may conflict with an automatically assigned user friendly name and give you incorrect device node names.

Attribute	Description
<b>path_grouping_policy</b>	<p>Specifies the default path grouping policy to apply to unspecified multipaths. Possible values include:</p> <p><b>failover</b> = 1 path per priority group</p> <p><b>multibus</b> = all valid paths in 1 priority group</p> <p><b>group_by_serial</b> = 1 priority group per detected serial number</p> <p><b>group_by_prio</b> = 1 priority group per path priority value</p> <p><b>group_by_node_name</b> = 1 priority group per target node name</p>
<b>path_selector</b>	<p>Specifies the default algorithm to use in determining what path to use for the next I/O operation. Possible values include:</p> <p><b>round-robin 0</b>: Loop through every path in the path group, sending the same amount of I/O to each.</p> <p><b>queue-length 0</b>: Send the next bunch of I/O down the path with the least number of outstanding I/O requests.</p> <p><b>service-time 0</b>: Send the next bunch of I/O down the path with the shortest estimated service time, which is determined by dividing the total size of the outstanding I/O to each path by its relative throughput.</p>
<b>failback</b>	<p>Manages path group failback.</p> <p>A value of <b>immediate</b> specifies immediate failback to the highest priority path group that contains active paths.</p> <p>A value of <b>manual</b> specifies that there should not be immediate failback but that failback can happen only with operator intervention.</p> <p>A value of <b>followover</b> specifies that automatic failback should be performed when the first path of a path group becomes active. This keeps a node from automatically failing back when another node requested the failover.</p> <p>A numeric value greater than zero specifies deferred failback, expressed in seconds.</p>

Attribute	Description
<b>prio</b>	<p>Specifies the default function to call to obtain a path priority value. For example, the ALUA bits in SPC-3 provide an exploitable <b>prio</b> value. Possible values include:</p> <p><b>const</b>: Set a priority of 1 to all paths.</p> <p><b>emc</b>: Generate the path priority for EMC arrays.</p> <p><b>alua</b>: Generate the path priority based on the SCSI-3 ALUA settings.</p> <p><b>ontap</b>: Generate the path priority for NetApp arrays.</p> <p><b>rdac</b>: Generate the path priority for LSI/Engenio RDAC controller.</p> <p><b>hp_sw</b>: Generate the path priority for Compaq/HP controller in active/standby mode.</p> <p><b>hds</b>: Generate the path priority for Hitachi HDS Modular storage arrays.</p>
<b>no_path_retry</b>	<p>A numeric value for this attribute specifies the number of times the system should attempt to use a failed path before disabling queuing.</p> <p>A value of <b>fail</b> indicates immediate failure, without queuing.</p> <p>A value of <b>queue</b> indicates that queuing should not stop until the path is fixed.</p>
<b>rr_min_io</b>	<p>Specifies the number of I/O requests to route to a path before switching to the next path in the current path group. This setting is only for systems running kernels older than 2.6.31. Newer systems should use <b>rr_min_io_rq</b>. The default value is 1000.</p>
<b>rr_min_io_rq</b>	<p>Specifies the number of I/O requests to route to a path before switching to the next path in the current path group, using request-based device-mapper-multipath. This setting should be used on systems running current kernels. On systems running kernels older than 2.6.31, use <b>rr_min_io</b>. The default value is 1.</p>
<b>rr_weight</b>	<p>If set to <b>priorities</b>, then instead of sending <b>rr_min_io</b> requests to a path before calling <b>path_selector</b> to choose the next path, the number of requests to send is determined by <b>rr_min_io</b> times the path's priority, as determined by the <b>prio</b> function. If set to <b>uniform</b>, all path weights are equal.</p>
<b>flush_on_last_del</b>	<p>If set to <b>yes</b>, then multipath will disable queuing when the last path to a device has been deleted.</p>



Attribute	Description
<b>user_friendly_names</b>	If set to <b>yes</b> , specifies that the system should use the <code>/etc/multipath/bindings</code> file to assign a persistent and unique alias to the multipath, in the form of <code>mpathn</code> . If set to <b>no</b> , specifies that the system should use the WWID as the alias for the multipath. In either case, what is specified here will be overridden by any device-specific aliases you specify in the <code>multipaths</code> section of the configuration file.
<b>delay_watch_checks</b>	(Red Hat Enterprise Linux Release 7.2 and later) If set to a value greater than 0, the <code>multipathd</code> daemon will watch paths that have recently become valid for the specified number of checks. If they fail again while they are being watched, when they next become valid they will not be used until they have stayed up for the number of consecutive checks specified with <code>delay_wait_checks</code> . This allows you to keep paths that may be unreliable from immediately being put back into use as soon as they come back online.
<b>delay_wait_checks</b>	(Red Hat Enterprise Linux Release 7.2 and later) If set to a value greater than 0, when a device that has recently come back online fails again within the number of checks specified with <code>delay_watch_checks</code> , the next time it comes back online it will be marked and delayed and it will not be used until it has passed the number of checks specified in <code>delay_wait_checks</code> .
<b>config_dir</b>	(Red Hat Enterprise Linux Release 7.2 and later) If set to anything other than "", multipath will search this directory alphabetically for files ending in ".conf" and it will read configuration information from them, just as if the information were in the <code>/etc/multipath.conf</code> file. This allows you to have one main configuration that you share between machines in addition to a separate machine-specific configuration file or files. The <code>config_dir</code> parameter must either be "" or a fully qualified directory name. This parameter can be set only in the main <code>/etc/multipath.conf</code> file and not in one of the files specified in the <code>config_dir</code> file itself. The default value is <code>/etc/multipath/conf.d</code> .
<b>deferred_remove</b>	If set to <b>yes</b> , multipathd will do a deferred remove instead of a regular remove when the last path device has been deleted. This ensures that if a multipathed device is in use when a regular remove is performed and the remove fails, the device will automatically be removed when the last user closes the device.

The following example shows multipath attributes specified in the configuration file for two specific multipath devices. The first device has a WWID of **3600508b4000156d70001200000b0000** and a symbolic name of **yellow**.

The second multipath device in the example has a WWID of **1DEC\_\_\_\_\_321816758474** and a symbolic name of **red**. In this example, the `rr_weight` attributes is set to **priorities**.

```

multipaths {
    multipath {
        wwid                3600508b4000156d70001200000b0000
        alias                yellow
        path_grouping_policy multibus
        path_selector        "round-robin 0"
        failback             manual
    }
}

```

```

        rr_weight          priorities
        no_path_retry      5
    }
    multipath {
        wwid                1DEC_____321816758474
        alias               red
        rr_weight           priorities
    }
}

```

## 4.5. Configuration File Devices

[Table 4.3, “Device Attributes”](#) shows the attributes that you can set for each individual storage device in the **devices** section of the **multipath.conf** configuration file. These attributes are used by DM Multipath unless they are overwritten by the attributes specified in the **multipaths** section of the **multipath.conf** file for paths that contain the device. These attributes override the attributes set in the **defaults** section of the **multipath.conf** file.

Many devices that support multipathing are included by default in a multipath configuration. For information on the default configuration value, including supported devices, run either of the following commands.

```

# multipathd show config
# multipath -t

```

You probably will not need to modify the values for these devices, but if you do you can overwrite the default values by including an entry in the configuration file for the device that overwrites those values. You can copy the device configuration defaults for the device that the **multipathd show config** command displays and override the values that you want to change.

To add a device that is not configured automatically by default to this section of the configuration file, you need to set the **vendor** and **product** parameters. You can find these values by looking at **/sys/block/device\_name/device/vendor** and **/sys/block/device\_name/device/model** where *device\_name* is the device to be multipathed, as in the following example:

```

# cat /sys/block/sda/device/vendor
WINSYS
# cat /sys/block/sda/device/model
SF2372

```

The additional parameters to specify depend on your specific device. If the device is active/active, you will usually not need to set additional parameters. You may want to set **path\_grouping\_policy** to **multibus**. Other parameters you may need to set are **no\_path\_retry** and **rr\_min\_io**, as described in [Table 4.3, “Device Attributes”](#).

If the device is active/passive, but it automatically switches paths with I/O to the passive path, you need to change the checker function to one that does not send I/O to the path to test if it is working (otherwise, your device will keep failing over). This almost always means that you set the **path\_checker** to **tur**; this works for all SCSI devices that support the Test Unit Ready command, which most do.

If the device needs a special command to switch paths, then configuring this device for multipath requires a hardware handler kernel module. The current available hardware handler is **emc**. If this is not sufficient for your device, you may not be able to configure the device for multipath.

Table 4.3. Device Attributes

Attribute	Description
<b>vendor</b>	Specifies the vendor name of the storage device to which the <b>device</b> attributes apply, for example <b>COMPAQ</b> .
<b>product</b>	Specifies the product name of the storage device to which the <b>device</b> attributes apply, for example <b>HSV110 (C)COMPAQ</b> .
<b>revision</b>	Specifies the product revision identifier of the storage device.
<b>product_blacklist</b>	Specifies a regular expression used to blacklist devices by product.
<b>hardware_handler</b>	<p>Specifies a module that will be used to perform hardware specific actions when switching path groups or handling I/O errors. Possible values include:</p> <ul style="list-style-type: none"> <li><b>1 emc</b>: hardware handler for EMC storage arrays.</li> <li><b>1 alua</b>: hardware handler for SCSI-3 ALUA arrays.</li> <li><b>1 hp_sw</b>: hardware handler for Compaq/HP controllers.</li> <li><b>1 rdac</b>: hardware handler for the LSI/Engenio RDAC controllers.</li> </ul>
<b>path_grouping_policy</b>	<p>Specifies the default path grouping policy to apply to unspecified multipaths. Possible values include:</p> <ul style="list-style-type: none"> <li><b>failover</b> = 1 path per priority group</li> <li><b>multibus</b> = all valid paths in 1 priority group</li> <li><b>group_by_serial</b> = 1 priority group per detected serial number</li> <li><b>group_by_prio</b> = 1 priority group per path priority value</li> <li><b>group_by_node_name</b> = 1 priority group per target node name</li> </ul>
<b>path_selector</b>	<p>Specifies the default algorithm to use in determining what path to use for the next I/O operation. Possible values include:</p> <ul style="list-style-type: none"> <li><b>round-robin 0</b>: Loop through every path in the path group, sending the same amount of I/O to each.</li> <li><b>queue-length 0</b>: Send the next bunch of I/O down the path with the least number of outstanding I/O requests.</li> <li><b>service-time 0</b>: Send the next bunch of I/O down the path with the shortest estimated service time, which is determined by dividing the total size of the outstanding I/O to each path by its relative throughput.</li> </ul>

Attribute	Description
<b>path_checker</b>	<p>Specifies the default method used to determine the state of the paths. Possible values include:</p> <p><b>readsector0</b>: Read the first sector of the device.</p> <p><b>tur</b>: Issue a <b>TEST UNIT READY</b> to the device.</p> <p><b>emc_clariion</b>: Query the EMC Clariion specific EVPD page 0xC0 to determine the path.</p> <p><b>hp_sw</b>: Check the path state for HP storage arrays with Active/Standby firmware.</p> <p><b>rdac</b>: Check the path stat for LSI/Engenio RDAC storage controller.</p> <p><b>directio</b>: Read the first sector with direct I/O.</p>
<b>features</b>	<p>The default extra features of multipath devices, using the format: "<i>number_of_features_plus_arguments feature1 ...</i>".</p> <p>Possible values for <b>features</b> include:</p> <p><b>queue_if_no_path</b>, which is the same as setting <b>no_path_retry</b> to <b>queue</b>. For information on issues that may arise when using this feature, see <a href="#">Section 5.7, "Issues with queue_if_no_path feature"</a>.</p> <p><b>retain_attached_hw_handler</b>: If this parameter is set to <b>yes</b> and the SCSI layer has already attached a hardware handler to the path device, multipath will not force the device to use the <b>hardware_handler</b> specified by the <b>multipath.conf</b> file. If the SCSI layer has not attached a hardware handler, multipath will continue to use its configured hardware handler as usual.</p> <p><b>pg_init_retries</b> <i>n</i>: Retry path group initialization up to <i>n</i> times before failing where <math>1 \leq n \leq 50</math>.</p> <p><b>pg_init_delay_msecs</b> <i>n</i>: Wait <i>n</i> milliseconds between path group initialization retries where <math>0 \leq n \leq 60000</math>.</p>

Attribute	Description
<b>prio</b>	<p>Specifies the default function to call to obtain a path priority value. For example, the ALUA bits in SPC-3 provide an exploitable <b>prio</b> value. Possible values include:</p> <p><b>const</b>: Set a priority of 1 to all paths.</p> <p><b>emc</b>: Generate the path priority for EMC arrays.</p> <p><b>alua</b>: Generate the path priority based on the SCSI-3 ALUA settings.</p> <p><b>ontap</b>: Generate the path priority for NetApp arrays.</p> <p><b>rdac</b>: Generate the path priority for LSI/Engenio RDAC controller.</p> <p><b>hp_sw</b>: Generate the path priority for Compaq/HP controller in active/standby mode.</p> <p><b>hds</b>: Generate the path priority for Hitachi HDS Modular storage arrays.</p>
<b>failback</b>	<p>Manages path group failback.</p> <p>A value of <b>immediate</b> specifies immediate failback to the highest priority path group that contains active paths.</p> <p>A value of <b>manual</b> specifies that there should not be immediate failback but that failback can happen only with operator intervention.</p> <p>A value of <b>followover</b> specifies that automatic failback should be performed when the first path of a path group becomes active. This keeps a node from automatically failing back when another node requested the failover.</p> <p>A numeric value greater than zero specifies deferred failback, expressed in seconds.</p>
<b>rr_weight</b>	<p>If set to <b>priorities</b>, then instead of sending <b>rr_min_io</b> requests to a path before calling <b>path_selector</b> to choose the next path, the number of requests to send is determined by <b>rr_min_io</b> times the path's priority, as determined by the <b>prio</b> function. If set to <b>uniform</b>, all path weights are equal.</p>
<b>no_path_retry</b>	<p>A numeric value for this attribute specifies the number of times the system should attempt to use a failed path before disabling queuing.</p> <p>A value of <b>fail</b> indicates immediate failure, without queuing.</p> <p>A value of <b>queue</b> indicates that queuing should not stop until the path is fixed.</p>

Attribute	Description
<b>rr_min_io</b>	Specifies the number of I/O requests to route to a path before switching to the next path in the current path group. This setting is only for systems running kernels older than 2.6.31. Newer systems should use <b>rr_min_io_rq</b> . The default value is 1000.
<b>rr_min_io_rq</b>	Specifies the number of I/O requests to route to a path before switching to the next path in the current path group, using request-based device-mapper-multipath. This setting should be used on systems running current kernels. On systems running kernels older than 2.6.31, use <b>rr_min_io</b> . The default value is 1.
<b>fast_io_fail_tmo</b>	The number of seconds the SCSI layer will wait after a problem has been detected on an FC remote port before failing I/O to devices on that remote port. This value should be smaller than the value of <b>dev_loss_tmo</b> . Setting this to <b>off</b> will disable the timeout.
<b>dev_loss_tmo</b>	The number of seconds the SCSI layer will wait after a problem has been detected on an FC remote port before removing it from the system. Setting this to infinity will set this to 2147483647 seconds, or 68 years.
<b>flush_on_last_del</b>	If set to <b>yes</b> , the <b>multipathd</b> daemon will disable queuing when the last path to a device has been deleted.
<b>user_friendly_names</b>	If set to <b>yes</b> , specifies that the system should use the <b>/etc/multipath/bindings</b> file to assign a persistent and unique alias to the multipath, in the form of <b>mpathn</b> . If set to <b>no</b> , specifies that the system should use the WWID as the alias for the multipath. In either case, what is specified here will be overridden by any device-specific aliases you specify in the <b>multipaths</b> section of the configuration file. The default value is <b>no</b> .
<b>retain_attached_hw_handler</b>	If this parameter is set to <b>yes</b> and the SCSI layer has already attached a hardware handler to the path device, multipath will not force the device to use the <b>hardware_handler</b> specified by the <b>multipath.conf</b> file. If the SCSI layer has not attached a hardware handler, multipath will continue to use its configured hardware handler as usual. The default value is <b>no</b> .
<b>detect_prio</b>	If this is set to <b>yes</b> , multipath will first check if the device supports ALUA, and if so it will automatically assign the device the <b>alua</b> prioritizer. If the device doesn't support ALUA, it will determine the prioritizer as it always does.
<b>delay_watch_checks</b>	(Red Hat Enterprise Linux Release 7.2 and later) If set to a value greater than 0, the <b>multipathd</b> daemon will watch paths that have recently become valid for the specified number of checks. If they fail again while they are being watched, when they next become valid they will not be used until they have stayed up for the number of consecutive checks specified with <b>delay_wait_checks</b> . This allows you to keep paths that may be unreliable from immediately being put back into use as soon as they come back online.
<b>delay_wait_checks</b>	(Red Hat Enterprise Linux Release 7.2 and later) If set to a value greater than 0, when a device that has recently come back online fails again within the number of checks specified with <b>delay_watch_checks</b> , the next time it comes back online it will be marked and delayed and it will not be used until it has passed the number of checks specified in <b>delay_wait_checks</b> .

Attribute	Description
<b>deferred_remove</b>	If set to <b>yes</b> , multipathd will do a deferred remove instead of a regular remove when the last path device has been deleted. This ensures that if a multipathed device is in use when a regular remove is performed and the remove fails, the device will automatically be removed when the last user closes the device.

The following example shows a **device** entry in the multipath configuration file.

```
# }
# device {
#   vendor    "COMPAQ  "
#   product   "MSA1000      "
#   path_grouping_policy multibus
#   path_checker tur
#   rr_weight priorities
# }
#}
```

## Chapter 5. DM Multipath Administration and Troubleshooting

This chapter provides information on administering DM Multipath on a running system.

### 5.1. Automatic Configuration File Generation with Multipath Helper

You can generate a basic configuration for multipath devices on Red Hat Enterprise Linux with the Multipath Helper application. The application gives you options to create multipath configurations with custom aliases, device blacklists, and settings for the characteristics of individual multipath devices. Upon completion, the application generates an installation script that includes the configuration parameters you selected and it provides a `multipath.conf` configuration file for review.

The Multipath Helper application can be found at <https://access.redhat.com/labsinfo/multipathhelper>.

### 5.2. Resizing an Online Multipath Device

If you need to resize an online multipath device, use the following procedure.

1. Resize your physical device.
2. Execute the following command to find the paths to the LUN:

```
# multipath -l
```

3. Resize your paths. For SCSI devices, writing a 1 to the `rescan` file for the device causes the SCSI driver to rescan, as in the following command:

```
# echo 1 > /sys/block/device_name/device/rescan
```

4. Resize your multipath device by executing the `multipathd resize` command:

```
# multipathd -k'resize map mpatha'
```

5. Resize the file system (assuming no LVM or DOS partitions are used):

```
# resize2fs /dev/mapper/mpatha
```

### 5.3. Moving root File Systems from a Single Path Device to a Multipath Device

If you have installed your system on a single-path device and later add another path to the root file system, you will need to move your root file system to a multipathed device. This section documents the procedure for moving from a single-path to a multipathed device.

After ensuring that you have installed the `device-mapper-multipath` package, perform the following procedure:

1. Execute the following command to create the `/etc/multipath.conf` configuration file, load the multipath module, and set `chkconfig` for the `multipathd` to `on`:



```
# mpathconf --enable
```

For further information on using the `mpathconf` command to set up multipathing, see [Section 3.1, “Setting Up DM Multipath”](#).

- If the `find_multipaths` configuration parameter is not set to `yes`, edit the `blacklist` and `blacklist_exceptions` sections of the `/etc/multipath.conf` file, as described in [Section 4.2, “Configuration File Blacklist”](#).
- In order for multipath to build a multipath device on top of the root device as soon as it is discovered, run the following command. This command also ensures that `find_multipaths` will allow the device, even if it only has one path.

```
# multipath -a root_devname
```

For example, if the root device is `/dev/sdb`, run the following command.

```
# multipath -a /dev/sdb
wwid '3600d02300069c9ce09d41c4ac9c53200' added
```

- To confirm that your configuration file is set up correctly, you can run the `multipath` command and search the output for a line of the following format. This indicates that the command failed to create the multipath device.

```
date wwid: ignoring map
```

For example, if the WWID of the device is `3600d02300069c9ce09d41c4ac9c53200`, you would see a line in the output such as the following:

```
# multipath
Oct 21 09:37:19 | 3600d02300069c9ce09d41c4ac9c53200: ignoring map
```

- To rebuild the `initramfs` file system with `multipath`, execute the `dracut` command with the following options:

```
# dracut --force -H --add multipath
```

- Shut the machine down.
- Configure the FC switch so that other paths are visible to the machine.
- Boot the machine.
- Check whether the root file system (`/`) is on the multipathed device.

## 5.4. Moving swap File Systems from a Single Path Device to a Multipath Device

By default, swap devices are set up as logical volumes. This does not require any special procedure for configuring them as multipath devices as long as you set up multipathing on the physical volumes that constitute the logical volume group. If your swap device is not an LVM volume, however, and it is mounted by device name, you may need to edit the `/etc/fstab` file to switch to the appropriate multipath device name.

1. Determine the WWID number of the swap device by running the `/sbin/multipath` command with the `-v3` option. The output from the command should show the swap device in the paths list.

You should look in the command output for a line of the following format, showing the swap device:

```
WWID H:B:T:L devname MAJOR:MINOR
```

For example, if your swap file system is set up on `sda` or one of its partitions, you would see a line in the output such as the following:

```
==== paths list ====
...
1ATA      WDC WD800JD-75MSA3          WD-WMAM9F
1:0:0:0 sda 8:0
...
```

2. Set up an alias for the swap device in the `/etc/multipath.conf` file:

```
multipaths {
  multipath {
    wwid WWID_of_swap_device
    alias swapdev
  }
}
```

3. Edit the `/etc/fstab` file and replace the old device path to the root device with the multipath device.

For example, if you had the following entry in the `/etc/fstab` file:

```
/dev/sda2 swap                ext4    defaults    0 0
```

You would change the entry to the following:

```
/dev/mapper/swapdev swap                ext4    defaults    0 0
```

## 5.5. The Multipath Daemon

If you find you have trouble implementing a multipath configuration, you should ensure that the multipath daemon is running, as described in [Chapter 3, Setting Up DM Multipath](#). The `multipathd` daemon must be running in order to use multipathed devices.

## 5.6. Issues with Large Number of LUNs

When a large number of LUNs are added to a node, using multipathed devices can significantly increase the time it takes for the `udev` device manager to create device nodes for them. If you experience this problem, you can correct it by deleting the following line in `/etc/udev/rules.d/40-multipath.rules`:

```
KERNEL!="dm-[0-9]*", ACTION=="add", PROGRAM=="bin/bash -c '/sbin/lsmode | /bin/grep ^dm_multipath'", RUN+="/sbin/multipath -v0 %M:%m"
```

This line causes the **udev** device manager to run **multipath** every time a block device is added to the node. Even with this line removed, the **multipathd** daemon will still automatically create multipathed devices, and **multipath** will still be called during the boot process for nodes with multipathed root file systems. The only change is that multipathed devices will not be automatically created when the **multipathd** daemon is not running, which should not be a problem for the vast majority of multipath users.

## 5.7. Issues with `queue_if_no_path` feature

If a multipath device is configured with features **"1 queue\_if\_no\_path"**, then any process that issues I/O will hang until one or more paths are restored. To avoid this, set the **no\_path\_retry *N*** parameter in the `/etc/multipath.conf` file (where *N* is the number of times the system should retry a path).

If you need to use the features **"1 queue\_if\_no\_path"** option and you experience the issue noted here, use the **dmsetup** command to edit the policy at runtime for a particular LUN (that is, for which all the paths are unavailable). For example, if you want to change the policy on the multipath device `mpathc` from **"queue\_if\_no\_path"** to **"fail\_if\_no\_path"**, execute the following command.

```
dmsetup message mpathc 0 "fail_if_no_path"
```

Note that you must specify the `mpathn` alias rather than the path.

## 5.8. Multipath Command Output

When you create, modify, or list a multipath device, you get a display of the current device setup. The format is as follows.

For each multipath device:

```
action_if_any: alias (wwid_if_different_from_alias)
dm_device_name_if_known vendor,product size=size features='features'
hwandler='hardware_handler' wp=write_permission_if_known
```

For each path group:

```
-+- policy='scheduling_policy' prio=prio_if_known
status=path_group_status_if_known
```

For each path:

```
`- host:channel:id:lun devnode major:minor dm_status_if_known
path_status online_status
```

For example, the output of a multipath command might appear as follows:

```
3600d0230000000000e13955cc3757800 dm-1 WINSYS,SF2372
size=269G features='0' hwandler='0' wp=rw
```

```
|+- policy='round-robin 0' prio=1 status=active
| ` 6:0:0:0 sdb 8:16 active ready running
`+- policy='round-robin 0' prio=1 status=enabled
  ` 7:0:0:0 sdf 8:80 active ready running
```

If the path is up and ready for I/O, the status of the path is **ready** or **ghost**. If the path is down, the status is **faulty** or **shaky**. The path status is updated periodically by the **multipathd** daemon based on the polling interval defined in the **/etc/multipath.conf** file.

The dm status is similar to the path status, but from the kernel's point of view. The dm status has two states: **failed**, which is analogous to **faulty**, and **active** which covers all other path states. Occasionally, the path state and the dm state of a device will temporarily not agree.

The possible values for **online\_status** are **running** and **offline**. A status of **offline** means that this SCSI device has been disabled.



### Note

When a multipath device is being created or modified, the path group status, the dm device name, the write permissions, and the dm status are not known. Also, the features are not always correct.

## 5.9. Multipath Queries with multipath Command

You can use the **-l** and **-ll** options of the **multipath** command to display the current multipath configuration. The **-l** option displays multipath topology gathered from information in **sysfs** and the device mapper. The **-ll** option displays the information the **-l** option displays in addition to all other available components of the system.

When displaying the multipath configuration, there are three verbosity levels you can specify with the **-v** option of the **multipath** command. Specifying **-v0** yields no output. Specifying **-v1** outputs the created or updated multipath names only, which you can then feed to other tools such as **kpartx**. Specifying **-v2** prints all detected paths, multipaths, and device maps.

The following example shows the output of a **multipath -l** command.

```
# multipath -l
3600d0230000000000e13955cc3757800 dm-1 WINSYS,SF2372
size=269G features='0' hwhandler='0' wp=rw
|+- policy='round-robin 0' prio=1 status=active
| ` 6:0:0:0 sdb 8:16 active ready running
`+- policy='round-robin 0' prio=1 status=enabled
  ` 7:0:0:0 sdf 8:80 active ready running
```

The following example shows the output of a **multipath -ll** command.

```
# multipath -ll
3600d0230000000000e13955cc3757801 dm-10 WINSYS,SF2372
size=269G features='0' hwhandler='0' wp=rw
|+- policy='round-robin 0' prio=1 status=enabled
| ` 19:0:0:1 sdc 8:32 active ready running
`+- policy='round-robin 0' prio=1 status=enabled
  ` 18:0:0:1 sdh 8:112 active ready running
```

```
3600d0230000000000e13955cc3757803 dm-2 WINSYS,SF2372
size=125G features='0' hwhandler='0' wp=rw
`-+- policy='round-robin 0' prio=1 status=active
  |- 19:0:0:3 sde 8:64 active ready running
  `-- 18:0:0:3 sdj 8:144 active ready running
```

## 5.10. Multipath Command Options

[Table 5.1, “Useful multipath Command Options”](#) describes some options of the `multipath` command that you may find useful.

**Table 5.1. Useful multipath Command Options**

Option	Description
<code>-l</code>	Display the current multipath configuration gathered from <code>sysfs</code> and the device mapper.
<code>-ll</code>	Display the current multipath configuration gathered from <code>sysfs</code> , the device mapper, and all other available components on the system.
<code>-f device</code>	Remove the named multipath device.
<code>-F</code>	Remove all unused multipath devices.
<code>-w device</code>	Remove the <code>wwid</code> of the specified device from the <code>wwids</code> file.
<code>-W</code>	Reset the <code>wwids</code> file to include only the current multipath devices.

Some `multipathd` commands include a `format` option followed by a wildcard. You can display a list of available wildcards with the following command.

```
# multipathd show wildcards
```

## 5.11. Determining Device Mapper Entries with the dmsetup Command

You can use the `dmsetup` command to find out which device mapper entries match the multipathed devices.

The following command displays all the device mapper devices and their major and minor numbers. The minor numbers determine the name of the dm device. For example, a minor number of 3 corresponds to the multipathed device `/dev/dm-3`.

```
# dmsetup ls
mpathd (253:4)
mpathp1 (253:12)
mpathfp1 (253:11)
mpathb (253:3)
mpathgp1 (253:14)
mpathhp1 (253:13)
mpatha (253:2)
mpathh (253:9)
mpathg (253:8)
VolGroup00-LogVol01 (253:1)
mpathf (253:7)
```

```
VolGroup00-LogVol100      (253:0)
mpathe (253:6)
mpathbp1 (253:10)
mpathd (253:5)
```

## 5.12. Troubleshooting with the multipathd Interactive Console

The **multipathd -k** command is an interactive interface to the **multipathd** daemon. Entering this command brings up an interactive multipath console. After executing this command, you can enter **help** to get a list of available commands, you can enter an interactive command, or you can enter **CTRL-D** to quit.

The **multipathd** interactive console can be used to troubleshoot problems you may be having with your system. For example, the following command sequence displays the multipath configuration, including the defaults, before exiting the console.

```
# multipathd -k
> > show config
> > CTRL-D
```

The following command sequence ensures that multipath has picked up any changes to the **multipath.conf**,

```
# multipathd -k
> > reconfigure
> > CTRL-D
```

Use the following command sequence to ensure that the path checker is working properly.

```
# multipathd -k
> > show paths
> > CTRL-D
```

## Appendix A. Revision History

<b>Revision 0.3-6</b>	<b>Mon Nov 9 2015</b>	<b>Steven Levine</b>
Preparing document for 7.2 GA publication		
<b>Revision 0.3-3</b>	<b>Wed Aug 19 2015</b>	<b>Steven Levine</b>
Preparing document for 7.2 Beta publication		
<b>Revision 0.2-7</b>	<b>Mon Feb 16 2015</b>	<b>Steven Levine</b>
Version for 7.1 GA release		
<b>Revision 0.2-6</b>	<b>Thu Dec 11 2014</b>	<b>Steven Levine</b>
Version for 7.1 Beta release		
<b>Revision 0.2-2</b>	<b>Tue Dec 2 2014</b>	<b>Steven Levine</b>
Resolves: #1162514, #1085979, #1129897 Small updates for technical clarification and correction.		
Resolves: #1117037 Documents -w and -W options of the multipath command.		
Resolves: #1117043 Documents new options for the features configuration parameter.		
Resolves: #1117499 Documents -force_sync configuration parameter.		
<b>Revision 0.2-1</b>	<b>Thu Oct 30 2014</b>	<b>Steven Levine</b>
Build to update procedure for moving root file system for 7.0		
<b>Revision 0.2-0</b>	<b>Wed Oct 22 2014</b>	<b>Steven Levine</b>
Updated procedure for moving root file system		
<b>Revision 0.1-22</b>	<b>Mon Jun 2 2014</b>	<b>Steven Levine</b>
Version for 7.0 GA release		
<b>Revision 0.1-20</b>	<b>Tue May 20 2014</b>	<b>Steven Levine</b>
Rebuild for style changes		
Resolves: #1069212, #1074858 Document new and updated parameters		
Resolves: #1040328 Small technical fixes		
<b>Revision 0.1-10</b>	<b>Wed Apr 9 2014</b>	<b>Steven Levine</b>
7.0 Beta update		
<b>Revision 0.1-3</b>	<b>Tue Nov 26 2013</b>	<b>Steven Levine</b>
Version for 7.0 pre-Beta		
<b>Revision 0.1-2</b>	<b>Tue Nov 12 2013</b>	<b>Steven Levine</b>

Removing RHEL 6 references.

**Revision 0.1-1**

**Wed Jan 16 2013**

**Steven Levine**

Branched from the Red Hat Enterprise Linux 6 version of the document

## Index

### Symbols

`/etc/multipath.conf` package, [Setting Up DM Multipath](#)

### A

#### active/active configuration

- definition, [Overview of DM Multipath](#)
- illustration, [Overview of DM Multipath](#)

#### active/passive configuration

- definition, [Overview of DM Multipath](#)
- illustration, [Overview of DM Multipath](#)

#### alias parameter, [Multipaths Device Configuration Attributes](#)

- configuration file, [Multipath Device Identifiers](#)

### B

#### blacklist

- configuration file, [Configuration File Blacklist](#)
- default devices, [Blacklisting By Device Name](#)
- device name, [Blacklisting By Device Name](#)
- device type, [Blacklisting By Device Type](#)
- WWID, [Blacklisting by WWID](#)

#### blacklist\_exceptions section

- multipath.conf file, [Blacklist Exceptions](#)

### C

#### checker\_timeout parameter, [Configuration File Defaults](#)

#### configuration file

- alias parameter, [Multipaths Device Configuration Attributes](#)
- blacklist, [Configuration File Blacklist](#)
- checker\_timeout parameter, [Configuration File Defaults](#)
- config\_dir parameter, [Multipaths Device Configuration Attributes](#)
- deferred\_remove parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)
- delay\_wait\_checks parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)
- delay\_watch\_checks parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)
- detect\_prio parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#)
- dev\_loss\_tmo parameter, [Configuration File Defaults](#), [Configuration File Devices](#)



- failback parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)
- fast\_io\_fail\_tmo parameter, [Configuration File Defaults](#), [Configuration File Devices](#)
- features parameter, [Configuration File Defaults](#), [Configuration File Devices](#)
- flush\_on\_last\_del parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)
- force\_sync parameter, [Configuration File Defaults](#)
- hardware\_handler parameter, [Configuration File Devices](#)
- hwtable\_string\_match parameter, [Configuration File Defaults](#)
- ignore\_new\_boot\_devs parameter, [Configuration File Defaults](#)
- log\_checker\_err parameter, [Configuration File Defaults](#)
- max\_fds parameter, [Configuration File Defaults](#)
- new\_bindings\_in\_boot parameter, [Configuration File Defaults](#)
- no\_path\_retry parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)
- overview, [Configuration File Overview](#)
- path\_checker parameter, [Configuration File Defaults](#), [Configuration File Devices](#)
- path\_grouping\_policy parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)
- path\_selector parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)
- polling-interval parameter, [Configuration File Defaults](#)
- prio parameter, [Configuration File Defaults](#), [Configuration File Devices](#)
- product parameter, [Configuration File Devices](#)
- product\_blacklist parameter, [Configuration File Devices](#)
- queue\_without\_daemon parameter, [Configuration File Defaults](#)
- retain\_attached\_hw\_handler parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#)
- retrigger\_delay parameter, [Configuration File Defaults](#)
- retrigger\_tries parameter, [Configuration File Defaults](#)
- revision parameter, [Configuration File Devices](#)
- rr\_min\_io parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#)
- rr\_weight parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)
- uid\_attribute parameter, [Configuration File Defaults](#)
- user\_friendly\_names parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)
- vendor parameter, [Configuration File Devices](#)
- verbosity parameter, [Configuration File Defaults](#)
- wwid parameter, [Multipaths Device Configuration Attributes](#)

## configuring

- DM Multipath, [Setting Up DM Multipath](#)

**config\_dir** parameter, [Multipaths Device Configuration Attributes](#)

## D

### defaults section

- multipath.conf file, [Configuration File Defaults](#)

**deferred\_remove** parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)

**delay\_wait\_checks** parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)

**delay\_watch\_checks** parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)

**detect\_prio** parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#)

**dev/mapper** directory, [Multipath Device Identifiers](#)

**device name**, [Multipath Device Identifiers](#)

**device-mapper-multipath** package, [Setting Up DM Multipath](#)

**devices**

- adding, [Configuring Storage Devices](#), [Configuration File Devices](#)

**devices section**

- multipath.conf file, [Configuration File Devices](#)

**dev\_loss\_tmo** parameter, [Configuration File Defaults](#), [Configuration File Devices](#)

**DM Multipath**

- and LVM, [Multipath Devices in Logical Volumes](#)
- components, [DM Multipath Components](#)
- configuration file, [The DM Multipath Configuration File](#)
- configuring, [Setting Up DM Multipath](#)
- definition, [Device Mapper Multipathing](#)
- device name, [Multipath Device Identifiers](#)
- devices, [Multipath Devices](#)
- failover, [Overview of DM Multipath](#)
- overview, [Overview of DM Multipath](#)
- redundancy, [Overview of DM Multipath](#)
- setup, [Setting Up DM Multipath](#)
- setup, overview, [DM Multipath Setup Overview](#)

**dm-n** devices, [Multipath Device Identifiers](#)

**dmsetup** command, determining device mapper entries, [Determining Device Mapper Entries with the dmsetup Command](#)

**dm\_multipath** kernel module, [DM Multipath Components](#)

## F

**failback** parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)

**failover**, [Overview of DM Multipath](#)

**fast\_io\_fail\_tmo** parameter, [Configuration File Defaults](#), [Configuration File Devices](#)

**features** parameter, [Configuration File Defaults](#), [Configuration File Devices](#)

**features**, new and changed, [New and Changed Features](#)

**flush\_on\_last\_del** parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)

**force\_sync** parameter, [Configuration File Defaults](#)

## H

**hardware\_handler** parameter, [Configuration File Devices](#)

**hwtable\_string\_match** parameter, [Configuration File Defaults](#)

## I

**ignore\_new\_boot\_devs** parameter, [Configuration File Defaults](#)

**initramfs**

- starting multipath, [Setting Up Multipathing in the initramfs File System](#)

## K

**kpartx** command, [DM Multipath Components](#)

## L

**local disks, ignoring**, [Ignoring Local Disks when Generating Multipath Devices](#)

**log\_checker\_err** parameter, [Configuration File Defaults](#)

**LVM physical volumes**

- multipath devices, [Multipath Devices in Logical Volumes](#)

**lvm.conf** file, [Multipath Devices in Logical Volumes](#)

## M

**max\_fds** parameter, [Configuration File Defaults](#)

**mpathconf** command, [DM Multipath Components](#)

**multipath** command, [DM Multipath Components](#)

- options, [Multipath Command Options](#)
- output, [Multipath Command Output](#)
- queries, [Multipath Queries with multipath Command](#)

**multipath daemon (multipathd)**, [The Multipath Daemon](#)

**multipath devices**, [Multipath Devices](#)

- logical volumes, [Multipath Devices in Logical Volumes](#)
- LVM physical volumes, [Multipath Devices in Logical Volumes](#)

**Multipath Helper**, [Automatic Configuration File Generation with Multipath Helper](#)

**multipath.conf** file, [Storage Array Support](#), [The DM Multipath Configuration File](#)

- blacklist\_exceptions section, [Blacklist Exceptions](#)
- defaults section, [Configuration File Defaults](#)
- devices section, [Configuration File Devices](#)
- multipaths section, [Multipaths Device Configuration Attributes](#)

**multipathd**

- command, [Troubleshooting with the multipathd Interactive Console](#)
- interactive console, [Troubleshooting with the multipathd Interactive Console](#)

**multipathd daemon**, [DM Multipath Components](#)

**multipathd start** command, [Setting Up DM Multipath](#)

**multipathed root file system**, [Moving root File Systems from a Single Path Device to a Multipath Device](#)

**multipathed swap file system**, [Moving swap File Systems from a Single Path Device to a Multipath Device](#)

**multipaths** section

- multipath.conf file, [Multipaths Device Configuration Attributes](#)

## N

**new\_bindings\_in\_boot** parameter, [Configuration File Defaults](#)

**no\_path\_retry** parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)

## O

### overview

- features, new and changed, [New and Changed Features](#)

## P

**path\_checker** parameter, [Configuration File Defaults](#), [Configuration File Devices](#)

**path\_grouping\_policy** parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)

**path\_selector** parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)

**polling\_interval** parameter, [Configuration File Defaults](#)

**prio** parameter, [Configuration File Defaults](#), [Configuration File Devices](#)

**product** parameter, [Configuration File Devices](#)

**product\_blacklist** parameter, [Configuration File Devices](#)

## Q

**queue\_without\_daemon** parameter, [Configuration File Defaults](#)

## R

resizing a multipath device, [Resizing an Online Multipath Device](#)

**retain\_attached\_hw\_handler** parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#)

**retrigger\_delay** parameter, [Configuration File Defaults](#)

**retrigger\_tries** parameter, [Configuration File Defaults](#)

**revision** parameter, [Configuration File Devices](#)

root file system, [Moving root File Systems from a Single Path Device to a Multipath Device](#)

**rr\_min\_io** parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#)

**rr\_weight** parameter, [Configuration File Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)

## S

### setup

- DM Multipath, [Setting Up DM Multipath](#)

storage array support, [Storage Array Support](#)

### storage arrays

- adding, [Configuring Storage Devices](#), [Configuration File Devices](#)

swap file system, [Moving swap File Systems from a Single Path Device to a Multipath Device](#)

## U

**uid\_attribute** parameter, [Configuration File Defaults](#)

**user\_friendly\_names** parameter, [Multipath Device Identifiers](#), [Configuration File](#)

[Defaults](#), [Multipaths Device Configuration Attributes](#), [Configuration File Devices](#)

## V

vendor parameter, [Configuration File Devices](#)

verbosity parameter, [Configuration File Defaults](#)

## W

World Wide Identifier (WWID), [Multipath Device Identifiers](#)

wwid parameter, [Multipaths Device Configuration Attributes](#)