



# **Red Hat Enterprise Linux 8.0 Beta**

## **Configuring and managing file systems**

Creating, modifying, and administering file systems in Red Hat Enterprise Linux 8



# Red Hat Enterprise Linux 8.0 Beta Configuring and managing file systems

---

Creating, modifying, and administering file systems in Red Hat Enterprise Linux 8

## Legal Notice

Copyright © 2019 Red Hat, Inc.

The text of and illustrations in this document are licensed by Red Hat under a Creative Commons Attribution–Share Alike 3.0 Unported license ("CC-BY-SA"). An explanation of CC-BY-SA is available at

<http://creativecommons.org/licenses/by-sa/3.0/>

. In accordance with CC-BY-SA, if you distribute this document or an adaptation of it, you must provide the URL for the original version.

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

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

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

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

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

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

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

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

All other trademarks are the property of their respective owners.

## Abstract

This documentation collection provides instructions on how to effectively manage file systems in Red Hat Enterprise Linux 8.

## Table of Contents

<b>THIS IS A BETA VERSION!</b> .....	<b>9</b>
<b>PROVIDING FEEDBACK ON RED HAT DOCUMENTATION</b> .....	<b>10</b>
<b>CHAPTER 1. GETTING STARTED WITH XFS</b> .....	<b>11</b>
1.1. CREATING AN XFS FILE SYSTEM	11
1.1.1. Creating an XFS file system with mkfs.xfs	11
Prerequisites	11
Procedure	11
Additional resources	12
1.1.2. Additional resources	12
1.2. BACKING UP AN XFS FILE SYSTEM	12
1.2.1. Features of XFS backup	12
Additional resources	12
1.2.2. Backing up an XFS file system with xfsdump	12
Prerequisites	13
Procedure	13
Additional resources	13
1.2.3. Additional resources	13
1.3. RESTORING AN XFS FILE SYSTEM FROM BACKUP	13
1.3.1. Features of restoring XFS from backup	14
Additional resources	14
1.3.2. Restoring an XFS file system from backup with xfsrestore	14
Prerequisites	14
Procedure	14
Additional resources	15
1.3.3. Informational messages when restoring an XFS backup from a tape	15
1.3.4. Additional resources	16
1.4. REPAIRING AN XFS FILE SYSTEM	16
1.4.1. Error-handling mechanisms in XFS	16
Unclean unmounts	16
Corruption	16
Additional resources	17
1.4.2. Repairing an XFS file system with xfs_repair	17
Procedure	17
Additional resources	18
1.5. INCREASING THE SIZE OF AN XFS FILE SYSTEM	18
1.5.1. Increasing the size of an XFS file system with xfs_growfs	18
Prerequisites	18
Procedure	18
Additional resources	19
<b>CHAPTER 2. GETTING STARTED WITH PARTITIONS</b> .....	<b>20</b>
2.1. VIEWING THE PARTITION TABLE	20
2.1.1. Viewing the partition table with parted	20
Procedure	20
Additional resources	20
2.1.2. Example output of parted print	20
2.2. CREATING A PARTITION TABLE ON A DISK	21
2.2.1. Considerations before modifying partitions on a disk	21
The maximum number of partitions	22
The maximum size of a partition	22

Size alignment	22
2.2.2. Comparison of partition table types	22
2.2.3. Creating a partition table on a disk with parted	23
Procedure	23
Additional resources	23
Next steps	24
2.3. CREATING A PARTITION	24
2.3.1. Considerations before modifying partitions on a disk	24
The maximum number of partitions	24
The maximum size of a partition	24
Size alignment	24
2.3.2. Partition types	25
Partition types or flags	25
Partition file system type	25
2.3.3. Creating a partition with parted	26
Prerequisites	26
Procedure	26
Additional resources	27
2.3.4. Setting a partition type with fdisk	27
Prerequisites	27
Procedure	27
2.4. REMOVING A PARTITION	28
2.4.1. Considerations before modifying partitions on a disk	28
The maximum number of partitions	28
The maximum size of a partition	29
Size alignment	29
2.4.2. Removing a partition with parted	29
Procedure	29
Additional resources	30
2.5. RESIZING A PARTITION	30
2.5.1. Considerations before modifying partitions on a disk	30
The maximum number of partitions	31
The maximum size of a partition	31
Size alignment	31
2.5.2. Resizing a partition with parted	31
Prerequisites	31
Procedure	32
Additional resources	33
<b>CHAPTER 3. OVERVIEW OF PERSISTENT NAMING ATTRIBUTES</b> .....	<b>34</b>
3.1. DISADVANTAGES OF NON-PERSISTENT NAMING ATTRIBUTES	34
3.2. FILE SYSTEM AND DEVICE IDENTIFIERS	34
File system identifiers	35
Device identifiers	35
Recommendations	35
3.3. DEVICE NAMES MANAGED BY THE UDEV MECHANISM IN /DEV/DISK/	35
3.3.1. File system identifiers	35
The UUID attribute in /dev/disk/by-uuid/	35
The Label attribute in /dev/disk/by-label/	36
3.3.2. Device identifiers	36
The WWID attribute in /dev/disk/by-id/	36
The Partition UUID attribute in /dev/disk/by-partuuid	37
The Path attribute in /dev/disk/by-path/	37

3.4. THE WORLD WIDE IDENTIFIER WITH DM MULTIPATH	37
3.5. LIMITATIONS OF THE UDEV DEVICE NAMING CONVENTION	38
3.6. LISTING PERSISTENT NAMING ATTRIBUTES	39
Procedure	39
3.7. MODIFYING PERSISTENT NAMING ATTRIBUTES	40
Prerequisites	40
Procedure	40
<b>CHAPTER 4. MOUNTING FILE SYSTEMS</b> .....	<b>41</b>
4.1. THE LINUX MOUNT MECHANISM	41
Additional resources	41
4.2. LISTING CURRENTLY MOUNTED FILE SYSTEMS	41
Procedure	41
Additional resources	42
4.3. MOUNTING A FILE SYSTEM WITH MOUNT	42
Prerequisites	42
Procedure	42
Additional resources	43
4.4. MOVING A MOUNT POINT	43
Procedure	43
Additional resources	43
4.5. UNMOUNTING A FILE SYSTEM WITH UMOUNT	43
Procedure	43
4.6. COMMON MOUNT OPTIONS	44
4.7. SHARING A MOUNT ON MULTIPLE MOUNT POINTS	45
4.7.1. Types of shared mounts	45
4.7.2. Creating a private mount point duplicate	45
Procedure	46
Additional resources	47
4.7.3. Creating a shared mount point duplicate	47
Procedure	47
Additional resources	48
4.7.4. Creating a slave mount point duplicate	48
Procedure	48
Additional resources	49
4.7.5. Preventing a mount point from being duplicated	49
Procedure	49
Additional resources	50
4.7.6. Related information	50
4.8. PERSISTENTLY MOUNTING FILE SYSTEMS	50
4.8.1. The /etc/fstab file	50
Additional resources	50
4.8.2. Adding a file system to /etc/fstab	51
Procedure	51
Additional resources	51
4.9. MOUNTING FILE SYSTEMS ON DEMAND	51
4.9.1. The autofs service	52
Additional resources	52
4.9.2. The autofs configuration files	52
The master map file	52
Map files	53
The amd map format	54
Additional resources	54

4.9.3. Configuring autofs mount points	54
Prerequisites	54
Procedure	54
4.9.4. Overriding or augmenting autofs site configuration files	54
Procedure	55
4.9.5. Using LDAP to store automounter maps	56
Prerequisites	56
Procedure	56
Additional resources	58
4.10. SETTING READ-ONLY PERMISSIONS FOR THE ROOT FILE SYSTEM	58
4.10.1. Files and directories that always retain write permissions	58
4.10.2. Configuring the root file system to mount with read-only permissions on boot	59
Procedure	59
Troubleshooting	60
<b>CHAPTER 5. MOUNTING AN SMB SHARE ON RED HAT ENTERPRISE LINUX</b>	<b>61</b>
5.1. PREREQUISITES	61
5.2. SUPPORTED SMB PROTOCOL VERSIONS	61
5.3. UNIX EXTENSIONS SUPPORT	61
5.4. MANUALLY MOUNTING AN SMB SHARE	62
Prerequisites	62
Procedure	62
5.5. MOUNTING AN SMB SHARE AUTOMATICALLY WHEN THE SYSTEM BOOTS	63
Prerequisites	63
Procedure	63
5.6. AUTHENTICATING TO AN SMB SHARE USING A CREDENTIALS FILE	63
Prerequisites	63
Procedure	63
5.7. PERFORMING A MULTI-USER SMB MOUNT	64
5.7.1. Prerequisites	64
5.7.2. Mounting a share with the multiuser option	64
Procedure	65
5.7.3. Verifying if an SMB share is mounted with the multiuser option	65
Procedure	65
5.7.4. Accessing a share as a user	65
5.8. FREQUENTLY USED MOUNT OPTIONS	65
<b>CHAPTER 6. MOUNTING NFS SHARES</b>	<b>67</b>
6.1. INTRODUCTION TO NFS	67
6.2. SUPPORTED NFS VERSIONS	67
Default NFS version	67
Features of minor NFS versions	67
The TCP and UDP protocols in NFSv3 and NFSv4	68
6.3. SERVICES REQUIRED BY NFS	68
The RPC services with NFSv4	69
Additional resources	69
6.4. NFS HOST NAME FORMATS	69
6.5. INSTALLING NFS	70
Procedure	70
6.6. DISCOVERING NFS EXPORTS	70
Procedure	70
Additional resources	71
6.7. MOUNTING AN NFS SHARE WITH MOUNT	71



---

Procedure	71
Additional resources	71
6.8. COMMON NFS MOUNT OPTIONS	71
Additional resources	73
6.9. RELATED INFORMATION	73
<b>CHAPTER 7. EXPORTING NFS SHARES</b> .....	<b>74</b>
7.1. INTRODUCTION TO NFS	74
7.2. SUPPORTED NFS VERSIONS	74
Default NFS version	74
Features of minor NFS versions	74
The TCP and UDP protocols in NFSv3 and NFSv4	75
7.3. SERVICES REQUIRED BY NFS	75
The RPC services with NFSv4	76
Additional resources	76
7.4. NFS HOST NAME FORMATS	76
7.5. NFS SERVER CONFIGURATION	77
7.5.1. The /etc/exports configuration file	77
Export entry	77
Default options	78
Default and overridden options	79
7.5.2. The exportfs utility	79
Common exportfs options	79
Additional resources	80
7.6. NFS AND RPCBIND	80
Additional resources	80
7.7. INSTALLING NFS	80
Procedure	81
7.8. STARTING THE NFS SERVER	81
Prerequisites	81
Procedure	81
Additional resources	81
7.9. TROUBLESHOOTING NFS AND RPCBIND	81
Procedure	81
Additional resources	82
7.10. CONFIGURING THE NFS SERVER TO RUN BEHIND A FIREWALL	82
Procedure	82
Additional resources	83
7.11. EXPORTING RPC QUOTA THROUGH A FIREWALL	83
Procedure	83
7.12. ENABLING NFS OVER RDMA (NFSORDMA)	84
Procedure	84
Additional resources	84
7.13. CONFIGURING AN NFSV4-ONLY SERVER	84
7.13.1. Benefits and drawbacks of an NFSv4-only server	84
7.13.2. NFS and rpcbind	85
Additional resources	85
7.13.3. Configuring the NFS server to support only NFSv4	85
Procedure	85
7.13.4. Verifying the NFSv4-only configuration	86
Procedure	86
7.14. RELATED INFORMATION	87

---

<b>CHAPTER 8. MANAGING LAYERED LOCAL STORAGE WITH STRATIS</b> .....	<b>88</b>
8.1. SETTING UP STRATIS FILE SYSTEMS	88
8.1.1. The purpose and features of Stratis	88
8.1.2. Components of a Stratis volume	88
8.1.3. Block devices usable with Stratis	89
Supported devices	89
Unsupported devices	90
Additional resources	90
8.1.4. Installing Stratis	90
Procedure	90
8.1.5. Creating a Stratis pool	90
Prerequisites	90
Procedure	90
Additional resources	91
Next steps	91
8.1.6. Creating a Stratis file system	91
Prerequisites	91
Procedure	91
Additional resources	91
Next steps	91
8.1.7. Mounting a Stratis file system	91
Prerequisites	91
Procedure	92
Additional resources	92
8.1.8. Persistently mounting a Stratis file system	92
Prerequisites	92
Procedure	92
Additional resources	93
8.1.9. Related information	93
8.2. EXTENDING A STRATIS VOLUME WITH ADDITIONAL BLOCK DEVICES	93
8.2.1. Components of a Stratis volume	93
8.2.2. Adding block devices to a Stratis pool	94
Prerequisites	94
Procedure	94
Additional resources	94
8.2.3. Related information	94
8.3. MONITORING STRATIS FILE SYSTEMS	94
8.3.1. Stratis sizes reported by different utilities	94
Additional resources	95
8.3.2. Displaying information about Stratis volumes	95
Prerequisites	95
Procedure	95
Additional resources	95
8.3.3. Related information	95
8.4. USING SNAPSHOTS ON STRATIS FILE SYSTEMS	96
8.4.1. Characteristics of Stratis snapshots	96
8.4.2. Creating a Stratis snapshot	96
Prerequisites	96
Procedure	96
Additional resources	96
8.4.3. Accessing the content of a Stratis snapshot	96
Prerequisites	97
Procedure	97

---

Additional resources	97
8.4.4. Reverting a Stratis file system to a previous snapshot	97
Prerequisites	97
Procedure	97
Additional resources	97
8.4.5. Removing a Stratis snapshot	98
Prerequisites	98
Procedure	98
Additional resources	98
8.4.6. Related information	98
8.5. REMOVING STRATIS FILE SYSTEMS	98
8.5.1. Components of a Stratis volume	98
8.5.2. Removing a Stratis file system	99
Prerequisites	99
Procedure	99
Additional resources	99
8.5.3. Removing a Stratis pool	100
Prerequisites	100
Procedure	100
Additional resources	100
8.5.4. Related information	100



## THIS IS A BETA VERSION!

Thank you for your interest in Red Hat Enterprise Linux 8.0 Beta. Be aware that:

- Beta code should not be used with production data or on production systems.
- Beta does not include a guarantee of support.
- Feedback and bug reports are welcome. Discussions with your account representative, partner contact, and Technical Account Manager (TAM) are also welcome.
- Upgrades to or from a Beta are not supported or recommended.

## PROVIDING FEEDBACK ON RED HAT DOCUMENTATION

We appreciate your input on our documentation. Please let us know how we could make it better. To do so:

- For simple comments on specific passages, make sure you are viewing the documentation in the Multi-page HTML format. Highlight the part of text that you want to comment on. Then, click the **Add Feedback** pop-up that appears below the highlighted text, and follow the displayed instructions.
- For submitting more complex feedback, create a Bugzilla ticket:
  1. Go to the [Bugzilla](#) website.
  2. As the Component, use **Documentation**.
  3. Fill in the **Description** field with your suggestion for improvement. Include a link to the relevant part(s) of documentation.
  4. Click **Submit Bug**.

# CHAPTER 1. GETTING STARTED WITH XFS

This is an overview of how to create and maintain XFS file systems.

## 1.1. CREATING AN XFS FILE SYSTEM

As a system administrator, you can create an XFS file system on a block device to enable it to store files and directories.

### 1.1.1. Creating an XFS file system with `mkfs.xfs`

This procedure describes how to create an XFS file system on a block device.

#### Prerequisites

- Create or reuse a partition on your disk.  
For information on creating MBR or GPT partitions, see [Section 2.2, “Creating a partition table on a disk”](#).

Alternatively, use an LVM or MD volume or a similar layer below XFS.

#### Procedure

1. To create the file system:

- If the device is a regular partition, an LVM volume, an MD volume, or a similar device, use the following command:

```
# mkfs.xfs block-device
```

- Replace *block-device* with the path to a partition or a logical volume. For example, `/dev/sdb1`, `/dev/disk/by-uuid/05e99ec8-def1-4a5e-8a9d-5945339ceb2a`, or `/dev/my-volgroup/my-lv`.
- In general, the default options are optimal for common use.
- When using `mkfs.xfs` on a block device containing an existing file system, add the `-f` option to overwrite that file system.
- To create the file system on a hardware RAID device, check if the system correctly detects the stripe geometry of the device:
  - If the stripe geometry information is correct, no additional options are needed. Create the file system:

```
# mkfs.xfs block-device
```

- If the information is incorrect, specify stripe geometry manually with the `su` and `sw` parameters of the `-d` option. The `su` parameter specifies the RAID chunk size, and the `sw` parameter specifies the number of data disks in the RAID device.

For example:

```
# mkfs.xfs -d su=64k,sw=4 /dev/sda3
```

2. Use the following command to wait for the system to register the new device node:

```
# udevadm settle
```

### Additional resources

- The `mkfs.xfs(8)` man page.

### 1.1.2. Additional resources

- The `mkfs.xfs(8)` man page.

## 1.2. BACKING UP AN XFS FILE SYSTEM

As a system administrator, you can use the `xfsdump` to back up an XFS file system into a file or on a tape. This provides a simple backup mechanism.

### 1.2.1. Features of XFS backup

This section describes key concepts and features of backing up an XFS file system with the `xfsdump` utility.

You can use the `xfsdump` utility to:

- Perform backups to regular file images.  
Only one backup can be written to a regular file.
- Perform backups to tape drives.  
The `xfsdump` utility also enables you to write multiple backups to the same tape. A backup can span multiple tapes.

To back up multiple file systems to a single tape device, simply write the backup to a tape that already contains an XFS backup. This appends the new backup to the previous one. By default, `xfsdump` never overwrites existing backups.

- Create incremental backups.  
The `xfsdump` utility uses dump levels to determine a base backup to which other backups are relative. Numbers from 0 to 9 refer to increasing dump levels. An incremental backup only backs up files that have changed since the last dump of a lower level:
  - To perform a full backup, perform a level 0 dump on the file system.
  - A level 1 dump is the first incremental backup after a full backup. The next incremental backup would be level 2, which only backs up files that have changed since the last level 1 dump; and so on, to a maximum of level 9.
- Exclude files from a backup using size, subtree, or inode flags to filter them.

### Additional resources

- The `xfsdump(8)` man page.

### 1.2.2. Backing up an XFS file system with `xfsdump`



This procedure describes how to back up the content of an XFS file system into a file or a tape.

### Prerequisites

- An XFS file system that you can back up.
- Another file system or a tape drive where you can store the backup.

### Procedure

- Use the following command to back up an XFS file system:

```
# xfsdump -l level [-L label] \  
-f backup-destination path-to-xfs-filesystem
```

- Replace *level* with the dump level of your backup. Use **0** to perform a full backup or **1** to **9** to perform consequent incremental backups.
- Replace *backup-destination* with the path where you want to store your backup. The destination can be a regular file, a tape drive, or a remote tape device. For example, **/backup-files/Data.xfsdump** for a file or **/dev/st0** for a tape drive.
- Replace *path-to-xfs-filesystem* with the mount point of the XFS file system you want to back up. For example, **/mnt/data/**. The file system must be mounted.
- When backing up multiple file systems and saving them on a single tape device, add a session label to each backup using the **-L *label*** option so that it is easier to identify them when restoring. Replace *label* with any name for your backup: for example, **backup\_data**.

#### Example 1.1. Backing up multiple XFS file systems

- To back up the content of XFS file systems mounted on the **/boot/** and **/data/** directories and save them as files in the **/backup-files/** directory:

```
# xfsdump -l 0 -f /backup-files/boot.xfsdump /boot  
# xfsdump -l 0 -f /backup-files/data.xfsdump /data
```

- To back up multiple file systems on a single tape device, add a session label to each backup using the **-L *label*** option:

```
# xfsdump -l 0 -L "backup_boot" -f /dev/st0 /boot  
# xfsdump -l 0 -L "backup_data" -f /dev/st0 /data
```

### Additional resources

- The **xfsdump(8)** man page.

### 1.2.3. Additional resources

- The **xfsdump(8)** man page.

## 1.3. RESTORING AN XFS FILE SYSTEM FROM BACKUP

As a system administrator, you can use the **xfsrestore** utility to restore XFS backup created with the **xfsdump** utility and stored in a file or on a tape.

### 1.3.1. Features of restoring XFS from backup

This section describes key concepts and features of restoring an XFS file system from backup with the **xfsrestore** utility.

The **xfsrestore** utility restores file systems from backups produced by **xfsdump**. The **xfsrestore** utility has two modes:

- The **simple** mode enables users to restore an entire file system from a level 0 dump. This is the default mode.
- The **cumulative** mode enables file system restoration from an incremental backup: that is, level 1 to level 9.

A unique *session ID* or *session label* identifies each backup. Restoring a backup from a tape containing multiple backups requires its corresponding session ID or label.

To extract, add, or delete specific files from a backup, enter the **xfsrestore** interactive mode. The interactive mode provides a set of commands to manipulate the backup files.

#### Additional resources

- The **xfsrestore(8)** man page.

### 1.3.2. Restoring an XFS file system from backup with xfsrestore

This procedure describes how to restore the content of an XFS file system from a file or tape backup.

#### Prerequisites

- A file or tape backup of XFS file systems, as described in [Section 1.2, “Backing up an XFS file system”](#).
- A storage device where you can restore the backup.

#### Procedure

- The command to restore the backup varies depending on whether you are restoring from a full backup or an incremental one, or are restoring multiple backups from a single tape device:

```
# xfsrestore [-r] [-S session-id] [-L session-label] [-i]
             -f backup-location restoration-path
```

- Replace *backup-location* with the location of the backup. This can be a regular file, a tape drive, or a remote tape device. For example, **/backup-files/Data.xfsdump** for a file or **/dev/st0** for a tape drive.
- Replace *restoration-path* with the path to the directory where you want to restore the file system. For example, **/mnt/data/**.
- To restore a file system from an incremental (level 1 to level 9) backup, add the **-r** option.

- o To restore a backup from a tape device that contains multiple backups, specify the backup using the **-S** or **-L** options.

The **-S** option lets you choose a backup by its session ID, while the **-L** option lets you choose by the session label. To obtain the session ID and session labels, use the **xfsrestore -I** command.

Replace *session-id* with the session ID of the backup. For example, **b74a3586-e52e-4a4a-8775-c3334fa8ea2c**. Replace *session-label* with the session label of the backup. For example, **my\_backup\_session\_label**.

- o To use **xfsrestore** interactively, use the **-i** option.  
The interactive dialog begins after **xfsrestore** finishes reading the specified device. Available commands in the interactive **xfsrestore** shell include **cd**, **ls**, **add**, **delete**, and **extract**; for a complete list of commands, use the **help** command.

### Example 1.2. Restoring Multiple XFS File Systems

- To restore the XFS backup files and save their content into directories under **/mnt/**:

```
# xfsrestore -f /backup-files/boot.xfsdump /mnt/boot/
# xfsrestore -f /backup-files/data.xfsdump /mnt/data/
```

- To restore from a tape device containing multiple backups, specify each backup by its session label or session ID:

```
# xfsrestore -L "backup_boot" -f /dev/st0 /mnt/boot/
# xfsrestore -S "45e9af35-efd2-4244-87bc-4762e476cbab" \
-f /dev/st0 /mnt/data/
```

### Additional resources

- The **xfsrestore(8)** man page.

### 1.3.3. Informational messages when restoring an XFS backup from a tape

When restoring a backup from a tape with backups from multiple file systems, the **xfsrestore** utility might issue messages. The messages inform you whether a match of the requested backup has been found when **xfsrestore** examines each backup on the tape in sequential order. For example:

```
xfsrestore: preparing drive
xfsrestore: examining media file 0
xfsrestore: inventory session uuid (8590224e-3c93-469c-a311-fc8f23029b2a)
does not match the media header's session uuid (7eda9f86-f1e9-4dfd-b1d4-
c50467912408)
xfsrestore: examining media file 1
xfsrestore: inventory session uuid (8590224e-3c93-469c-a311-fc8f23029b2a)
does not match the media header's session uuid (7eda9f86-f1e9-4dfd-b1d4-
c50467912408)
[...]
```

The informational messages keep appearing until the matching backup is found.

### 1.3.4. Additional resources

- The `xfsrestore(8)` man page.

## 1.4. REPAIRING AN XFS FILE SYSTEM

As a system administrator, you can repair a corrupted XFS file system.

### 1.4.1. Error-handling mechanisms in XFS

This section describes how XFS handles various kinds of errors in the file system.

#### Unclean unmounts

Journaling maintains a transactional record of metadata changes that happen on the file system.

In the event of a system crash, power failure, or other unclean unmount, XFS uses the journal (also called log) to recover the file system. The kernel performs journal recovery when mounting the XFS file system.

#### Corruption

In this context, *corruption* means errors on the file system caused by, for example:

- Hardware faults
- Bugs in storage firmware, device drivers, the software stack, or the file system itself
- Problems that cause parts of the file system to be overwritten by something outside of the file system

When XFS detects corruption in the file system or the file-system metadata, it shuts down the file system and reports the incident in the system log. Note that if the corruption occurred on the file system hosting the `/var` directory, these logs will not be available after a reboot.

#### Example 1.3. System log entry reporting an XFS corruption

```
# dmesg --notime | tail -15

XFS (loop0): Mounting V5 Filesystem
XFS (loop0): Metadata CRC error detected at
xfs_agi_read_verify+0xcb/0xf0 [xfs], xfs_agi block 0x2
XFS (loop0): Unmount and run xfs_repair
XFS (loop0): First 128 bytes of corrupted metadata buffer:
00000000027b3b56: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
.....
0000000005f9abc7a: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
.....
0000000005b0aef35: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
.....
000000000da9d2ded: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
.....
0000000001e265b07: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
.....
0000000006a40df69: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
.....
0000000000b272907: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
```

```

.....
00000000e484aac5: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
.....
XFS (loop0): metadata I/O error in "xfs_trans_read_buf_map" at daddr 0x2
len 1 error 74
XFS (loop0): xfs_imap_lookup: xfs_ialloc_read_agi() returned error -117,
agno 0
XFS (loop0): Failed to read root inode 0x80, error 11

```

User-space utilities usually report the *Input/output error* message when trying to access a corrupted XFS file system. Mounting an XFS file system with a corrupted log results in the following error message:

```

mount: /mount-point: mount(2) system call failed: Structure needs
cleaning.

```

You must manually use the **xfs\_repair** utility to repair the corruption. Unlike other file system repair utilities, **xfs\_repair** does not run at boot time, even when an XFS file system was not cleanly unmounted. In the event of an unclean unmount, XFS simply replays the log at mount time, ensuring a consistent file system; **xfs\_repair** cannot repair an XFS file system with a dirty log without remounting it first.

### Additional resources

- The **xfs\_repair(8)** man page provides a detailed list of XFS corruption checks.

## 1.4.2. Repairing an XFS file system with xfs\_repair

This procedure repairs a corrupted XFS file system using the **xfs\_repair** utility.

### Procedure

1. Clear the log by remounting the file system:

```

# mount file-system
# umount file-system

```

2. Use the **xfs\_repair** utility to repair the unmounted file system:

- If the mount succeeded, no additional options are required:

```

# xfs_repair block-device

```

- If the mount failed with the *Structure needs cleaning* error, the log is corrupted and cannot be replayed. Use the **-L** option (*force log zeroing*) to clear the log:

**WARNING**

This command causes all metadata updates in progress at the time of the crash to be lost, which might cause significant file system damage and data loss. This should be used only as a last resort.

```
# xfs_repair -L block-device
```

3. Mount the file system:


```
# mount file-system
```

**Additional resources**

- The `xfs_repair(8)` man page.

## 1.5. INCREASING THE SIZE OF AN XFS FILE SYSTEM

As a system administrator, you can increase the size of an XFS file system to utilize larger storage capacity.

**IMPORTANT**

It is not currently possible to decrease the size of XFS file systems.

### 1.5.1. Increasing the size of an XFS file system with `xfs_growfs`

This procedure describes how to grow an XFS file system using the `xfs_growfs` utility.

**Prerequisites**

- Ensure that the underlying block device is of an appropriate size to hold the resized file system later. Use the appropriate resizing methods for the affected block device.
- Mount the XFS file system.

**Procedure**

- While the XFS file system is mounted, use the `xfs_growfs` utility to increase its size:

```
# xfs_growfs file-system -D new-size
```

- Replace *file-system* with the mount point of the XFS file system.
- With the `-D` option, replace *new-size* with the desired new size of the file system specified in the number of file system blocks.  
To find out the block size in kB of a given XFS file system, use the `xfs_info` utility:

```
# xfs_info block-device
```

```
...  
data      =                bsize=4096  
...
```

- Without the **-D** option, **xfs\_growfs** grows the file system to the maximum size supported by the underlying device.

### Additional resources

- The **xfs\_growfs(8)** man page.

## CHAPTER 2. GETTING STARTED WITH PARTITIONS

As a system administrator, you can use the following procedures to create, delete, and modify various types of disk partitions.

For an overview of the advantages and disadvantages to using partitions on block devices, see the following KBase article: <https://access.redhat.com/solutions/163853>.

### 2.1. VIEWING THE PARTITION TABLE

As a system administrator, you can display the partition table of a block device to see the partition layout and details about individual partitions.

#### 2.1.1. Viewing the partition table with parted

This procedure describes how to view the partition table on a block device using the **parted** utility.

##### Procedure

1. Start the interactive **parted** shell:

```
#parted block-device
```

- Replace *block-device* with the path to the device you want to examine: for example, **/dev/sda**.

2. View the partition table:

```
(parted) print
```

3. Optionally, use the following command to switch to another device you want to examine next:

```
(parted) select block-device
```

##### Additional resources

- The **parted(8)** man page.

#### 2.1.2. Example output of parted print

This section provides an example output of the **print** command in the **parted** shell and describes fields in the output.

##### Example 2.1. Output of the print command

```
Model: ATA SAMSUNG MZNLN256 (scsi)
Disk /dev/sda: 256GB
Sector size (logical/physical): 512B/512B
Partition Table: msdos
Disk Flags:
```

Number	Start	End	Size	Type	File system	Flags
1	1049kB	269MB	268MB	primary	xfs	boot



2	269MB	34.6GB	34.4GB	primary
3	34.6GB	45.4GB	10.7GB	primary
4	45.4GB	256GB	211GB	extended
5	45.4GB	256GB	211GB	logical

Following is a description of the fields:

**Model: ATA SAMSUNG MZNLN256 (scsi)**

The disk type, manufacturer, model number, and interface.

**Disk /dev/sda: 256GB**

The disk label type.

**Number**

The partition number. For example, the partition with minor number 1 corresponds to `/dev/sda1`.

**Start and End**

The location on the device where the partition starts and ends.

**Type**

Valid types are metadata, free, primary, extended, or logical.

**File system**

The file system type. If the **File system** field of a device shows no value, this means that its file system type is unknown. The **parted** utility cannot recognize the file system on encrypted devices.

**Flags**

Lists the flags set for the partition. Available flags are **boot**, **root**, **swap**, **hidden**, **raid**, **lvm**, or **lba**.

## 2.2. CREATING A PARTITION TABLE ON A DISK

As a system administrator, you can format a block device with different types of partition tables to enable using partitions on the device.



**WARNING**

Formatting a block device with a partition table deletes all data stored on the device.

### 2.2.1. Considerations before modifying partitions on a disk

This section lists key points to consider before creating, removing, or resizing partitions.



**NOTE**

This section does not cover the DASD partition table, which is specific to the IBM System z architecture. For information on DASD, see the [What you should know about DASD](#) article at the IBM Knowledge Center.

### The maximum number of partitions

The number of partitions on a device is limited by the type of the partition table:

- On a device formatted with the **Master Boot Record (MBR)** partition table, you can have either:
  - Up to four primary partitions, or
  - Up to three primary partitions, one extended partition, and multiple logical partitions within the extended.
- On a device formatted with the **GUID Partition Table (GPT)**, the maximum number of partitions is 128. While the GPT specification allows for more partitions by growing the area reserved for the partition table, common practice used by the **parted** utility is to limit it to enough area for 128 partitions.

### The maximum size of a partition

The size of a partition on a device is limited by the type of the partition table:

- On a device formatted with the **Master Boot Record (MBR)** partition table, the maximum size is 2TiB.
- On a device formatted with the **GUID Partition Table (GPT)**, the maximum size is 8ZiB.

If you want to create a partition larger than 2TiB, the disk must be formatted with GPT.

### Size alignment

The **parted** utility enables you to specify partition size using multiple different suffixes:

#### MiB, GiB, or TiB

Size expressed in powers of 2.

- The starting point of the partition is aligned to the exact sector specified by size.
- The ending point is aligned to the specified size minus 1 sector.

#### MB, GB, or TB

Size expressed in powers of 10.

The starting and ending point is aligned within one half of the specified unit: for example,  $\pm 500\text{KB}$  when using the MB suffix.

## 2.2.2. Comparison of partition table types

This section compares the properties of different types of partition tables that you can create on a block device.

**Table 2.1. Partition table types**

Partition table	Maximum number of partitions	Maximum partition size
Master Boot Record (MBR)	4 primary, or 3 primary and 12 logical inside an extended partition	2TiB

Partition table	Maximum number of partitions	Maximum partition size
GUID Partition Table (GPT)	128	8ZiB

### 2.2.3. Creating a partition table on a disk with parted

This procedure describes how to format a block device with a partition table using the **parted** utility.

#### Procedure

1. Start the interactive **parted** shell:

```
# parted block-device
```

- Replace *block-device* with the path to the device where you want to create a partition table: for example, `/dev/sda`.

2. Determine if there already is a partition table on the device:

```
(parted) print
```

If the device already contains partitions, they will be deleted in the next steps.

3. Create the new partition table:

```
(parted) mklabel table-type
```

- Replace *table-type* with with the intended partition table type:
  - **msdos** for MBR
  - **gpt** for GPT

#### Example 2.2. Creating a GPT table

For example, to create a GPT table on the disk, use:

```
(parted) mklabel gpt
```

The changes start taking place as soon as you enter this command, so review it before executing it.

4. View the partition table to confirm that the partition table exists:

```
(parted) print
```

5. Exit the **parted** shell:

```
(parted) quit
```

#### Additional resources

- The **parted(8)** man page.

### Next steps

- Create partitions on the device. See [Section 2.3, “Creating a partition”](#) for details.

## 2.3. CREATING A PARTITION

As a system administrator, you can create new partitions on a disk.

### 2.3.1. Considerations before modifying partitions on a disk

This section lists key points to consider before creating, removing, or resizing partitions.



#### NOTE

This section does not cover the DASD partition table, which is specific to the IBM System z architecture. For information on DASD, see the [What you should know about DASD](#) article at the IBM Knowledge Center.

#### The maximum number of partitions

The number of partitions on a device is limited by the type of the partition table:

- On a device formatted with the **Master Boot Record (MBR)** partition table, you can have either:
  - Up to four primary partitions, or
  - Up to three primary partitions, one extended partition, and multiple logical partitions within the extended.
- On a device formatted with the **GUID Partition Table (GPT)**, the maximum number of partitions is 128. While the GPT specification allows for more partitions by growing the area reserved for the partition table, common practice used by the **parted** utility is to limit it to enough area for 128 partitions.

#### The maximum size of a partition

The size of a partition on a device is limited by the type of the partition table:

- On a device formatted with the **Master Boot Record (MBR)** partition table, the maximum size is 2TiB.
- On a device formatted with the **GUID Partition Table (GPT)**, the maximum size is 8ZiB.

If you want to create a partition larger than 2TiB, the disk must be formatted with GPT.

#### Size alignment

The **parted** utility enables you to specify partition size using multiple different suffixes:

#### MiB, GiB, or TiB

Size expressed in powers of 2.

- The starting point of the partition is aligned to the exact sector specified by size.
- The ending point is aligned to the specified size minus 1 sector.

## MB, GB, or TB

Size expressed in powers of 10.

The starting and ending point is aligned within one half of the specified unit: for example,  $\pm 500\text{KB}$  when using the MB suffix.

## 2.3.2. Partition types

This section describes different attributes that specify the type of a partition.

### Partition types or flags

The partition type, or flag, is used by a running system only rarely. However, the partition type matters to on-the-fly generators, such as **systemd-gpt-auto-generator**, which use the partition type to, for example, automatically identify and mount devices.

- The **parted** utility provides some control of partition types by mapping the partition type to *flags*. The parted utility can handle only certain partition types: for example LVM, swap, or RAID.
- The **fdisk** utility supports the full range of partition types by specifying hexadecimal codes.

### Partition file system type

The **parted** utility optionally accepts a file system type argument when creating a partition. The value is used to:

- Set the partition flags on MBR, or
- Set the partition UUID type on GPT. For example, the **swap**, **fat**, or **hfs** file system types set different GUIDs. The default value is the Linux Data GUID.

The argument does not modify the file system on the partition in any way. It only differentiates between the supported flags or GUIDs.

The following file system types are supported:

- **xf**s
- **ext2**
- **ext3**
- **ext4**
- **fat16**
- **fat32**
- **hfs**
- **hfs+**
- **linux-swap**
- **ntfs**
- **reiserfs**

### 2.3.3. Creating a partition with parted

This procedure describes how to create a new partition on a block device using the **parted** utility.

#### Prerequisites

- There is a partition table on the disk. For details on how to format the disk, see [Section 2.2, “Creating a partition table on a disk”](#).
- If the partition you want to create is larger than 2TiB, the disk must be formatted with the GUID Partition Table (GPT).

#### Procedure

1. Start the interactive **parted** shell:

```
# parted block-device
```

- Replace *block-device* with the path to the device where you want to create a partition: for example, **/dev/sda**.

2. View the current partition table to determine if there is enough free space:

```
(parted) print
```

- If there is not enough free space, you can resize an existing partition. For more information, see [Section 2.5, “Resizing a partition”](#).
- From the partition table, determine:
  - The start and end points of the new partition
  - On MBR, what partition type it should be.

3. Create the new partition:

```
(parted) mkpart part-type name fs-type start end
```

- Replace *part-type* with **primary**, **logical**, or **extended** based on what you decided from the partition table. This applies only to the MBR partition table.
- Replace *name* with an arbitrary partition name. This is required for GPT partition tables.
- Replace *fs-type* with any one of **xfs**, **ext2**, **ext3**, **ext4**, **fat16**, **fat32**, **hfs**, **hfs+**, **linux-swap**, **ntfs**, or **reiserfs**. The *fs-type* parameter is optional. Note that **parted** does not create the file system on the partition.
- Replace *start* and *end* with the sizes that determine the starting and ending points of the partition, counting from the beginning of the disk. You can use size suffixes, such as **512MiB**, **20GiB**, or **1.5TiB**. The default size megabytes.

#### Example 2.3. Creating a small primary partition

For example, to create a primary partition from 1024MiB until 2048MiB on an MBR table, use:

```
(parted) mkpart primary 1024MiB 2048MiB
```

```

The changes start taking place as soon as you enter this command, so review it before
executing it.

```

4. View the partition table to confirm that the created partition is in the partition table with the correct partition type, file system type, and size:

```

(parted) print

```

5. Exit the **parted** shell:

```

(parted) quit

```

6. Use the following command to wait for the system to register the new device node:

```

# udevadm settle

```

7. Verify that the kernel recognizes the new partition:

```

# cat /proc/partitions

```

### Additional resources

- The **parted(8)** man page.

### 2.3.4. Setting a partition type with fdisk

This procedure describes how to set a partition type, or flag, using the **fdisk** utility.

#### Prerequisites

- There is a partition on the disk.

#### Procedure

1. Start the interactive **fdisk** shell:

```

# fdisk block-device

```

- Replace *block-device* with the path to the device where you want to set a partition type: for example, **/dev/sda**.

2. View the current partition table to determine the minor partition number:

```

Command (m for help): print

```

You can see the current partition type in the **Type** column and its corresponding type ID in the **Id** column.

3. Enter the partition type command and select a partition using its minor number:

```

Command (m for help): type
Partition number (1,2,3 default 3): 2

```

- 
- Optionally, list the available hexadecimal codes:

```
Hex code (type L to list all codes): L
```

- Set the partition type:

```
Hex code (type L to list all codes): 8e
```

- Write your changes and exit the **fdisk** shell:

```
Command (m for help): write
The partition table has been altered.
Syncing disks.
```

- Verify your changes:

```
# fdisk --list block-device
```

## 2.4. REMOVING A PARTITION

As a system administrator, you can remove a disk partition that is no longer used to free up disk space.



### WARNING

Removing a partition deletes all data stored on the partition.

### 2.4.1. Considerations before modifying partitions on a disk

This section lists key points to consider before creating, removing, or resizing partitions.



### NOTE

This section does not cover the DASD partition table, which is specific to the IBM System z architecture. For information on DASD, see the [What you should know about DASD](#) article at the IBM Knowledge Center.

#### The maximum number of partitions

The number of partitions on a device is limited by the type of the partition table:

- On a device formatted with the **Master Boot Record (MBR)** partition table, you can have either:
  - Up to four primary partitions, or
  - Up to three primary partitions, one extended partition, and multiple logical partitions within the extended.
- On a device formatted with the **GUID Partition Table (GPT)**, the maximum number of partitions



is 128. While the GPT specification allows for more partitions by growing the area reserved for the partition table, common practice used by the **parted** utility is to limit it to enough area for 128 partitions.

### The maximum size of a partition

The size of a partition on a device is limited by the type of the partition table:

- On a device formatted with the **Master Boot Record (MBR)** partition table, the maximum size is 2TiB.
- On a device formatted with the **GUID Partition Table (GPT)**, the maximum size is 8ZiB.

If you want to create a partition larger than 2TiB, the disk must be formatted with GPT.

### Size alignment

The **parted** utility enables you to specify partition size using multiple different suffixes:

#### MiB, GiB, or TiB

Size expressed in powers of 2.

- The starting point of the partition is aligned to the exact sector specified by size.
- The ending point is aligned to the specified size minus 1 sector.

#### MB, GB, or TB

Size expressed in powers of 10.

The starting and ending point is aligned within one half of the specified unit: for example,  $\pm 500\text{KB}$  when using the MB suffix.

## 2.4.2. Removing a partition with parted

This procedure describes how to remove a disk partition using the **parted** utility.

### Procedure

1. Start the interactive **parted** shell:

```
# parted block-device
```

- Replace *block-device* with the path to the device where you want to remove a partition: for example, **/dev/sda**.

2. View the current partition table to determine the minor number of the partition to remove:

```
(parted) print
```

3. Remove the partition:

```
(parted) rm minor-number
```

- Replace *minor-number* with the minor number of the partition you want to remove: for example, **3**.

The changes start taking place as soon as you enter this command, so review it before executing it.

4. Confirm that the partition is removed from the partition table:

```
(parted) print
```

5. Exit the **parted** shell:

```
(parted) quit
```

6. Verify that the kernel knows the partition is removed:

```
# cat /proc/partitions
```

7. Remove the partition from the **/etc/fstab** file if it is present. Find the line that declares the removed partition, and remove it from the file.

8. Regenerate mount units so that your system registers the new **/etc/fstab** configuration:

```
# systemctl daemon-reload
```

9. If you have deleted a swap partition or removed pieces of LVM, regenerate GRUB configuration:

- On a BIOS-based system:

```
# grub2-mkconfig --output=/etc/grub2.cfg
```

- On a UEFI-based system:

```
# grub2-mkconfig --output=/etc/grub2-efi.cfg
```

10. To register the changes in the early boot system, rebuild the **initramfs** file system:

```
# dracut --force --verbose
```

### Additional resources

- The **parted(8)** man page

## 2.5. RESIZING A PARTITION

As a system administrator, you can extend a partition to utilize unused disk space, or shrink a partition to use its capacity for different purposes.

### 2.5.1. Considerations before modifying partitions on a disk

This section lists key points to consider before creating, removing, or resizing partitions.



## NOTE

This section does not cover the DASD partition table, which is specific to the IBM System z architecture. For information on DASD, see the [What you should know about DASD](#) article at the IBM Knowledge Center.

### The maximum number of partitions

The number of partitions on a device is limited by the type of the partition table:

- On a device formatted with the **Master Boot Record (MBR)** partition table, you can have either:
  - Up to four primary partitions, or
  - Up to three primary partitions, one extended partition, and multiple logical partitions within the extended.
- On a device formatted with the **GUID Partition Table (GPT)**, the maximum number of partitions is 128. While the GPT specification allows for more partitions by growing the area reserved for the partition table, common practice used by the **parted** utility is to limit it to enough area for 128 partitions.

### The maximum size of a partition

The size of a partition on a device is limited by the type of the partition table:

- On a device formatted with the **Master Boot Record (MBR)** partition table, the maximum size is 2TiB.
- On a device formatted with the **GUID Partition Table (GPT)**, the maximum size is 8ZiB.

If you want to create a partition larger than 2TiB, the disk must be formatted with GPT.

### Size alignment

The **parted** utility enables you to specify partition size using multiple different suffixes:

#### MiB, GiB, or TiB

Size expressed in powers of 2.

- The starting point of the partition is aligned to the exact sector specified by size.
- The ending point is aligned to the specified size minus 1 sector.

#### MB, GB, or TB

Size expressed in powers of 10.

The starting and ending point is aligned within one half of the specified unit: for example,  $\pm 500\text{KB}$  when using the MB suffix.

## 2.5.2. Resizing a partition with parted

This procedure resizes a disk partition using the **parted** utility.

### Prerequisites

- If you want to shrink a partition, back up the data that are stored on it.

**WARNING**

Shrinking a partition might result in data loss on the partition.

- If you want to resize a partition to be larger than 2TiB, the disk must be formatted with the GUID Partition Table (GPT). For details on how to format the disk, see [Section 2.2, “Creating a partition table on a disk”](#).

**Procedure**

1. If you want to shrink the partition, shrink the file system on it first so that it is not larger than the resized partition. Note that XFS does not support shrinking.
2. Start the interactive **parted** shell:

```
# parted block-device
```

- Replace *block-device* with the path to the device where you want to resize a partition: for example, **/dev/sda**.

3. View the current partition table:

```
(parted) print
```

From the partition table, determine:

- The minor number of the partition
- The location of the existing partition and its new ending point after resizing

4. Resize the partition:

```
(parted) resizepart minor-number new-end
```

- Replace *minor-number* with the minor number of the partition that you are resizing: for example, **3**.
- Replace *new-end* with the size that determines the new ending point of the resized partition, counting from the beginning of the disk. You can use size suffixes, such as **512MiB**, **20GiB**, or **1.5TiB**. The default size megabytes.

**Example 2.4. Extending a partition**

For example, to extend a partition located at the beginning of the disk to be 2GiB in size, use:

```
(parted) resizepart 1 2GiB
```

The changes start taking place as soon as you enter this command, so review it before executing it.

5. View the partition table to confirm that the resized partition is in the partition table with the correct size:

```
┆ (parted) print
```

6. Exit the **parted** shell:

```
┆ (parted) quit
```

7. Verify that the kernel recognizes the new partition:

```
┆ # cat /proc/partitions
```

8. If you extended the partition, extend the file system on it as well. See (reference) for details.

### Additional resources

- The **parted(8)** man page.

## CHAPTER 3. OVERVIEW OF PERSISTENT NAMING ATTRIBUTES

As a system administrator, you need to refer to storage volumes using persistent naming attributes to build storage setups that are reliable over multiple system boots.

### 3.1. DISADVANTAGES OF NON-PERSISTENT NAMING ATTRIBUTES

Red Hat Enterprise Linux provides a number of ways to identify storage devices. It is important to use the correct option to identify each device when used in order to avoid inadvertently accessing the wrong device, particularly when installing to or reformatting drives.

Traditionally, non-persistent names in the form of `/dev/sd(major number)(minor number)` are used on Linux to refer to storage devices. The major and minor number range and associated **sd** names are allocated for each device when it is detected. This means that the association between the major and minor number range and associated **sd** names can change if the order of device detection changes.

Such a change in the ordering might occur in the following situations:

- The parallelization of the system boot process detects storage devices in a different order with each system boot.
- A disk fails to power up or respond to the SCSI controller. This results in it not being detected by the normal device probe. The disk is not accessible to the system and subsequent devices will have their major and minor number range, including the associated **sd** names shifted down. For example, if a disk normally referred to as **sdb** is not detected, a disk that is normally referred to as **sdc** would instead appear as **sdb**.
- A SCSI controller (host bus adapter, or HBA) fails to initialize, causing all disks connected to that HBA to not be detected. Any disks connected to subsequently probed HBAs are assigned different major and minor number ranges, and different associated **sd** names.
- The order of driver initialization changes if different types of HBAs are present in the system. This causes the disks connected to those HBAs to be detected in a different order. This might also occur if HBAs are moved to different PCI slots on the system.
- Disks connected to the system with Fibre Channel, iSCSI, or FCoE adapters might be inaccessible at the time the storage devices are probed, due to a storage array or intervening switch being powered off, for example. This might occur when a system reboots after a power failure, if the storage array takes longer to come online than the system take to boot. Although some Fibre Channel drivers support a mechanism to specify a persistent SCSI target ID to WWPN mapping, this does not cause the major and minor number ranges, and the associated **sd** names to be reserved; it only provides consistent SCSI target ID numbers.

These reasons make it undesirable to use the major and minor number range or the associated **sd** names when referring to devices, such as in the `/etc/fstab` file. There is the possibility that the wrong device will be mounted and data corruption might result.

Occasionally, however, it is still necessary to refer to the **sd** names even when another mechanism is used, such as when errors are reported by a device. This is because the Linux kernel uses **sd** names (and also SCSI host/channel/target/LUN tuples) in kernel messages regarding the device.

### 3.2. FILE SYSTEM AND DEVICE IDENTIFIERS

This sections explains the difference between persistent attributes identifying file systems and block devices.

### File system identifiers

File system identifiers are tied to a particular file system created on a block device. The identifier is also stored as part of the file system. If you copy the file system to a different device, it still carries the same file system identifier. On the other hand, if you rewrite the device, such as by formatting it with the **mkfs** utility, the device loses the attribute.

File system identifiers include:

- Unique identifier (UUID)
- Label

### Device identifiers

Device identifiers are tied to a block device: for example, a disk or a partition. If you rewrite the device, such as by formatting it with the **mkfs** utility, the device keeps the attribute, because it is not stored in the file system.

Device identifiers include:

- World Wide Identifier (WWID)
- Partition UUID
- Serial number

### Recommendations

- Some file systems, such as logical volumes, span multiple devices. Red Hat recommends accessing these file systems using file system identifiers rather than device identifiers.

## 3.3. DEVICE NAMES MANAGED BY THE UDEV MECHANISM IN /DEV/DISK/

This section lists different kinds of persistent naming attributes that the **udev** service provides in the **/dev/disk/** directory.

The **udev** mechanism is used for all types of devices in Linux, not just for storage devices. In the case of storage devices, Red Hat Enterprise Linux contains **udev** rules that create symbolic links in the **/dev/disk/** directory. This enables you to refer to storage devices by:

- Their content
- A unique identifier
- Their serial number.

Although **udev** naming attributes are persistent, in that they do not change on their own across system reboots, some are also configurable.

### 3.3.1. File system identifiers

#### The UUID attribute in **/dev/disk/by-uuid/**

Entries in this directory provide a symbolic name that refers to the storage device by a **unique identifier** (UUID) in the content (that is, the data) stored on the device. For example:

```
/dev/disk/by-uuid/3e6be9de-8139-11d1-9106-a43f08d823a6
```

You can use the UUID to refer to the device in the `/etc/fstab` file using the following syntax:

```
UUID=3e6be9de-8139-11d1-9106-a43f08d823a6
```

You can configure the UUID attribute when creating a file system, and you can also change it later on.

### The Label attribute in `/dev/disk/by-label/`

Entries in this directory provide a symbolic name that refers to the storage device by a **label** in the content (that is, the data) stored on the device.

For example:

```
/dev/disk/by-label/Boot
```

You can use the label to refer to the device in the `/etc/fstab` file using the following syntax:

```
LABEL=Boot
```

You can configure the Label attribute when creating a file system, and you can also change it later on.

## 3.3.2. Device identifiers

### The WWID attribute in `/dev/disk/by-id/`

The World Wide Identifier (WWID) is a persistent, **system-independent identifier** that the SCSI Standard requires from all SCSI devices. The WWID identifier is guaranteed to be unique for every storage device, and independent of the path that is used to access the device. The identifier is a property of the device but is not stored in the content (that is, the data) on the devices.

This identifier can be obtained by issuing a SCSI Inquiry to retrieve the Device Identification Vital Product Data (page `0x83`) or Unit Serial Number (page `0x80`).

Red Hat Enterprise Linux automatically maintains the proper mapping from the WWID-based device name to a current `/dev/sd` name on that system. Applications can use the `/dev/disk/by-id/` name to reference the data on the disk, even if the path to the device changes, and even when accessing the device from different systems.

#### Example 3.1. WWID mappings

WWID symlink	Non-persistent device	Note
<code>/dev/disk/by-id/scsi-3600508b400105e210000900000490000</code>	<code>/dev/sda</code>	A device with a page <code>0x83</code> identifier
<code>/dev/disk/by-id/scsi-SSEAGATE_ST373453LW_3HW1RHM6</code>	<code>/dev/sdb</code>	A device with a page <code>0x80</code> identifier



WWID symlink	Non-persistent device	Note
<code>/dev/disk/by-id/ata-SAMSUNG_MZNLN256MHQ-000L7_S2WDX0J336519-part3</code>	<code>/dev/sdc3</code>	A disk partition

In addition to these persistent names provided by the system, you can also use **udev** rules to implement persistent names of your own, mapped to the WWID of the storage.

### The Partition UUID attribute in `/dev/disk/by-partuuid`

The Partition UUID (PARTUUID) attribute identifies partitions as defined by GPT partition table.

#### Example 3.2. Partition UUID mappings

PARTUUID symlink	Non-persistent device
<code>/dev/disk/by-partuuid/4cd1448a-01</code>	<code>/dev/sda1</code>
<code>/dev/disk/by-partuuid/4cd1448a-02</code>	<code>/dev/sda2</code>
<code>/dev/disk/by-partuuid/4cd1448a-03</code>	<code>/dev/sda3</code>

### The Path attribute in `/dev/disk/by-path/`

This attribute provides a symbolic name that refers to the storage device by the **hardware path** used to access the device.



#### WARNING

The Path attribute is unreliable, and Red Hat does not recommend using it.

## 3.4. THE WORLD WIDE IDENTIFIER WITH DM MULTIPATH

This section describes the mapping between the World Wide Identifier (WWID) and non-persistent device names in a Device Mapper Multipath configuration.

If there are multiple paths from a system to a device, DM Multipath uses the WWID to detect this. DM Multipath then presents a single "pseudo-device" in the `/dev/mapper/wwid` directory, such as `/dev/mapper/3600508b400105df70000e00000ac0000`.

The command `multipath -l` shows the mapping to the non-persistent identifiers:

- **Host:Channel:Target:LUN**
- **/dev/sd** name
- **major:minor** number


### Example 3.3. WWID mappings in a multipath configuration

An example output of the `multipath -l` command:

```
3600508b400105df70000e00000ac0000 dm-2 vendor,product
[size=20G][features=1 queue_if_no_path][hwandler=0][rw]
\_ round-robin 0 [prio=0][active]
  \_ 5:0:1:1 sdc 8:32 [active][undef]
  \_ 6:0:1:1 sdg 8:96 [active][undef]
\_ round-robin 0 [prio=0][enabled]
  \_ 5:0:0:1 sdb 8:16 [active][undef]
  \_ 6:0:0:1 sdf 8:80 [active][undef]
```

DM Multipath automatically maintains the proper mapping of each WWID-based device name to its corresponding `/dev/sd` name on the system. These names are persistent across path changes, and they are consistent when accessing the device from different systems.

When the `user_friendly_names` feature of DM Multipath is used, the WWID is mapped to a name of the form `/dev/mapper/mpathN`. By default, this mapping is maintained in the file `/etc/multipath/bindings`. These `mpathN` names are persistent as long as that file is maintained.



#### IMPORTANT

If you use `user_friendly_names`, then additional steps are required to obtain consistent names in a cluster.

## 3.5. LIMITATIONS OF THE UDEV DEVICE NAMING CONVENTION

The following are some limitations of the `udev` naming convention:

- It is possible that the device might not be accessible at the time the query is performed because the `udev` mechanism might rely on the ability to query the storage device when the `udev` rules are processed for a `udev` event. This is more likely to occur with Fibre Channel, iSCSI or FCoE storage devices when the device is not located in the server chassis.
- The kernel might send `udev` events at any time, causing the rules to be processed and possibly causing the `/dev/disk/by-*/` links to be removed if the device is not accessible.
- There might be a delay between when the `udev` event is generated and when it is processed, such as when a large number of devices are detected and the user-space `udev` service takes some amount of time to process the rules for each one. This might cause a delay between when the kernel detects the device and when the `/dev/disk/by-*/` names are available.
- External programs such as `blkid` invoked by the rules might open the device for a brief period of time, making the device inaccessible for other uses.

## 3.6. LISTING PERSISTENT NAMING ATTRIBUTES

This procedure describes how to find out the persistent naming attributes of non-persistent storage devices.

### Procedure

- To list the UUID and Label attributes, use the **lsblk** utility:

```
$ lsblk --fs storage-device
```

For example:

#### Example 3.4. Viewing the UUID and Label of a file system

```
$ lsblk --fs /dev/sda1

NAME FSTYPE LABEL UUID
MOUNTPOINT
sda1 xfs      Boot  afa5d5e3-9050-48c3-acc1-bb30095f3dc4 /boot
```

- To list the PARTUUID attribute, use the **lsblk** utility with the **--output +PARTUUID** option:

```
$ lsblk --output +PARTUUID
```

For example:

#### Example 3.5. Viewing the PARTUUID attribute of a partition

```
$ lsblk --output +PARTUUID /dev/sda1

NAME MAJ:MIN RM  SIZE RO  TYPE MOUNTPOINT PARTUUID
sda1   8:1    0  512M  0  part /boot      4cd1448a-01
```

- To list the WWID attribute, examine the targets of symbolic links in the **/dev/disk/by-id/** directory. For example:

#### Example 3.6. Viewing the WWID of all storage devices on the system

```
$ file /dev/disk/by-id/*

/dev/disk/by-id/ata-QEMU_HARDDISK_QM00001:
symbolic link to ../../sda
/dev/disk/by-id/ata-QEMU_HARDDISK_QM00001-part1:
symbolic link to ../../sda1
/dev/disk/by-id/ata-QEMU_HARDDISK_QM00001-part2:
symbolic link to ../../sda2
/dev/disk/by-id/dm-name-rhel_rhel8-root:
symbolic link to ../../dm-0
/dev/disk/by-id/dm-name-rhel_rhel8-swap:
symbolic link to ../../dm-1
/dev/disk/by-id/dm-uuid-LVM-
```


```

QIWtEHtXGobe5bewlIUdIVK0z5ofkgFhP0RMFsNyySVihqEl2cWwBR7MjXJo1D6g:
symbolic link to ../../dm-1
/dev/disk/by-id/dm-uuid-LVM-
QIWtEHtXGobe5bewlIUdIVK0z5ofkgFhXqH2M45hD2H9nAf2qfWSr1RLhzfMyOKd:
symbolic link to ../../dm-0
/dev/disk/by-id/lvm-pv-uuid-at1r2Y-vuMo-ueoH-CpMG-4JuH-AhEF-
wu4QQm:                               symbolic link to ../../sda2

```

### 3.7. MODIFYING PERSISTENT NAMING ATTRIBUTES

This procedure describes how to change the UUID or Label persistent naming attribute of a file system.



#### NOTE

Changing **udev** attributes happens in the background and might take a long time. The **udevadm settle** command waits until the change is fully registered, which ensures that your next command will be able to utilize the new attribute correctly.

In the following commands:

- Replace *new-uuid* with the UUID you want to set; for example, **1cdfbc07-1c90-4984-b5ec-f61943f5ea50**. You can generate a UUID using the **uuidgen** command.
- Replace *new-label* with a label; for example, **backup\_data**.

#### Prerequisites

- If you are modifying the attributes of an XFS file system, unmount it first.

#### Procedure

- To change the UUID or Label attributes of an **XFS** file system, use the **xfs\_admin** utility:

```

# xfs_admin -U new-uuid -L new-label storage-device
# udevadm settle

```

- To change the UUID or Label attributes of an **ext4**, **ext3**, or **ext2** file system, use the **tune2fs** utility:

```

# tune2fs -U new-uuid -L new-label storage-device
# udevadm settle

```

- To change the UUID or Label attributes of a swap volume, use the **swaponlabel** utility:

```

# swaponlabel --uuid new-uuid --label new-label swap-device
# udevadm settle

```

## CHAPTER 4. MOUNTING FILE SYSTEMS

As a system administrator, you can mount file systems on your system to access data on them.

### 4.1. THE LINUX MOUNT MECHANISM

This section explains basic concepts of mounting file systems on Linux.

On Linux, UNIX, and similar operating systems, file systems on different partitions and removable devices (CDs, DVDs, or USB flash drives for example) can be attached to a certain point (the mount point) in the directory tree, and then detached again. While a file system is mounted on a directory, the original content of the directory is not accessible.

Note that Linux does not prevent you from mounting a file system to a directory with a file system already attached to it.

When mounting, you can identify the device by:

- a universally unique identifier (UUID): for example, **UUID=34795a28-ca6d-4fd8-a347-73671d0c19cb**
- a volume label: for example, **LABEL=home**
- a full path to a non-persistent block device: for example, **/dev/sda3**

When you mount a file system using the **mount** command without all required information, that is without the device name, the target directory, or the file system type, the **mount** utility reads the content of the **/etc/fstab** file to check if the given file system is listed there. The **/etc/fstab** file contains a list of device names and the directories in which the selected file systems are set to be mounted as well as the file system type and mount options. Therefore, when mounting a file system that is specified in **/etc/fstab**, the following command syntax is sufficient:

- Mounting by the mount point:

```
# mount directory
```

- Mounting by the block device:

```
# mount device
```

#### Additional resources

- The **mount(8)** man page.
- For information on how to list persistent naming attributes such as the UUID, see [Section 3.6](#), “Listing persistent naming attributes”.

### 4.2. LISTING CURRENTLY MOUNTED FILE SYSTEMS

This procedure describes how to list all currently mounted file systems on the command line.

#### Procedure

- To list all mounted file systems, use the **findmnt** utility:

```
$ findmnt
```

- To limit the listed file systems only to a certain file system type, add the **--types** option:

```
$ findmnt --types fs-type
```

For example:

#### Example 4.1. Listing only XFS file systems

```
$ findmnt --types xfs

TARGET SOURCE
FSTYPE OPTIONS
/ /dev/mapper/luks-5564ed00-6aac-4406-bfb4-c59bf5de48b5 xfs
rw,relatime
└─/boot /dev/sda1
xfs rw,relatime
└─/home /dev/mapper/luks-9d185660-7537-414d-b727-d92ea036051e xfs
rw,relatime
```

#### Additional resources

- The **findmnt(8)** man page.

## 4.3. MOUNTING A FILE SYSTEM WITH MOUNT

This procedure describes how to mount a file system using the **mount** utility.

#### Prerequisites

- Make sure that no file system is already mounted on your chosen mount point:

```
$ findmnt mount-point
```

#### Procedure

1. To attach a certain file system, use the **mount** utility:

```
# mount device mount-point
```

#### Example 4.2. Mounting an XFS file system

For example, to mount a local XFS file system identified by UUID:

```
# mount UUID=ea74bbec-536d-490c-b8d9-5b40bbd7545b /mnt/data
```

2. If **mount** cannot recognize the file system type automatically, specify it using the **--types** option:

```
# mount --types type device mount-point
```

### Example 4.3. Mounting an NFS file system

For example, to mount a remote NFS file system:

```
# mount --types nfs4 host:/remote-export /mnt/nfs
```

### Additional resources

- The `mount(8)` man page.

## 4.4. MOVING A MOUNT POINT

This procedure describes how to change the mount point of a mounted file system to a different directory.

### Procedure

1. To change the directory in which a file system is mounted:

```
# mount --move old-directory new-directory
```

### Example 4.4. Moving a home file system

For example, to move the file system mounted in the `/mnt/userdirs/` directory to the `/home/` mount point:

```
# mount --move /mnt/userdirs /home
```

2. Verify that the file system has been moved as expected:

```
$ findmnt
$ ls old-directory
$ ls new-directory
```

### Additional resources

- The `mount(8)` man page.

## 4.5. UNMOUNTING A FILE SYSTEM WITH Umount

This procedure describes how to unmount a file system using the `umount` utility.

### Procedure

1. Try unmounting the file system using either of the following commands:
  - By mount point:

```
# umount mount-point
```

- By device:

```
# umount device
```

If the command fails with an error similar to the following, it means that the file system is in use because of a process is using resources on it:

```
umount: /run/media/user/FlashDrive: target is busy.
```

2. If the file system is in use, use the **fuser** utility to determine which processes are accessing it. For example:

```
$ fuser --mount /run/media/user/FlashDrive
/run/media/user/FlashDrive: 18351
```

Afterwards, terminate the processes using the file system and try unmounting it again.

## 4.6. COMMON MOUNT OPTIONS

This section lists some commonly used options of the **mount** utility.

You can use these options in the following syntax:

```
# mount --options option1,option2,option3 device mount-point
```

**Table 4.1. Common mount options**

Option	Description
<b>async</b>	Enables asynchronous input and output operations on the file system.
<b>auto</b>	Enables the file system to be mounted automatically using the <b>mount -a</b> command.
<b>defaults</b>	Provides an alias for the <b>async, auto, dev, exec, nouser, rw, suid</b> options.
<b>exec</b>	Allows the execution of binary files on the particular file system.
<b>loop</b>	Mounts an image as a loop device.
<b>noauto</b>	Default behavior disables the automatic mount of the file system using the <b>mount -a</b> command.
<b>noexec</b>	Disallows the execution of binary files on the particular file system.



Option	Description
<b>nouser</b>	Disallows an ordinary user (that is, other than root) to mount and unmount the file system.
<b>remount</b>	Remounts the file system in case it is already mounted.
<b>ro</b>	Mounts the file system for reading only.
<b>rw</b>	Mounts the file system for both reading and writing.
<b>user</b>	Allows an ordinary user (that is, other than root) to mount and unmount the file system.

## 4.7. SHARING A MOUNT ON MULTIPLE MOUNT POINTS

As a system administrator, you can duplicate mount points to make the file systems accessible from multiple directories.

### 4.7.1. Types of shared mounts

There are multiple types of shared mounts that you can use. The difference between them is what happens when you mount another file system under one of the shared mount points. The shared mounts are implemented using the *shared subtrees* functionality.

The types are:

#### Private mount

This type does not receive or forward any propagation events.

When you mount another file system under either the duplicate or the original mount point, it is not reflected in the other.

#### Shared mount

This type creates an exact replica of a given mount point.

When a mount point is marked as a shared mount, any mount within the original mount point is reflected in it, and vice versa.

This is the default mount type of the root file system.

#### Slave mount

This type creates a limited duplicate of a given mount point.

When a mount point is marked as a slave mount, any mount within the original mount point is reflected in it, but no mount within a slave mount is reflected in its original.

#### Unbindable mount

This type prevents the given mount point from being duplicated whatsoever.

### 4.7.2. Creating a private mount point duplicate

This procedure duplicates a mount point as a private mount. File systems that you later mount under the duplicate or the original mount point are not reflected in the other.

## Procedure

1. Create a virtual file system (VFS) node from the original mount point:

```
# mount --bind original-dir original-dir
```

2. Mark the original mount point as private:

```
# mount --make-private original-dir
```

Alternatively, to change the mount type for the selected mount point and all mount points under it, use the **--make-rprivate** option instead of **--make-private**.

3. Create the duplicate:

```
# mount --bind original-dir duplicate-dir
```

### Example 4.5. Duplicating /media into /mnt as a private mount point

1. Create a VFS node from the **/media** directory:

```
# mount --bind /media /media
```

2. Mark the **/media** directory as private:

```
# mount --make-private /media
```

3. Create its duplicate in **/mnt**:

```
# mount --bind /media /mnt
```

4. It is now possible to verify that **/media** and **/mnt** share content but none of the mounts within **/media** appear in **/mnt**. For example, if the CD-ROM drive contains non-empty media and the **/media/cdrom/** directory exists, use:

```
# mount /dev/cdrom /media/cdrom
# ls /media/cdrom
EFI  GPL  isolinux  LiveOS
# ls /mnt/cdrom
#
```

5. It is also possible to verify that file systems mounted in the **/mnt** directory are not reflected in **/media**. For instance, if a non-empty USB flash drive that uses the **/dev/sdc1** device is plugged in and the **/mnt/flashdisk/** directory is present, use:

```
# mount /dev/sdc1 /mnt/flashdisk
# ls /media/flashdisk
# ls /mnt/flashdisk
en-US  publican.cfg
```

■

**Additional resources**

- The **mount (8)** man page.

**4.7.3. Creating a shared mount point duplicate**

This procedure duplicates a mount point as a shared mount. File systems that you later mount under the original directory or the duplicate are always reflected in the other.

**Procedure**

1. Create a virtual file system (VFS) node from the original mount point:

```
# mount --bind original-dir original-dir
```

2. Mark the original mount point as shared:

```
# mount --make-shared original-dir
```

Alternatively, to change the mount type for the selected mount point and all mount points under it, use the **--make-rshared** option instead of **--make-shared**.

3. Create the duplicate:

```
# mount --bind original-dir duplicate-dir
```

**Example 4.6. Duplicating /media into /mnt as a shared mount point**

To make the **/media** and **/mnt** directories share the same content:

1. Create a VFS node from the **/media** directory:

```
# mount --bind /media /media
```

2. Mark the **/media** directory as shared:

```
# mount --make-shared /media
```

3. Create its duplicate in **/mnt**:

```
# mount --bind /media /mnt
```

4. It is now possible to verify that a mount within **/media** also appears in **/mnt**. For example, if the CD-ROM drive contains non-empty media and the **/media/cdrom/** directory exists, use:

```
# mount /dev/cdrom /media/cdrom
# ls /media/cdrom
EFI  GPL  isolinux  LiveOS
# ls /mnt/cdrom
EFI  GPL  isolinux  LiveOS
```

- Similarly, it is possible to verify that any file system mounted in the `/mnt` directory is reflected in `/media`. For instance, if a non-empty USB flash drive that uses the `/dev/sdc1` device is plugged in and the `/mnt/flashdisk/` directory is present, use:

```
# mount /dev/sdc1 /mnt/flashdisk
# ls /media/flashdisk
en-US  publican.cfg
# ls /mnt/flashdisk
en-US  publican.cfg
```

### Additional resources

- The `mount(8)` man page.

### 4.7.4. Creating a slave mount point duplicate

This procedure duplicates a mount point as a slave mount. File systems that you later mount under the original mount point are reflected in the duplicate but not the other way around.

#### Procedure

- Create a virtual file system (VFS) node from the original mount point:

```
# mount --bind original-dir original-dir
```

- Mark the original mount point as shared:

```
# mount --make-shared original-dir
```

Alternatively, to change the mount type for the selected mount point and all mount points under it, use the `--make-rshared` option instead of `--make-shared`.

- Create the duplicate and mark it as slave:

```
# mount --bind original-dir duplicate-dir
# mount --make-slave duplicate-dir
```

#### Example 4.7. Duplicating `/media` into `/mnt` as a slave mount point

This example shows how to get the content of the `/media` directory to appear in `/mnt` as well, but without any mounts in the `/mnt` directory to be reflected in `/media`.

- Create a VFS node from the `/media` directory:

```
# mount --bind /media /media
```

- Mark the `/media` directory as shared:

```
# mount --make-shared /media
```

- Create its duplicate in `/mnt` and mark it as slave:

```
# mount --bind /media /mnt
# mount --make-slave /mnt
```

- Verify that a mount within **/media** also appears in **/mnt**. For example, if the CD-ROM drive contains non-empty media and the **/media/cdrom/** directory exists, use:

```
# mount /dev/cdrom /media/cdrom
# ls /media/cdrom
EFI  GPL  isolinux  LiveOS
# ls /mnt/cdrom
EFI  GPL  isolinux  LiveOS
```

- Also verify that file systems mounted in the **/mnt** directory are not reflected in **/media**. For instance, if a non-empty USB flash drive that uses the **/dev/sdc1** device is plugged in and the **/mnt/flashdisk/** directory is present, use:

```
# mount /dev/sdc1 /mnt/flashdisk
# ls /media/flashdisk
# ls /mnt/flashdisk
en-US  publican.cfg
```

### Additional resources

- The **mount (8)** man page.

### 4.7.5. Preventing a mount point from being duplicated

This procedure marks a mount point as unbindable so that it is not possible to duplicate it in another mount point.

#### Procedure

- To change the type of a mount point to an unbindable mount, use:

```
# mount --bind mount-point mount-point
# mount --make-unbindable mount-point
```

Alternatively, to change the mount type for the selected mount point and all mount points under it, use the **--make-runbindable** option instead of **--make-unbindable**.

Any subsequent attempt to make a duplicate of this mount fails with the following error:

```
# mount --bind mount-point duplicate-dir

mount: wrong fs type, bad option, bad superblock on mount-point,
missing codepage or helper program, or other error
In some cases useful info is found in syslog - try
dmesg | tail  or so
```

#### Example 4.8. Preventing /media from being duplicated

- To prevent the **/media** directory from being shared, use:

```
# mount --bind /media /media
# mount --make-unbindable /media
```

### Additional resources

- The **mount (8)** man page.

### 4.7.6. Related information

- The *Shared subtrees* article on Linux Weekly News: <https://lwn.net/Articles/159077/>.

## 4.8. PERSISTENTLY MOUNTING FILE SYSTEMS

As a system administrator, you can persistently mount file systems to configure non-removable storage.

### 4.8.1. The `/etc/fstab` file

This section describes the `/etc/fstab` configuration file, which controls persistent mount points of file systems. Using `/etc/fstab` is the recommended way to persistently mount file systems.

Each line in the `/etc/fstab` file defines a mount point of a file system. It includes six fields separated by white space:

1. The block device identified by a persistent attribute or a path in the `/dev` directory.
2. The directory where the device will be mounted.
3. The file system on the device.
4. Mount options for the file system. The option **defaults** means that the partition is mounted at boot time with default options. This section also recognizes **systemd** mount unit options in the **x-systemd.option** format.
5. Backup option for the **dump** utility.
6. Check order for the **fsck** utility.

#### Example 4.9. The `/boot` file system in `/etc/fstab`

Block device	Mount point	File system	Options	Backup	Check
UUID=ea74bbec-536d-490c-b8d9-5b40bbd7545b	/boot	xfs	defaults	0	0

The **systemd** service automatically generates mount units from entries in `/etc/fstab`.

### Additional resources

- The **fstab(5)** man page.
- The *fstab* section of the **systemd.mount(5)** man page.

### 4.8.2. Adding a file system to */etc/fstab*

This procedure describes how to configure persistent mount point for a file system in the */etc/fstab* configuration file.

#### Procedure

1. Find out the UUID attribute of the file system:

```
$ lsblk --fs storage-device
```

For example:

#### Example 4.10. Viewing the UUID of a partition

```
$ lsblk --fs /dev/sda1

NAME FSTYPE LABEL UUID
MOUNTPOINT
sda1 xfs      Boot  ea74bbec-536d-490c-b8d9-5b40bbd7545b /boot
```

2. If the mount point directory does not exist, create it:

```
# mkdir --parents mount-point
```

3. As root, edit the */etc/fstab* file and add a line for the file system, identified by the UUID.

For example:

#### Example 4.11. The */boot* mount point in */etc/fstab*

```
UUID=ea74bbec-536d-490c-b8d9-5b40bbd7545b /boot xfs defaults 0 0
```

4. Regenerate mount units so that your system registers the new configuration:

```
# systemctl daemon-reload
```

5. Try mounting the file system to verify that the configuration works:

```
# mount mount-point
```

#### Additional resources

- Other persistent attributes that you can use to identify the file system: [Section 3.3, “Device names managed by the udev mechanism in \*/dev/disk/\*”](#)

## 4.9. MOUNTING FILE SYSTEMS ON DEMAND

As a system administrator, you can configure file systems, such as NFS, to mount automatically on demand.

### 4.9.1. The **autofs** service

This section explains the benefits and basic concepts of the **autofs** service, used to mount file systems on demand.

One drawback of permanent mounting using the `/etc/fstab` configuration is that, regardless of how infrequently a user accesses the mounted file system, the system must dedicate resources to keep the mounted file system in place. This might affect system performance when, for example, the system is maintaining NFS mounts to many systems at one time.

An alternative to `/etc/fstab` is to use the kernel-based **autofs** service. It consists of the following components:

- A kernel module that implements a file system, and
- A user-space service that performs all of the other functions.

The **autofs** service can mount and unmount file systems automatically (on-demand), therefore saving system resources. It can be used to mount file systems such as NFS, AFS, SMBFS, CIFS, and local file systems.

#### Additional resources

- The **autofs(8)** man page.

### 4.9.2. The **autofs** configuration files

This section describes the usage and syntax of configuration files used by the **autofs** service.

#### The master map file

The **autofs** service uses `/etc/auto.master` (master map) as its default primary configuration file. This can be changed to use another supported network source and name using the **autofs** configuration in the `/etc/autofs.conf` configuration file in conjunction with the Name Service Switch (NSS) mechanism.

All on-demand mount points must be configured in the master map. Mount point, host name, exported directory, and options can all be specified in a set of files (or other supported network sources) rather than configuring them manually for each host.

The master map file lists mount points controlled by **autofs**, and their corresponding configuration files or network sources known as automount maps. The format of the master map is as follows:

```
mount-point map-name options
```

The variables used in this format are:

#### **mount-point**

The **autofs** mount point; for example, `/mnt/data/`.

#### **map-file**

The map source file, which contains a list of mount points and the file system location from which those mount points should be mounted.



**options**

If supplied, these apply to all entries in the given map, if they do not themselves have options specified.

**Example 4.12. The /etc/auto.master file**

The following is a sample line from `/etc/auto.master` file:

```
/mnt/data /etc/auto.data
```

**Map files**

Map files configure the properties of individual on-demand mount points.

The automounter creates the directories if they do not exist. If the directories exist before the automounter was started, the automounter will not remove them when it exits. If a timeout is specified, the directory is automatically unmounted if the directory is not accessed for the timeout period.

The general format of maps is similar to the master map. However, the options field appears between the mount point and the location instead of at the end of the entry as in the master map:

```
mount-point options location
```

The variables used in this format are:

**mount-point**

This refers to the **autofs** mount point. This can be a single directory name for an indirect mount or the full path of the mount point for direct mounts. Each direct and indirect map entry key (*mount-point*) can be followed by a space separated list of offset directories (subdirectory names each beginning with `/`) making them what is known as a multi-mount entry.

**options**

When supplied, these are the mount options for the map entries that do not specify their own options. This field is optional.

**location**

This refers to the file system location such as a local file system path (preceded with the Sun map format escape character `:` for map names beginning with `/`), an NFS file system or other valid file system location.

**Example 4.13. A map file**

The following is a sample from a map file; for example, `/etc/auto.misc`:

```
payroll -fstype=nfs4 personnel:/dev/disk/by-uuid/52b94495-e106-4f29-
b868-fe6f6c2789b1
sales -fstype=xfss :/dev/disk/by-uuid/5564ed00-6aac-4406-bfb4-
c59bf5de48b5
```

The first column in the map file indicates the **autofs** mount point: **sales** and **payroll** from the server called **personnel**. The second column indicates the options for the **autofs** mount. The third column indicates the source of the mount.

Following the given configuration, the **autofs** mount points will be **/home/payroll** and **/home/sales**. The **-fstype=** option is often omitted and is generally not needed for correct operation.

Using the given configuration, if a process requires access to an **autofs** unmounted directory such as **/home/payroll/2006/July.sxc**, the **autofs** service automatically mounts the directory.

### The amd map format

The **autofs** service recognizes map configuration in the **amd** format as well. This is useful if you want to reuse existing automounter configuration written for the **am-utils** service, which has been removed from Red Hat Enterprise Linux.

However, Red Hat recommends using the simpler **autofs** format described in the previous sections.

### Additional resources

- The **autofs(5)**, **autofs.conf(5)**, and **auto.master(5)** man pages.
- For details on the **amd** map format, see the **/usr/share/doc/autofs/README.amd-maps** file, which is provided by the **autofs** package.

## 4.9.3. Configuring autofs mount points

This procedure describes how to configure on-demand mount points using the **autofs** service.

### Prerequisites

- Install the **autofs** package:

```
# yum install autofs
```

- Start and enable the **autofs** service:

```
# systemctl enable --now autofs
```

### Procedure

1. Create a map file for the on-demand mount point, located at **/etc/auto.*identifier***. Replace *identifier* with a name that identifies the mount point.
2. In the map file, fill in the mount point, options, and location fields as described in [Section 4.9.2, “The autofs configuration files”](#).
3. Register the map file in the master map file, as described in [Section 4.9.2, “The autofs configuration files”](#).
4. Try accessing content in the on-demand directory:

```
$ ls automounted-directory
```

## 4.9.4. Overriding or augmenting autofs site configuration files

It is sometimes useful to override site defaults for a specific mount point on a client system.

**Example 4.14. Initial conditions**

For example, consider the following conditions:

- Automounter maps are stored in NIS and the `/etc/nsswitch.conf` file has the following directive:

```
automount:    files nis
```

- The `auto.master` file contains:

```
+auto.master
```

- The NIS `auto.master` map file contains:

```
/home auto.home
```

- The NIS `auto.home` map contains:

```
beth    fileserver.example.com:/export/home/beth
joe     fileserver.example.com:/export/home/joe
*       fileserver.example.com:/export/home/&
```

- The file map `/etc/auto.home` does not exist.

**Procedure****Example 4.15. Mounting home directories from a different server**

Given the preceding conditions, let's assume that the client system needs to override the NIS map `auto.home` and mount home directories from a different server.

- In this case, the client needs to use the following `/etc/auto.master` map:

```
/home /etc/auto.home
+auto.master
```

- The `/etc/auto.home` map contains the entry:

```
*    labserver.example.com:/export/home/&
```

Because the automounter only processes the first occurrence of a mount point, the `/home` directory contains the content of `/etc/auto.home` instead of the NIS `auto.home` map.

**Example 4.16. Augmenting auto.home with only selected entries**

Alternatively, to augment the site-wide `auto.home` map with just a few entries:

1. Create an `/etc/auto.home` file map, and in it put the new entries. At the end, include the NIS `auto.home` map. Then the `/etc/auto.home` file map looks similar to:

```
mydir someserver:/export/mydir
+auto.home
```

2. With these NIS **auto.home** map conditions, listing the content of the **/home** directory outputs:

```
$ ls /home

beth joe mydir
```

This last example works as expected because **autofs** does not include the contents of a file map of the same name as the one it is reading. As such, **autofs** moves on to the next map source in the **nsswitch** configuration.

#### 4.9.5. Using LDAP to store automounter maps

This procedure configures **autofs** to store automounter maps in LDAP configuration rather than in **autofs** map files.

##### Prerequisites

- LDAP client libraries must be installed on all systems configured to retrieve automounter maps from LDAP. On Red Hat Enterprise Linux, the **openldap** package should be installed automatically as a dependency of the **autofs** package.

##### Procedure

1. To configure LDAP access, modify the **/etc/openldap/ldap.conf** file. Ensure that the **BASE**, **URI**, and **schema** options are set appropriately for your site.
2. The most recently established schema for storing automount maps in LDAP is described by the **rfc2307bis** draft. To use this schema, set it in the **/etc/autofs.conf** configuration file by removing the comment characters from the schema definition. For example:

##### Example 4.17. Setting autofs configuration

```
DEFAULT_MAP_OBJECT_CLASS="automountMap"
DEFAULT_ENTRY_OBJECT_CLASS="automount"
DEFAULT_MAP_ATTRIBUTE="automountMapName"
DEFAULT_ENTRY_ATTRIBUTE="automountKey"
DEFAULT_VALUE_ATTRIBUTE="automountInformation"
```

3. Ensure that all other schema entries are commented in the configuration. The **automountKey** attribute replaces the **cn** attribute in the **rfc2307bis** schema. Following is an example of an LDAP Data Interchange Format (LDIF) configuration:

##### Example 4.18. LDF Configuration

```
# extended LDIF
#
# LDAPv3
# base <> with scope subtree
# filter: (&(objectclass=automountMap))
```

```

(automountMapName=auto.master))
# requesting: ALL
#

# auto.master, example.com
dn: automountMapName=auto.master,dc=example,dc=com
objectClass: top
objectClass: automountMap
automountMapName: auto.master

# extended LDIF
#
# LDAPv3
# base <automountMapName=auto.master,dc=example,dc=com> with scope
subtree
# filter: (objectclass=automount)
# requesting: ALL
#

# /home, auto.master, example.com
dn: automountMapName=auto.master,dc=example,dc=com
objectClass: automount
cn: /home

automountKey: /home
automountInformation: auto.home

# extended LDIF
#
# LDAPv3
# base <> with scope subtree
# filter: (&(objectclass=automountMap)
(automountMapName=auto.home))
# requesting: ALL
#

# auto.home, example.com
dn: automountMapName=auto.home,dc=example,dc=com
objectClass: automountMap
automountMapName: auto.home

# extended LDIF
#
# LDAPv3
# base <automountMapName=auto.home,dc=example,dc=com> with scope
subtree
# filter: (objectclass=automount)
# requesting: ALL
#

# foo, auto.home, example.com
dn: automountKey=foo,automountMapName=auto.home,dc=example,dc=com
objectClass: automount
automountKey: foo
automountInformation: filer.example.com:/export/foo

```

```
# /, auto.home, example.com
dn: automountKey=/,automountMapName=auto.home,dc=example,dc=com
objectClass: automount
automountKey: /
automountInformation: filer.example.com:/export/&
```

### Additional resources

- The `rfc2307bis` draft: <https://tools.ietf.org/html/draft-howard-rfc2307bis>.

## 4.10. SETTING READ-ONLY PERMISSIONS FOR THE ROOT FILE SYSTEM

Sometimes, you need to mount the root file system (/) with read-only permissions. Example use cases include enhancing security or ensuring data integrity after an unexpected system power-off.

### 4.10.1. Files and directories that always retain write permissions

For the system to function properly, some files and directories need to retain write permissions. When the root file system is mounted in read-only mode, these files are mounted in RAM using the `tmpfs` temporary file system.

The default set of such files and directories is read from the `/etc/rwtab` file, which contains:

```
dirs /var/cache/man
dirs /var/gdm
<content truncated>

empty /tmp
empty /var/cache/foomatic
<content truncated>

files /etc/adjtime
files /etc/ntp.conf
<content truncated>
```

Entries in the `/etc/rwtab` file follow this format:

```
copy-method    path
```

In this syntax:

- Replace *copy-method* with one of the keywords specifying how the file or directory is copied to `tmpfs`.
- Replace *path* with the path to the file or directory.

The `/etc/rwtab` file recognizes the following ways in which a file or directory can be copied to `tmpfs`:

#### **empty**

An empty path is copied to `tmpfs`. For example:

```
empty /tmp
```

### dirs

A directory tree is copied to **tmpfs**, empty. For example:

```
dirs /var/run
```

### files

A file or a directory tree is copied to **tmpfs** intact. For example:

```
files /etc/resolv.conf
```

The same format applies when adding custom paths to **/etc/rwtab.d/**.

## 4.10.2. Configuring the root file system to mount with read-only permissions on boot

With this procedure, the root file system is mounted read-only on all following boots.

### Procedure

1. In the **/etc/sysconfig/readonly-root** file, set the **READONLY** option to **yes**:

```
# Set to 'yes' to mount the file systems as read-only.
READONLY=yes
```

2. Add the **ro** option in the root entry (**/**) in the **/etc/fstab** file:

```
/dev/mapper/luks-c376919e... / xfs x-systemd.device-timeout=0,ro
1 1
```

3. Add the **ro** option to the **GRUB\_CMDLINE\_LINUX** directive in the **/etc/default/grub** file and ensure that the directive does not contain **rw**:


```
GRUB_CMDLINE_LINUX="rhgb quiet... ro"
```

4. Recreate the GRUB2 configuration file:

```
# grub2-mkconfig -o /boot/grub2/grub.cfg
```

5. If you need to add files and directories to be mounted with write permissions in the **tmpfs** file system, create a text file in the **/etc/rwtab.d/** directory and put the configuration there. For example, to mount the **/etc/example/file** file with write permissions, add this line to the **/etc/rwtab.d/example** file:

```
files /etc/example/file
```



## IMPORTANT

Changes made to files and directories in **tmpfs** do not persist across boots.

6. Reboot the system to apply the changes.

## Troubleshooting

- If you mount the root file system with read-only permissions by mistake, you can remount it with read-and-write permissions again using the following command:

```
# mount -o remount,rw /
```



## CHAPTER 5. MOUNTING AN SMB SHARE ON RED HAT ENTERPRISE LINUX

The Server Message Block (SMB) protocol implements an application-layer network protocol used to access resources on a server, such as file shares and shared printers.



### NOTE

In the context of SMB, you can find mentions about the Common Internet File System (CIFS) protocol, which is a dialect of SMB. Both the SMB and CIFS protocol are supported, and the kernel module and utilities involved in mounting SMB and CIFS shares both use the name **cifs**.

This section describes how to mount shares from an SMB server. For details about setting up an SMB server on Red Hat Enterprise Linux using Samba, see the section about using Samba in the [Configuring and deploying different types of servers](#) guide.

### 5.1. PREREQUISITES

On Microsoft Windows, SMB is implemented by default. On Red Hat Enterprise Linux, the **cifs.ko** file system module of the kernel provides support for mounting SMB shares. Therefore install the **cifs-utils** package:

```
# yum install cifs-utils
```

The **cifs-utils** package provides utilities to:

- Mount SMB and CIFS shares
- Manage NT Lan Manager (NTLM) credentials in the kernel's keyring
- Set and display Access Control Lists (ACL) in a security descriptor on SMB and CIFS shares

### 5.2. SUPPORTED SMB PROTOCOL VERSIONS

The **cifs.ko** kernel module supports the following SMB protocol versions:

- SMB 1
- SMB 2.0
- SMB 2.1
- SMB 3.0



### NOTE

Depending on the protocol version, not all SMB features are implemented.

### 5.3. UNIX EXTENSIONS SUPPORT

Samba uses the **CAP\_UNIX** capability bit in the SMB protocol to provide the UNIX extensions feature. These extensions are also supported by the **cifs.ko** kernel module. However, both Samba and the kernel module support UNIX extensions only in the SMB 1 protocol.

To use UNIX extensions:

1. Set the **server min protocol** parameter in the **[global]** section in the **/etc/samba/smb.conf** file to **NT1**. This is the default on Samba servers.
2. Mount the share using the SMB 1 protocol by providing the **-o vers=1.0** option to the mount command. For example:

```
# mount -t cifs -o vers=1.0,username=user_name
//server_name/share_name /mnt/
```

By default, the kernel module uses SMB 2 or the highest later protocol version supported by the server. Passing the **-o vers=1.0** option to the **mount** command forces that the kernel module uses the SMB 1 protocol that is required for using UNIX extensions.

To verify if UNIX extensions are enabled, display the options of the mounted share:

```
# mount
...
//server/share on /mnt type cifs (... ,unix, ...)
```

If the **unix** entry is displayed in the list of mount options, UNIX extensions are enabled.

## 5.4. MANUALLY MOUNTING AN SMB SHARE

If you only require an SMB share to be temporary mounted, you can mount it manually using the **mount** utility.



### NOTE

Manually mounted shares are not mounted automatically again when you reboot the system. To configure that Red Hat Enterprise Linux automatically mounts the share when the system boots, see [Section 5.5, “Mounting an SMB share automatically when the system boots”](#).

### Prerequisites

- The **cifs-utils** package is installed.

### Procedure

To manually mount an SMB share, use the **mount** utility with the **-t cifs** parameter:

```
# mount -t cifs -o username=user_name //server_name/share_name /mnt/
Password for user_name@//server_name/share_name: password
```

In the **-o** parameter, you can specify options that are used to mount the share. For details, see [Section 5.8, “Frequently used mount options”](#) and the **OPTIONS** section in the **mount.cifs(8)** man page.

**Example 5.1. Mounting a share using an encrypted SMB 3.0 connection**

To mount the `\\server\example\` share as the `DOMAIN\Administrator` user over an encrypted SMB 3.0 connection into the `/mnt/` directory:

```
# mount -t cifs -o username=DOMAIN\Administrator,seal,vers=3.0
//server/example /mnt/
Password for DOMAIN\Administrator@//server_name/share_name: password
```

**5.5. MOUNTING AN SMB SHARE AUTOMATICALLY WHEN THE SYSTEM BOOTS**

If access to a mounted SMB share is permanently required on a server, mount the share automatically at boot time.

**Prerequisites**

- The `cifs-utils` package is installed.

**Procedure**

To mount an SMB share automatically when the system boots, add an entry for the share to the `/etc/fstab` file. For example:

```
//server_name/share_name /mnt cifs credentials=/root/smb.cred 0 0
```

**IMPORTANT**

To enable the system to mount a share automatically, you must store the user name, password, and domain name in a credentials file. For details, see [Section 5.6, “Authenticating to an SMB share using a credentials file”](#).

In the fourth field of the row in the `/etc/fstab`, specify mount options, such as the path to the credentials file. For details, see [Section 5.8, “Frequently used mount options”](#) and the **OPTIONS** section in the `mount.cifs(8)` man page.

To verify that the share mounts successfully, enter:

```
# mount /mnt/
```

**5.6. AUTHENTICATING TO AN SMB SHARE USING A CREDENTIALS FILE**

In certain situations, such as when mounting a share automatically at boot time, a share should be mounted without entering the user name and password. To implement this, create a credentials file.

**Prerequisites**

- The `cifs-utils` package is installed.

**Procedure**

1. Create a file, such as `/root/smb.cred`, and specify the user name, password, and domain name that file:

```
username=user_name  
password=password  
domain=domain_name
```

2. Set the permissions to only allow the owner to access the file:

```
# chown user_name /root/smb.cred  
# chmod 600 /root/smb.cred
```

You can now pass the `credentials=file_name` mount option to the `mount` utility or use it in the `/etc/fstab` file to mount the share without being prompted for the user name and password.

## 5.7. PERFORMING A MULTI-USER SMB MOUNT

The credentials you provide to mount a share determine the access permissions on the mount point by default. For example, if you use the `DOMAIN\example` user when you mount a share, all operations on the share will be executed as this user, regardless which local user performs the operation.

However, in certain situations, the administrator wants to mount a share automatically when the system boots, but users should perform actions on the share's content using their own credentials. The `multiuser` mount options lets you configure this scenario.



### IMPORTANT

To use the `multiuser` mount option, you must additionally set the `sec` mount option to a security type that supports providing credentials in a non-interactive way, such as `krb5` or the `ntlmssp` option with a credentials file. For details, see [Section 5.7.4, “Accessing a share as a user”](#).

The `root` user mounts the share using the `multiuser` option and an account that has minimal access to the contents of the share. Regular users can then provide their user name and password to the current session's kernel keyring using the `cifscreds` utility. If the user accesses the content of the mounted share, the kernel uses the credentials from the kernel keyring instead of the one initially used to mount the share.

Using this feature consists of the following steps:

- [Mount a share with the `multiuser` option.](#)
- [Optionally, verify if the share was successfully mounted with the `multiuser` option.](#)
- [Access the share as a user.](#)

### 5.7.1. Prerequisites

- The `cifs-utils` package is installed.

### 5.7.2. Mounting a share with the `multiuser` option

Before users can access the share with their own credentials, mount the share as the `root` user using

an account with limited permissions.

### Procedure

To mount a share automatically with the **multiuser** option when the system boots:

1. Create the entry for the share in the `/etc/fstab` file. For example:

```
//server_name/share_name /mnt cifs
multiuser,sec=ntlmssp,credentials=/root/smb.cred 0 0
```

2. Mount the share:

```
# mount /mnt/
```

If you do not want to mount the share automatically when the system boots, mount it manually by passing **-o multiuser,sec=security\_type** to the **mount** command. For details about mounting an SMB share manually, see [Section 5.4, “Manually mounting an SMB share”](#).

### 5.7.3. Verifying if an SMB share is mounted with the **multiuser** option

To verify if a share is mounted with the **multiuser** option, display the mount options.

#### Procedure

```
# mount
...
//server_name/share_name on /mnt type cifs (sec=ntlmssp,multiuser,...)
```

If the **multiuser** entry is displayed in the list of mount options, the feature is enabled.

### 5.7.4. Accessing a share as a user

If an SMB share is mounted with the **multiuser** option, users can provide their credentials for the server to the kernel's keyring:

```
# cifscreds add -u SMB_user_name server_name
Password: password
```

When the user performs operations in the directory that contains the mounted SMB share, the server applies the file system permissions for this user, instead of the one initially used when the share was mounted.



#### NOTE

Multiple users can perform operations using their own credentials on the mounted share at the same time.

## 5.8. FREQUENTLY USED MOUNT OPTIONS

When you mount an SMB share, the mount options determine:

- How the connection will be established with the server. For example, which SMB protocol version is used when connecting to the server.

- How the share will be mounted into the local file system. For example, if the system overrides the remote file and directory permissions to enable multiple local users to access the content on the server.

To set multiple options in the fourth field of the `/etc/fstab` file or in the `-o` parameter of a mount command, separate them with commas. For example, see [Section 5.7.2, “Mounting a share with the `multiuser` option”](#).

The following list gives frequently used mount options:

Option	Description
<code>credentials=file_name</code>	Sets the path to the credentials file. See <a href="#">Section 5.6, “Authenticating to an SMB share using a credentials file”</a>
<code>dir_mode=mode</code>	Sets the directory mode if the server does not support CIFS UNIX extensions.
<code>file_mode=mode</code>	Sets the file mode if the server does not support CIFS UNIX extensions.
<code>password=password</code>	Sets the password used to authenticate to the SMB server. Alternatively, specify a credentials file using the <b>credentials</b> option.
<code>seal</code>	Enables encryption support for connections using SMB 3.0 or a later protocol version. Therefore, use <b>seal</b> together with the <b>vers</b> mount option set to <b>3.0</b> or later. See <a href="#">Example 5.1, “Mounting a share using an encrypted SMB 3.0 connection”</a> .
<code>sec=security_mode</code>	<p>Sets the security mode, such as <b>ntlmsspi</b>, to enable NTLMv2 password hashing and enabled packet signing. For a list of supported values, see the option’s description in the <b>mount.cifs(8)</b> man page.</p> <p>If the server does not support the <b>ntlmv2</b> security mode, use <b>sec=ntlmssp</b>, which is the default.</p> <p>For security reasons, do not use the insecure <b>ntlm</b> security mode.</p>
<code>username=user_name</code>	Sets the user name used to authenticate to the SMB server. Alternatively, specify a credentials file using the <b>credentials</b> option.
<code>vers=SMB_protocol_version</code>	Sets the SMB protocol version used for the communication with the server.

For a complete list, see the **OPTIONS** section in the **mount.cifs(8)** man page.

## CHAPTER 6. MOUNTING NFS SHARES

As a system administrator, you can mount remote NFS shares on your system to access shared data.

### 6.1. INTRODUCTION TO NFS

This section explains the basic concepts of the NFS service.

A Network File System (NFS) allows remote hosts to mount file systems over a network and interact with those file systems as though they are mounted locally. This enables you to consolidate resources onto centralized servers on the network.

The NFS server refers to the `/etc/exports` configuration file to determine whether the client is allowed to access any exported file systems. Once verified, all file and directory operations are available to the user.

### 6.2. SUPPORTED NFS VERSIONS

This section lists versions of NFS supported in Red Hat Enterprise Linux and their features.

Currently, Red Hat Enterprise Linux 8 supports the following major versions of NFS:

- NFS version 3 (NFSv3) supports safe asynchronous writes and is more robust at error handling than the previous NFSv2; it also supports 64-bit file sizes and offsets, allowing clients to access more than 2 GB of file data.
- NFS version 4 (NFSv4) works through firewalls and on the Internet, no longer requires an `rpcbind` service, supports Access Control Lists (ACLs), and utilizes stateful operations.

NFS version 2 (NFSv2) is no longer supported by Red Hat.

#### Default NFS version

The default NFS version in Red Hat Enterprise Linux 8 is 4.2. NFS clients attempt to mount using NFSv4.2 by default, and fall back to NFSv4.1 when the server does not support NFSv4.2. The mount later fall back to NFSv4.0 and then to NFSv3.

#### Features of minor NFS versions

Following are the features of NFSv4.2 in Red Hat Enterprise Linux 8:

##### Server-side copy

Enables the NFS client to efficiently copy data without wasting network resources using the `copy_file_range()` system call.

##### Sparse files

Enables files to have one or more *holes*, which are unallocated or uninitialized data blocks consisting only of zeroes. The `lseek()` operation in NFSv4.2 supports `seek_hole()` and `seek_data()`, which enables applications to map out the location of holes in the sparse file.

##### Space reservation

Permits storage servers to reserve free space, which prohibits servers to run out of space. NFSv4.2 supports the `allocate()` operation to reserve space, the `deallocate()` operation to unreserve space, and the `fallocate()` operation to preallocate or deallocate space in a file.

##### Labeled NFS

Enforces data access rights and enables SELinux labels between a client and a server for individual files on an NFS file system.

### Layout enhancements

Provides the **layoutstats()** operation, which enables some Parallel NFS (pNFS) servers to collect better performance statistics.

Following are the features of NFSv4.1:

- Enhances performance and security of network, and also includes client-side support for pNFS.
- No longer requires a separate TCP connection for callbacks, which allows an NFS server to grant delegations even when it cannot contact the client: for example, when NAT or a firewall interferes.
- Provides exactly once semantics (except for reboot operations), preventing a previous issue whereby certain operations sometimes returned an inaccurate result if a reply was lost and the operation was sent twice.

### The TCP and UDP protocols in NFSv3 and NFSv4

NFSv4 requires the Transmission Control Protocol (TCP) running over an IP network. NFSv3 can also use the User Datagram Protocol (UDP) running over an IP network to provide a stateless network connection between the client and server.

When using NFSv3 with UDP, the stateless UDP connection (under normal conditions) has less protocol overhead than TCP. This can translate into better performance on very clean, non-congested networks. However, because UDP is stateless, if the server goes down unexpectedly, UDP clients continue to saturate the network with requests for the server. In addition, when a frame is lost with UDP, the entire RPC request must be retransmitted; with TCP, only the lost frame needs to be resent. For these reasons, TCP is the preferred protocol when connecting to an NFS server.

## 6.3. SERVICES REQUIRED BY NFS

This section lists system services that are required for running an NFS server or mounting NFS shares. Red Hat Enterprise Linux starts these services automatically.

Red Hat Enterprise Linux uses a combination of kernel-level support and service processes to provide NFS file sharing. All NFS versions rely on Remote Procedure Calls (RPC) between clients and servers. To share or mount NFS file systems, the following services work together depending on which version of NFS is implemented:

### **nfsd**

The NFS server to service requests for shared NFS file systems.

### **rpcbind**

Accepts port reservations from local RPC services. These ports are then made available (or advertised) so the corresponding remote RPC services can access them. The **rpcbind** service responds to requests for RPC services and sets up connections to the requested RPC service. This is not used with NFSv4.

### **rpc.mountd**

This process is used by an NFS server to process **MOUNT** requests from NFSv3 clients. It checks that the requested NFS share is currently exported by the NFS server, and that the client is allowed to access it. If the mount request is allowed, the **nfs-mountd** service replies with a Success status and provides the File-Handle for this NFS share back to the NFS client.

### **rpc.nfsd**



This process enables explicit NFS versions and protocols the server advertises to be defined. It works with the Linux kernel to meet the dynamic demands of NFS clients, such as providing server threads each time an NFS client connects. This process corresponds to the **nfs-server** service.

### lockd

This is a kernel thread that runs on both clients and servers. It implements the Network Lock Manager (NLM) protocol, which enables NFSv3 clients to lock files on the server. It is started automatically whenever the NFS server is run and whenever an NFS file system is mounted.

### rpc.statd

This process implements the Network Status Monitor (NSM) RPC protocol, which notifies NFS clients when an NFS server is restarted without being gracefully brought down. The **rpc-statd** service is started automatically by the **nfs-server** service, and does not require user configuration. This is not used with NFSv4.

### rpc.rquotad

This process provides user quota information for remote users. The **rpc-rquotad** service is started automatically by the **nfs-server** service and does not require user configuration.

### rpc.idmapd

This process provides NFSv4 client and server upcalls, which map between on-the-wire NFSv4 names (strings in the form of **user@domain**) and local UIDs and GIDs. For **idmapd** to function with NFSv4, the **/etc/idmapd.conf** file must be configured. At a minimum, the **Domain** parameter should be specified, which defines the NFSv4 mapping domain. If the NFSv4 mapping domain is the same as the DNS domain name, this parameter can be skipped. The client and server must agree on the NFSv4 mapping domain for ID mapping to function properly.

Only the NFSv4 server uses **rpc.idmapd**, which is started by the **nfs-idmapd** service. The NFSv4 client uses the keyring-based **nfsidmap** utility, which is called by the kernel on-demand to perform ID mapping. If there is a problem with **nfsidmap**, the client falls back to using **rpc.idmapd**.

## The RPC services with NFSv4

The mounting and locking protocols have been incorporated into the NFSv4 protocol. The server also listens on the well-known TCP port 2049. As such, NFSv4 does not need to interact with **rpcbind**, **lockd**, and **rpc-statd** services. The **nfs-mountd** service is still required on the NFS server to set up the exports, but is not involved in any over-the-wire operations.

## Additional resources

- To configure an NFSv4-only server, which does not require **rpcbind**, see [Section 7.13](#), “Configuring an NFSv4-only server”.

## 6.4. NFS HOST NAME FORMATS

This section describes different formats that you can use to specify a host when mounting or exporting an NFS share.

You can specify the host in the following formats:

### Single machine

Either of the following:

- A fully-qualified domain name (that can be resolved by the server)
- Host name (that can be resolved by the server)

- An IP address.

### Series of machines specified with wildcards

You can use the `*` or `?` characters to specify a string match.

Wildcards are not to be used with IP addresses; however, they might accidentally work if reverse DNS lookups fail. When specifying wildcards in fully qualified domain names, dots (`.`) are not included in the wildcard. For example, `*.example.com` includes `one.example.com` but does not include `one.two.example.com`.

### IP networks

Either of the following formats is valid:

- `a.b.c.d/z`, where `a.b.c.d` is the network and `z` is the number of bits in the netmask; for example `192.168.0.0/24`.
- `a.b.c.d/netmask`, where `a.b.c.d` is the network and `netmask` is the netmask; for example, `192.168.100.8/255.255.255.0`.

### Netgroups

The `@group-name` format, where `group-name` is the NIS netgroup name.

## 6.5. INSTALLING NFS

This procedure installs all packages necessary to mount or export NFS shares.

### Procedure

- Install the `nfs-utils` package:

```
# yum install nfs-utils
```

## 6.6. DISCOVERING NFS EXPORTS

This procedure discovers which file systems a given NFSv3 or NFSv4 server exports.

### Procedure

- With any server that supports NFSv3, use the `showmount` utility:

```
$ showmount --exports my-server  
  
Export list for my-server  
/exports/foo  
/exports/bar
```

- With any server that supports NFSv4, mount the root directory and look around:

```
# mount my-server:/ /mnt/  
# ls /mnt/  
  
exports
```

```
# ls /mnt/exports/
foo
bar
```

On servers that support both NFSv4 and NFSv3, both methods work and give the same results.

### Additional resources

- The **showmount (8)** man page.

## 6.7. MOUNTING AN NFS SHARE WITH MOUNT

This procedure mounts an NFS share exported from a server using the **mount** utility.

### Procedure

- To mount an NFS share, use the following command:

```
# mount -t nfs -o options host:/remote/export /local/directory
```

This command uses the following variables:

#### **options**

A comma-delimited list of mount options.

#### **host**

The host name, IP address, or fully qualified domain name of the server exporting the file system you wish to mount.

#### **/remote/export**

The file system or directory being exported from the server, that is, the directory you wish to mount.

#### **/local/directory**

The client location where */remote/export* is mounted.

### Additional resources

- [Section 6.8, “Common NFS mount options”](#)
- [Section 6.4, “NFS host name formats”](#)
- [Section 4.3, “Mounting a file system with mount”](#)
- The **mount (8)** man page

## 6.8. COMMON NFS MOUNT OPTIONS

This section lists options commonly used when mounting NFS shares. These options can be used with manual mount commands, */etc/fstab* settings, and **autofs**.

### Common NFS mount options

**lookupcache=mode**

Specifies how the kernel should manage its cache of directory entries for a given mount point. Valid arguments for *mode* are **all**, **none**, or **positive**.

**nfsvers=version**

Specifies which version of the NFS protocol to use, where *version* is **3**, **4**, **4.0**, **4.1**, or **4.2**. This is useful for hosts that run multiple NFS servers, or to disable retrying a mount with lower versions. If no version is specified, NFS uses the highest version supported by the kernel and the **mount** utility. The option **vers** is identical to **nfsvers**, and is included in this release for compatibility reasons.

**noacl**

Turns off all ACL processing. This may be needed when interfacing with older versions of Red Hat Enterprise Linux, Red Hat Linux, or Solaris, because the most recent ACL technology is not compatible with older systems.

**noLOCK**

Disables file locking. This setting is sometimes required when connecting to very old NFS servers.

**noexec**

Prevents execution of binaries on mounted file systems. This is useful if the system is mounting a non-Linux file system containing incompatible binaries.

**nosuid**

Disables the **set-user-identifier** and **set-group-identifier** bits. This prevents remote users from gaining higher privileges by running a **setuid** program.

**port=num**

Specifies the numeric value of the NFS server port. If *num* is **0** (the default value), then **mount** queries the **rpcbind** service on the remote host for the port number to use. If the NFS service on the remote host is not registered with its **rpcbind** service, the standard NFS port number of TCP 2049 is used instead.

**rsize=num and wsize=num**

These options set the maximum number of bytes to be transferred in a single NFS read or write operation.

There is no fixed default value for **rsize** and **wsize**. By default, NFS uses the largest possible value that both the server and the client support. In Red Hat Enterprise Linux 8, the client and server maximum is 1,048,576 bytes. For more details, see the [What are the default and maximum values for rsize and wsize with NFS mounts?](#) KBase article.

**sec=mode**

Security flavors to use for accessing files on the mounted export.

The default setting is **sec=sys**, which uses local UNIX UIDs and GIDs. These use **AUTH\_SYS** to authenticate NFS operations.

Other options include:

- **sec=krb5** uses Kerberos V5 instead of local UNIX UIDs and GIDs to authenticate users.
- **sec=krb5i** uses Kerberos V5 for user authentication and performs integrity checking of NFS operations using secure checksums to prevent data tampering.
- **sec=krb5p** uses Kerberos V5 for user authentication, integrity checking, and encrypts NFS traffic to prevent traffic sniffing. This is the most secure setting, but it also involves the most performance overhead.

**tcp**

Instructs the NFS mount to use the TCP protocol.

**udp**

Instructs the NFS mount to use the UDP protocol.

**Additional resources**

- The **mount(8)** man page
- The **nfs(5)** man page

**6.9. RELATED INFORMATION**

- The Linux NFS wiki: <http://linux-nfs.org>
- To mount NFS shares persistently, see [Section 4.8, “Persistently mounting file systems”](#).
- To mount NFS shares on demand, see [Section 4.9, “Mounting file systems on demand”](#).

## CHAPTER 7. EXPORTING NFS SHARES

As a system administrator, you can use the NFS server to share a directory on your system over network.

### 7.1. INTRODUCTION TO NFS

This section explains the basic concepts of the NFS service.

A Network File System (NFS) allows remote hosts to mount file systems over a network and interact with those file systems as though they are mounted locally. This enables you to consolidate resources onto centralized servers on the network.

The NFS server refers to the `/etc/exports` configuration file to determine whether the client is allowed to access any exported file systems. Once verified, all file and directory operations are available to the user.

### 7.2. SUPPORTED NFS VERSIONS

This section lists versions of NFS supported in Red Hat Enterprise Linux and their features.

Currently, Red Hat Enterprise Linux 8 supports the following major versions of NFS:

- NFS version 3 (NFSv3) supports safe asynchronous writes and is more robust at error handling than the previous NFSv2; it also supports 64-bit file sizes and offsets, allowing clients to access more than 2 GB of file data.
- NFS version 4 (NFSv4) works through firewalls and on the Internet, no longer requires an `rpcbind` service, supports Access Control Lists (ACLs), and utilizes stateful operations.

NFS version 2 (NFSv2) is no longer supported by Red Hat.

#### Default NFS version

The default NFS version in Red Hat Enterprise Linux 8 is 4.2. NFS clients attempt to mount using NFSv4.2 by default, and fall back to NFSv4.1 when the server does not support NFSv4.2. The mount later fall back to NFSv4.0 and then to NFSv3.

#### Features of minor NFS versions

Following are the features of NFSv4.2 in Red Hat Enterprise Linux 8:

##### Server-side copy

Enables the NFS client to efficiently copy data without wasting network resources using the `copy_file_range()` system call.

##### Sparse files

Enables files to have one or more *holes*, which are unallocated or uninitialized data blocks consisting only of zeroes. The `lseek()` operation in NFSv4.2 supports `seek_hole()` and `seek_data()`, which enables applications to map out the location of holes in the sparse file.

##### Space reservation

Permits storage servers to reserve free space, which prohibits servers to run out of space. NFSv4.2 supports the `allocate()` operation to reserve space, the `deallocate()` operation to unreserve space, and the `fallocate()` operation to preallocate or deallocate space in a file.

##### Labeled NFS

Enforces data access rights and enables SELinux labels between a client and a server for individual files on an NFS file system.

### Layout enhancements

Provides the **layoutstats()** operation, which enables some Parallel NFS (pNFS) servers to collect better performance statistics.

Following are the features of NFSv4.1:

- Enhances performance and security of network, and also includes client-side support for pNFS.
- No longer requires a separate TCP connection for callbacks, which allows an NFS server to grant delegations even when it cannot contact the client: for example, when NAT or a firewall interferes.
- Provides exactly once semantics (except for reboot operations), preventing a previous issue whereby certain operations sometimes returned an inaccurate result if a reply was lost and the operation was sent twice.

### The TCP and UDP protocols in NFSv3 and NFSv4

NFSv4 requires the Transmission Control Protocol (TCP) running over an IP network. NFSv3 can also use the User Datagram Protocol (UDP) running over an IP network to provide a stateless network connection between the client and server.

When using NFSv3 with UDP, the stateless UDP connection (under normal conditions) has less protocol overhead than TCP. This can translate into better performance on very clean, non-congested networks. However, because UDP is stateless, if the server goes down unexpectedly, UDP clients continue to saturate the network with requests for the server. In addition, when a frame is lost with UDP, the entire RPC request must be retransmitted; with TCP, only the lost frame needs to be resent. For these reasons, TCP is the preferred protocol when connecting to an NFS server.

## 7.3. SERVICES REQUIRED BY NFS

This section lists system services that are required for running an NFS server or mounting NFS shares. Red Hat Enterprise Linux starts these services automatically.

Red Hat Enterprise Linux uses a combination of kernel-level support and service processes to provide NFS file sharing. All NFS versions rely on Remote Procedure Calls (RPC) between clients and servers. To share or mount NFS file systems, the following services work together depending on which version of NFS is implemented:

### **nfsd**

The NFS server to service requests for shared NFS file systems.

### **rpcbind**

Accepts port reservations from local RPC services. These ports are then made available (or advertised) so the corresponding remote RPC services can access them. The **rpcbind** service responds to requests for RPC services and sets up connections to the requested RPC service. This is not used with NFSv4.

### **rpc.mountd**

This process is used by an NFS server to process **MOUNT** requests from NFSv3 clients. It checks that the requested NFS share is currently exported by the NFS server, and that the client is allowed to access it. If the mount request is allowed, the **nfs-mountd** service replies with a Success status and provides the File-Handle for this NFS share back to the NFS client.

### **rpc.nfsd**

This process enables explicit NFS versions and protocols the server advertises to be defined. It works with the Linux kernel to meet the dynamic demands of NFS clients, such as providing server threads each time an NFS client connects. This process corresponds to the **nfs-server** service.

### **lockd**

This is a kernel thread that runs on both clients and servers. It implements the Network Lock Manager (NLM) protocol, which enables NFSv3 clients to lock files on the server. It is started automatically whenever the NFS server is run and whenever an NFS file system is mounted.

### **rpc.statd**

This process implements the Network Status Monitor (NSM) RPC protocol, which notifies NFS clients when an NFS server is restarted without being gracefully brought down. The **rpc-statd** service is started automatically by the **nfs-server** service, and does not require user configuration. This is not used with NFSv4.

### **rpc.rquotad**

This process provides user quota information for remote users. The **rpc-rquotad** service is started automatically by the **nfs-server** service and does not require user configuration.

### **rpc.idmapd**

This process provides NFSv4 client and server upcalls, which map between on-the-wire NFSv4 names (strings in the form of **user@domain**) and local UIDs and GIDs. For **idmapd** to function with NFSv4, the **/etc/idmapd.conf** file must be configured. At a minimum, the **Domain** parameter should be specified, which defines the NFSv4 mapping domain. If the NFSv4 mapping domain is the same as the DNS domain name, this parameter can be skipped. The client and server must agree on the NFSv4 mapping domain for ID mapping to function properly.

Only the NFSv4 server uses **rpc.idmapd**, which is started by the **nfs-idmapd** service. The NFSv4 client uses the keyring-based **nfsidmap** utility, which is called by the kernel on-demand to perform ID mapping. If there is a problem with **nfsidmap**, the client falls back to using **rpc.idmapd**.

## **The RPC services with NFSv4**

The mounting and locking protocols have been incorporated into the NFSv4 protocol. The server also listens on the well-known TCP port 2049. As such, NFSv4 does not need to interact with **rpcbind**, **lockd**, and **rpc-statd** services. The **nfs-mountd** service is still required on the NFS server to set up the exports, but is not involved in any over-the-wire operations.

## **Additional resources**

- To configure an NFSv4-only server, which does not require **rpcbind**, see [Section 7.13, “Configuring an NFSv4-only server”](#).

## **7.4. NFS HOST NAME FORMATS**

This section describes different formats that you can use to specify a host when mounting or exporting an NFS share.

You can specify the host in the following formats:

### **Single machine**

Either of the following:

- A fully-qualified domain name (that can be resolved by the server)
- Host name (that can be resolved by the server)



- An IP address.

### Series of machines specified with wildcards

You can use the `*` or `?` characters to specify a string match.

Wildcards are not to be used with IP addresses; however, they might accidentally work if reverse DNS lookups fail. When specifying wildcards in fully qualified domain names, dots (`.`) are not included in the wildcard. For example, `*.example.com` includes `one.example.com` but does not include `one.two.example.com`.

### IP networks

Either of the following formats is valid:

- `a.b.c.d/z`, where `a.b.c.d` is the network and `z` is the number of bits in the netmask; for example `192.168.0.0/24`.
- `a.b.c.d/netmask`, where `a.b.c.d` is the network and `netmask` is the netmask; for example, `192.168.100.8/255.255.255.0`.

### Netgroups

The `@group-name` format, where `group-name` is the NIS netgroup name.

## 7.5. NFS SERVER CONFIGURATION

This section describes the syntax and options of two ways to configure exports on an NFS server:

- Manually editing the `/etc/exports` configuration file
- Using the `exportfs` utility on the command line

### 7.5.1. The `/etc/exports` configuration file

The `/etc/exports` file controls which file systems are exported to remote hosts and specifies options. It follows the following syntax rules:

- Blank lines are ignored.
- To add a comment, start a line with the hash mark (`#`).
- You can wrap long lines with a backslash (`\`).
- Each exported file system should be on its own individual line.
- Any lists of authorized hosts placed after an exported file system must be separated by space characters.
- Options for each of the hosts must be placed in parentheses directly after the host identifier, without any spaces separating the host and the first parenthesis.

### Export entry

Each entry for an exported file system has the following structure:

```
export host(options)
```

It is also possible to specify multiple hosts, along with specific options for each host. To do so, list them on the same line as a space-delimited list, with each host name followed by its respective options (in parentheses), as in:

```
export host1(options1) host2(options2) host3(options3)
```

In this structure:

**export**

The directory being exported

**host**

The host or network to which the export is being shared

**options**

The options to be used for host

**Example 7.1. A simple /etc/exports file**

In its simplest form, the **/etc/exports** file only specifies the exported directory and the hosts permitted to access it:

```
/exported/directory bob.example.com
```

Here, **bob.example.com** can mount **/exported/directory/** from the NFS server. Because no options are specified in this example, NFS uses default options.

**IMPORTANT**

The format of the **/etc/exports** file is very precise, particularly in regards to use of the space character. Remember to always separate exported file systems from hosts and hosts from one another with a space character. However, there should be no other space characters in the file except on comment lines.

For example, the following two lines do not mean the same thing:

```
/home bob.example.com(rw)
/home bob.example.com (rw)
```

The first line allows only users from **bob.example.com** read and write access to the **/home** directory. The second line allows users from **bob.example.com** to mount the directory as read-only (the default), while the rest of the world can mount it read/write.

**Default options**

The default options for an export entry are:

**ro**

The exported file system is read-only. Remote hosts cannot change the data shared on the file system. To allow hosts to make changes to the file system (that is, read and write), specify the **rw** option.

**sync**

The NFS server will not reply to requests before changes made by previous requests are written to disk. To enable asynchronous writes instead, specify the option **async**.

### wdelay

The NFS server will delay writing to the disk if it suspects another write request is imminent. This can improve performance as it reduces the number of times the disk must be accessed by separate write commands, thereby reducing write overhead. To disable this, specify the **no\_wdelay** option, which is available only if the default sync option is also specified.

### root\_squash

This prevents root users connected remotely (as opposed to locally) from having root privileges; instead, the NFS server assigns them the user ID **nfsnobody**. This effectively "squashes" the power of the remote root user to the lowest local user, preventing possible unauthorized writes on the remote server. To disable root squashing, specify the **no\_root\_squash** option.

To squash every remote user (including root), use the **all\_squash** option. To specify the user and group IDs that the NFS server should assign to remote users from a particular host, use the **anonuid** and **anongid** options, respectively, as in:

```
export host(anonuid=uid,anongid=gid)
```

Here, *uid* and *gid* are user ID number and group ID number, respectively. The **anonuid** and **anongid** options enable you to create a special user and group account for remote NFS users to share.

By default, access control lists (ACLs) are supported by NFS under Red Hat Enterprise Linux. To disable this feature, specify the **no\_acl** option when exporting the file system.

### Default and overridden options

Each default for every exported file system must be explicitly overridden. For example, if the **rw** option is not specified, then the exported file system is shared as read-only. The following is a sample line from **/etc/exports** which overrides two default options:

```
/another/exported/directory 192.168.0.3(rw,async)
```

In this example, **192.168.0.3** can mount **/another/exported/directory/** read and write, and all writes to disk are asynchronous.

## 7.5.2. The **exportfs** utility

The **exportfs** utility enables the root user to selectively export or unexport directories without restarting the NFS service. When given the proper options, the **exportfs** utility writes the exported file systems to **/var/lib/nfs/xtab**. Because the **nfs-mountd** service refers to the **xtab** file when deciding access privileges to a file system, changes to the list of exported file systems take effect immediately.

### Common **exportfs** options

The following is a list of commonly-used options available for **exportfs**:

**-r**

Causes all directories listed in **/etc/exports** to be exported by constructing a new export list in **/etc/lib/nfs/xtab**. This option effectively refreshes the export list with any changes made to **/etc/exports**.

**-a**

Causes all directories to be exported or unexported, depending on what other options are passed to **exportfs**. If no other options are specified, **exportfs** exports all file systems specified in **/etc/exports**.

**-o file-systems**

Specifies directories to be exported that are not listed in **/etc/exports**. Replace *file-systems* with additional file systems to be exported. These file systems must be formatted in the same way they are specified in **/etc/exports**. This option is often used to test an exported file system before adding it permanently to the list of exported file systems.

**-i**

Ignores **/etc/exports**; only options given from the command line are used to define exported file systems.

**-u**

Unexports all shared directories. The command **exportfs -ua** suspends NFS file sharing while keeping all NFS services up. To re-enable NFS sharing, use **exportfs -r**.

**-v**

Verbose operation, where the file systems being exported or unexported are displayed in greater detail when the **exportfs** command is executed.

If no options are passed to the **exportfs** utility, it displays a list of currently exported file systems.

### Additional resources

- For information on different methods for specifying host names, see [Section 7.4, “NFS host name formats”](#).
- For a complete list of export options, see the **exports(5)** man page.
- For more information about the **exportfs** utility, see the **exportfs(8)** man page.

## 7.6. NFS AND RPCBIND

This section explains the purpose of the **rpcbind** service, which is required by NFSv3.

The **rpcbind** service maps Remote Procedure Call (RPC) services to the ports on which they listen. RPC processes notify **rpcbind** when they start, registering the ports they are listening on and the RPC program numbers they expect to serve. The client system then contacts **rpcbind** on the server with a particular RPC program number. The **rpcbind** service redirects the client to the proper port number so it can communicate with the requested service.

Because RPC-based services rely on **rpcbind** to make all connections with incoming client requests, **rpcbind** must be available before any of these services start.

Access control rules for **rpcbind** affect all RPC-based services. Alternatively, it is possible to specify access control rules for each of the NFS RPC daemons.

### Additional resources

- For the precise syntax of access control rules, see the **rpc.mountd(8)** and **rpc.statd(8)** man pages.

## 7.7. INSTALLING NFS

This procedure installs all packages necessary to mount or export NFS shares.

### Procedure

- Install the **nfs-utils** package:

```
# yum install nfs-utils
```

## 7.8. STARTING THE NFS SERVER

This procedure describes how to start the NFS server, which is required to export NFS shares.

### Prerequisites

- For servers that support NFSv2 or NFSv3 connections, the **rpcbind** service must be running. To verify that **rpcbind** is active, use the following command:

```
$ systemctl status rpcbind
```

If the service is stopped, start and enable it:

```
$ systemctl enable --now rpcbind
```

### Procedure

- To start the NFS server and enable it to start automatically at boot, use the following command:

```
# systemctl enable --now nfs-server
```

### Additional resources

- To configure an NFSv4-only server, which does not require **rpcbind**, see [Section 7.13, “Configuring an NFSv4-only server”](#).

## 7.9. TROUBLESHOOTING NFS AND RPCBIND

Because the **rpcbind** service provides coordination between RPC services and the port numbers used to communicate with them, it is useful to view the status of current RPC services using **rpcbind** when troubleshooting. The **rpcinfo** utility shows each RPC-based service with port numbers, an RPC program number, a version number, and an IP protocol type (TCP or UDP).

### Procedure

1. To make sure the proper NFS RPC-based services are enabled for **rpcbind**, use the following command:

```
# rpcinfo -p
```

#### Example 7.2. rpcinfo -p command output

The following is sample output from this command:

```
program vers proto  port  service
```

```

100000 4 tcp 111 portmapper
100000 3 tcp 111 portmapper
100000 2 tcp 111 portmapper
100000 4 udp 111 portmapper
100000 3 udp 111 portmapper
100000 2 udp 111 portmapper
100005 1 udp 20048 mountd
100005 1 tcp 20048 mountd
100005 2 udp 20048 mountd
100005 2 tcp 20048 mountd
100005 3 udp 20048 mountd
100005 3 tcp 20048 mountd
100024 1 udp 37769 status
100024 1 tcp 49349 status
100003 3 tcp 2049 nfs
100003 4 tcp 2049 nfs
100227 3 tcp 2049 nfs_acl
100021 1 udp 56691 nlockmgr
100021 3 udp 56691 nlockmgr
100021 4 udp 56691 nlockmgr
100021 1 tcp 46193 nlockmgr
100021 3 tcp 46193 nlockmgr
100021 4 tcp 46193 nlockmgr

```

If one of the NFS services does not start up correctly, **rpcbind** will be unable to map RPC requests from clients for that service to the correct port.

2. In many cases, if NFS is not present in **rpcinfo** output, restarting NFS causes the service to correctly register with **rpcbind** and begin working:

```
# systemctl restart nfs-server
```

### Additional resources

- For more information and a list of **rpcinfo** options, see the **rpcinfo(8)** man page.
- To configure an NFSv4-only server, which does not require **rpcbind**, see [Section 7.13, “Configuring an NFSv4-only server”](#).

## 7.10. CONFIGURING THE NFS SERVER TO RUN BEHIND A FIREWALL

NFS requires the **rpcbind** service, which dynamically assigns ports for RPC services and can cause issues for configuring firewall rules. This procedure describes how to configure the NFS server to work behind a firewall.

### Procedure

1. To allow clients to access NFS shares behind a firewall, set which ports the RPC services run on in the **[mountd]** section of the **/etc/nfs.conf** file:

```
[mountd]
port=port-number
```

This adds the **-p *port-number*** option to the `rpc.mount` command line: `rpc.mount -p port-number`.

- To allow NFSv4.0 callbacks to pass through firewalls, set `/proc/sys/fs/nfs/nfs_callback_tcpport` and allow the server to connect to that port on the client.

This step is not needed for NFSv4.1 or higher, and the other ports for `mountd`, `statd`, and `lockd` are not required in a pure NFSv4 environment.

- To specify the ports to be used by the RPC service `nlockmgr`, set the port number for the `nlm_tcpport` and `nlm_udpport` options in the `/etc/modprobe.d/lockd.conf` file.
- Restart the NFS server:

```
# systemctl restart nfs-server
```

If NFS fails to start, check `/var/log/messages`. Commonly, NFS fails to start if you specify a port number that is already in use.

- Confirm the changes have taken effect:

```
# rpcinfo -p
```

### Additional resources

- To configure an NFSv4-only server, which does not require `rpcbind`, see [Section 7.13](#), “Configuring an NFSv4-only server”.

## 7.11. EXPORTING RPC QUOTA THROUGH A FIREWALL

If you export a file system that uses disk quotas, you can use the quota Remote Procedure Call (RPC) service to provide disk quota data to NFS clients.

### Procedure

- Enable and start the `rpc-rquotad` service:

```
# systemctl enable --now rpc-rquotad
```



#### NOTE

The `rpc-rquotad` service is, if enabled, started automatically after starting the `nfs-server` service.

- To make the quota RPC service accessible behind a firewall, the UDP or TCP port 875 need to be open. The default port number is defined in the `/etc/services` file. You can override the default port number by appending **-p *port-number*** to the `RPCRQUOTADOPTS` variable in the `/etc/sysconfig/rpc-rquotad` file.
- By default, remote hosts can only read quotas. If you want to allow clients to set quotas, append the **-S** option to the `RPCRQUOTADOPTS` variable in the `/etc/sysconfig/rpc-rquotad` file.

- Restart **rpc-rquotad** for the changes in the `/etc/sysconfig/rpc-rquotad` file to take effect:

```
# systemctl restart rpc-rquotad
```

## 7.12. ENABLING NFS OVER RDMA (NFSORDMA)

The remote direct memory access (RDMA) service works automatically in Red Hat Enterprise Linux 8 if there is RDMA-capable hardware present.

### Procedure

- Install the **rdma** and **rdma-core** packages:

```
# yum install rdma rdma-core
```

- To enable automatic loading of NFSoRDMA *server* modules, add the **SVCRDMA\_LOAD=yes** option on a new line in the `/etc/rdma/rdma.conf` configuration file.

The **rdma=20049** option in the `[nfsd]` section of the `/etc/nfs.conf` file specifies the port number on which the NFSoRDMA service listens for clients. The RFC 5667 standard specifies that servers must listen on port **20049** when providing NFSv4 services over RDMA.

The `/etc/rdma/rdma.conf` file contains a line that sets the **XPRTRDMA\_LOAD=yes** option by default, which requests the **rdma** service to load the NFSoRDMA *client* module.

- Restart the **nfs-server** service:

```
# systemctl restart nfs-server
```

### Additional resources

- The RFC 5667 standard: <https://tools.ietf.org/html/rfc5667>.

## 7.13. CONFIGURING AN NFSV4-ONLY SERVER

As an NFS server administrator, you can configure the NFS server to support only NFSv4, which minimizes the number of open ports and running services on the system.

### 7.13.1. Benefits and drawbacks of an NFSv4-only server

This section explains the benefits and drawbacks of configuring the NFS server to only support NFSv4.

By default, the NFS server supports NFSv2, NFSv3, and NFSv4 connections in Red Hat Enterprise Linux 8. However, you can also configure NFS to support only NFS version 4.0 and later. This minimizes the number of open ports and running services on the system, because NFSv4 does not require the **rpcbind** service to listen on the network.

When your NFS server is configured as NFSv4-only, clients attempting to mount shares using NFSv2 or NFSv3 fail with an error like the following:

```
Requested NFS version or transport protocol is not supported.
```



Optionally, you can also disable listening for the **RPCBIND**, **MOUNT**, and **NSM** protocol calls, which are not necessary in the NFSv4-only case.

The effects of disabling these additional options are:

- Clients that attempt to mount shares from your server using NFSv2 or NFSv3 become unresponsive.
- The NFS server itself is unable to mount NFSv2 and NFSv3 file systems.

### 7.13.2. NFS and rpcbind

This section explains the purpose of the **rpcbind** service, which is required by NFSv3.

The **rpcbind** service maps Remote Procedure Call (RPC) services to the ports on which they listen. RPC processes notify **rpcbind** when they start, registering the ports they are listening on and the RPC program numbers they expect to serve. The client system then contacts **rpcbind** on the server with a particular RPC program number. The **rpcbind** service redirects the client to the proper port number so it can communicate with the requested service.

Because RPC-based services rely on **rpcbind** to make all connections with incoming client requests, **rpcbind** must be available before any of these services start.

Access control rules for **rpcbind** affect all RPC-based services. Alternatively, it is possible to specify access control rules for each of the NFS RPC daemons.

#### Additional resources

- For the precise syntax of access control rules, see the **rpc.mountd(8)** and **rpc.statd(8)** man pages.

### 7.13.3. Configuring the NFS server to support only NFSv4

This procedure describes how to configure your NFS server to support only NFS version 4.0 and later.

#### Procedure

1. Disable NFSv2 and NFSv3 by adding the following lines to the **[nfsd]** section of the **/etc/nfs.conf** configuration file:

```
[nfsd]
vers2=no
vers3=no
```

2. Optionally, disable listening for the **RPCBIND**, **MOUNT**, and **NSM** protocol calls, which are not necessary in the NFSv4-only case. Disable related services:

```
# systemctl mask --now rpc-statd.service rpcbind.service
rpcbind.socket
```

3. Restart the NFS server:

```
# systemctl restart nfs-server
```

The changes take effect as soon as you start or restart the NFS server.

### 7.13.4. Verifying the NFSv4-only configuration

This procedure describes how to verify that your NFS server is configured in the NFSv4-only mode by using the **netstat** utility.

#### Procedure

- Use the **netstat** utility to list services listening on the TCP and UDP protocols:

```
# netstat --listening --tcp --udp
```

#### Example 7.3. Output on an NFSv4-only server

The following is an example **netstat** output on an NFSv4-only server; listening for **RPCBIND**, **MOUNT**, and **NSM** is also disabled. Here, **nfs** is the only listening NFS service:

```
# netstat --listening --tcp --udp

Active Internet connections (only servers)
Proto Recv-Q Send-Q Local Address           Foreign Address         State
tcp        0      0 0.0.0.0:ssh             0.0.0.0:*               LISTEN
tcp        0      0 0.0.0.0:nfs             0.0.0.0:*               LISTEN
tcp6       0      0 [::]:ssh               [::]:*                  LISTEN
tcp6       0      0 [::]:nfs                [::]:*                  LISTEN
udp        0      0 localhost.locald:bootpc 0.0.0.0:*
```

#### Example 7.4. Output before configuring an NFSv4-only server

In comparison, the **netstat** output before configuring an NFSv4-only server includes the **sunrpc** and **mountd** services:

```
# netstat --listening --tcp --udp

Active Internet connections (only servers)
Proto Recv-Q Send-Q Local Address           Foreign Address         State
tcp        0      0 0.0.0.0:ssh             0.0.0.0:*               LISTEN
tcp        0      0 0.0.0.0:40189           0.0.0.0:*               LISTEN
tcp        0      0 0.0.0.0:46813           0.0.0.0:*               LISTEN
tcp        0      0 0.0.0.0:nfs             0.0.0.0:*               LISTEN
tcp        0      0 0.0.0.0:sunrpc           0.0.0.0:*               LISTEN
tcp        0      0 0.0.0.0:mountd           0.0.0.0:*
```

```

LISTEN
tcp6      0      0 [::]:ssh          [::]:*
LISTEN
tcp6      0      0 [::]:51227        [::]:*
LISTEN
tcp6      0      0 [::]:nfs          [::]:*
LISTEN
tcp6      0      0 [::]:sunrpc       [::]:*
LISTEN
tcp6      0      0 [::]:mountd       [::]:*
LISTEN
tcp6      0      0 [::]:45043        [::]:*
LISTEN
udp       0      0 localhost:1018    0.0.0.0:*
udp       0      0 localhost.locald:bootpc 0.0.0.0:*
udp       0      0 0.0.0.0:mountd   0.0.0.0:*
udp       0      0 0.0.0.0:46672    0.0.0.0:*
udp       0      0 0.0.0.0:sunrpc   0.0.0.0:*
udp       0      0 0.0.0.0:33494    0.0.0.0:*
udp6     0      0 [::]:33734        [::]:*
udp6     0      0 [::]:mountd       [::]:*
udp6     0      0 [::]:sunrpc       [::]:*
udp6     0      0 [::]:40243        [::]:*

```

## 7.14. RELATED INFORMATION

- The Linux NFS wiki: <http://linux-nfs.org>

# CHAPTER 8. MANAGING LAYERED LOCAL STORAGE WITH STRATIS

You can easily set up and manage complex storage configurations integrated by the Stratis high-level system.

## 8.1. SETTING UP STRATIS FILE SYSTEMS

As a system administrator, you can enable and set up the Stratis volume-managing file system on your system to easily manage layered storage.

### 8.1.1. The purpose and features of Stratis

Stratis is a local storage-management solution for Linux. It is focused on simplicity and ease of use, and gives you access to advanced storage features.

Stratis makes the following activities easier:

- Initial configuration of storage
- Making changes later
- Using advanced storage features

Stratis is a hybrid user-and-kernel local storage management system that supports advanced storage features. The central concept of Stratis is a storage *pool*. This pool is created from one or more local disks or partitions, and volumes are created from the pool.

The pool enables many useful features, such as:

- File system snapshots
- Thin provisioning
- Tiering

### 8.1.2. Components of a Stratis volume

Externally, Stratis presents the following volume components in the command-line interface and the API:

#### **blockdev**

Block devices, such as a disk or a disk partition.

#### **pool**

Composed of one or more block devices.

A pool has a fixed total size, equal to the size of the block devices.

The pool contains most Stratis layers, such as the non-volatile data cache using the **dm-cache** target.

Stratis creates a **/stratis/my-pool/** directory for each pool. This directory contains links to devices that represent Stratis file systems in the pool.

#### **filesystem**

Each pool can contain one or more file systems, which store files.

File systems are thinly provisioned and do not have a fixed total size. The actual size of a file system grows with the data stored on it. If the size of the data approaches the virtual size of the file system, Stratis grows the thin volume and the file system automatically.

The file systems are formatted with XFS.



### IMPORTANT

Stratis tracks information about file systems created using Stratis that XFS is not aware of, and changes made using XFS do not automatically create updates in Stratis. Users must not reformat or reconfigure XFS file systems that are managed by Stratis.

Stratis creates links to file systems at the `/stratis/my-pool/my-fs` path.



### NOTE

Stratis uses many Device Mapper devices, which show up in `dmsetup` listings and the `/proc/partitions` file. Similarly, the `lsblk` command output reflects the internal workings and layers of Stratis.

## 8.1.3. Block devices usable with Stratis

This section lists storage devices that you can use for Stratis.

### Supported devices

Stratis pools have been tested to work on these types of block devices:

- LUKS
- LVM logical volumes
- MDRAID
- DM Multipath
- iSCSI
- HDDs and SSDs
- NVMe devices



### WARNING

In the current version, Stratis does not handle failures in hard drives or other hardware. If you create a Stratis pool over multiple hardware devices, you increase the risk of data loss because multiple devices must be operational to access the data.

### Unsupported devices

Because Stratis contains a thin-provisioning layer, Red Hat does not recommend placing a Stratis pool on block devices that are already thinly-provisioned.

### Additional resources

- For iSCSI and other block devices requiring network, see the **systemd.mount(5)** man page for information on the **\_netdev** mount option.

## 8.1.4. Installing Stratis

This procedure installs all packages necessary to use Stratis.

### Procedure

1. Install packages that provide the Stratis service and command-line utilities:

```
# yum install stratisd stratis-cli
```

2. Make sure that the **stratisd** service is enabled:

```
# systemctl enable --now stratisd
```

## 8.1.5. Creating a Stratis pool

This procedure creates a Stratis pool from one or more block devices.

### Prerequisites

- Stratis is installed. See [Section 8.1.4, “Installing Stratis”](#).
- The **stratisd** service is running.
- The block devices on which you are creating a Stratis pool are not in use and not mounted.
- The block devices on which you are creating a Stratis pool are at least 1 GiB in size each.

### Procedure

1. If the selected block device contains file system, partition table, or RAID signatures, erase them:

```
# wipefs --all block-device
```

Replace *block-device* with the path to a block device, such as **/dev/sdb**.

2. To create a Stratis pool on the block device, use:

```
# stratis pool create my-pool block-device
```

- Replace *my-pool* with an arbitrary name for the pool.
- Replace *block-device* with the path to the empty or wiped block device, such as **/dev/sdb**.

To create a pool from more than one block device, list them all on the command line:

```
# stratis pool create my-pool device-1 device-2 device-n
```

- To verify, list all pools on your system:

```
# stratis pool list
```

#### Additional resources

- The **stratis(8)** man page

#### Next steps

- Create a Stratis file system on the pool. See [Section 8.1.6, “Creating a Stratis file system”](#).

### 8.1.6. Creating a Stratis file system

This procedure creates a Stratis file system on an existing Stratis pool.

#### Prerequisites

- Stratis is installed. See [Section 8.1.4, “Installing Stratis”](#).
- The **stratisd** service is running.
- You have created a Stratis pool. See [Section 8.1.5, “Creating a Stratis pool”](#).

#### Procedure

- To create a Stratis file system on a pool, use:

```
# stratis fs create my-pool my-fs
```

- Replace *my-pool* with the name of your existing Stratis pool.
- Replace *my-fs* with an arbitrary name for the file system.

- To verify, list file systems within the pool:

```
# stratis fs list my-pool
```

#### Additional resources

- The **stratis(8)** man page

#### Next steps

- Mount the Stratis file system. See [Section 8.1.7, “Mounting a Stratis file system”](#).

### 8.1.7. Mounting a Stratis file system

This procedure mounts an existing Stratis file system to access the content.

#### Prerequisites

- Stratis is installed. See [Section 8.1.4, “Installing Stratis”](#).

- The **stratisd** service is running.
- You have created a Stratis file system. See [Section 8.1.6, “Creating a Stratis file system”](#).

### Procedure

- To mount the file system, use the entries that Stratis maintains in the **/stratis/** directory:

```
# mount /stratis/my-pool/my-fs mount-point
```

The file system is now mounted on the *mount-point* directory and ready to use.

### Additional resources

- The **mount (8)** man page

## 8.1.8. Persistently mounting a Stratis file system

This procedure persistently mounts a Stratis file system so that it is available automatically after booting the system.

### Prerequisites

- Stratis is installed. See [Section 8.1.4, “Installing Stratis”](#).
- The **stratisd** service is running.
- You have created a Stratis file system. See [Section 8.1.6, “Creating a Stratis file system”](#).

### Procedure

1. Determine the UUID attribute of the file system:

```
$ lsblk --output=UUID /stratis/my-pool/my-fs
```

For example:

#### Example 8.1. Viewing the UUID of Stratis file system

```
$ lsblk --output=UUID /stratis/my-pool/fs1
```

```
UUID
a1f0b64a-4ebb-4d4e-9543-b1d79f600283
```

2. If the mount point directory does not exist, create it:

```
# mkdir --parents mount-point
```

3. As root, edit the **/etc/fstab** file and add a line for the file system, identified by the UUID. Use **xfs** as the file system type and add the **x-systemd.requires=stratisd.service** option. For example:

#### Example 8.2. The /fs1 mount point in /etc/fstab



```

| | UUID=a1f0b64a-4ebb-4d4e-9543-b1d79f600283 /fs1 xfs defaults,x-
| | systemd.requires=stratisd.service 0 0

```

4. Regenerate mount units so that your system registers the new configuration:

```

| # systemctl daemon-reload

```

5. Try mounting the file system to verify that the configuration works:

```

| # mount mount-point

```

### Additional resources

- [Section 4.8, “Persistently mounting file systems”](#)

### 8.1.9. Related information

- The *Stratis Storage* website: <https://stratis-storage.github.io/>

## 8.2. EXTENDING A STRATIS VOLUME WITH ADDITIONAL BLOCK DEVICES

You can attach additional block devices to a Stratis pool to provide more storage capacity for Stratis file systems.

### 8.2.1. Components of a Stratis volume

Externally, Stratis presents the following volume components in the command-line interface and the API:

#### **blockdev**

Block devices, such as a disk or a disk partition.

#### **pool**

Composed of one or more block devices.

A pool has a fixed total size, equal to the size of the block devices.

The pool contains most Stratis layers, such as the non-volatile data cache using the **dm-cache** target.

Stratis creates a **/stratis/my-pool/** directory for each pool. This directory contains links to devices that represent Stratis file systems in the pool.

#### **filesystem**

Each pool can contain one or more file systems, which store files.

File systems are thinly provisioned and do not have a fixed total size. The actual size of a file system grows with the data stored on it. If the size of the data approaches the virtual size of the file system, Stratis grows the thin volume and the file system automatically.

The file systems are formatted with XFS.



## IMPORTANT

Stratis tracks information about file systems created using Stratis that XFS is not aware of, and changes made using XFS do not automatically create updates in Stratis. Users must not reformat or reconfigure XFS file systems that are managed by Stratis.

Stratis creates links to file systems at the `/stratis/my-pool/my-fs` path.



## NOTE

Stratis uses many Device Mapper devices, which show up in `dmsetup` listings and the `/proc/partitions` file. Similarly, the `lsblk` command output reflects the internal workings and layers of Stratis.

## 8.2.2. Adding block devices to a Stratis pool

This procedure adds one or more block devices to a Stratis pool to be usable by Stratis file systems.

### Prerequisites

- Stratis is installed. See [Section 8.1.4, “Installing Stratis”](#).
- The `stratisd` service is running.
- The block devices that you are adding to the Stratis pool are not in use and not mounted.
- The block devices that you are adding to the Stratis pool are at least 1 GiB in size each.

### Procedure

- To add one or more block devices to the pool, use:

```
# stratis pool add-data my-pool device-1 device-2 device-n
```

### Additional resources

- The `stratis(8)` man page

## 8.2.3. Related information

- The *Stratis Storage* website: <https://stratis-storage.github.io/>

## 8.3. MONITORING STRATIS FILE SYSTEMS

As a Stratis user, you can view information about Stratis volumes on your system to monitor their state and free space.

### 8.3.1. Stratis sizes reported by different utilities

This section explains the difference between Stratis sizes reported by standard utilities such as `df` and the `stratis` utility.

Standard Linux utilities such as **df** report the size of the XFS file system layer on Stratis, which is 1 TiB. This is not useful information, because the actual storage usage of Stratis is less due to thin provisioning, and also because Stratis automatically grows the file system when the XFS layer is close to full.



## IMPORTANT

Regularly monitor the amount of data written to your Stratis file systems, which is reported as the *Total Physical Used* value. Make sure it does not exceed the *Total Physical Size* value.

### Additional resources

- The **stratis(8)** man page

## 8.3.2. Displaying information about Stratis volumes

This procedure lists statistics about your Stratis volumes, such as the total, used, and free size or file systems and block devices belonging to a pool.

### Prerequisites

- Stratis is installed. See [Section 8.1.4, “Installing Stratis”](#).
- The **stratisd** service is running.

### Procedure

- To display information about all **block devices** used for Stratis on your system:

```
# stratis blockdev

Pool Name  Device Node      Physical Size  State  Tier
my-pool    /dev/sdb         9.10 TiB      In-use  Data
```

- To display information about all Stratis **pools** on your system:

```
# stratis pool

Name      Total Physical Size  Total Physical Used
my-pool   9.10 TiB             598 MiB
```

- To display information about all Stratis **file systems** on your system:

```
# stratis filesystem

Pool Name  Name  Used      Created                Device
my-pool    my-fs 546 MiB   Nov 08 2018 08:03     /stratis/my-pool/my-fs
```

### Additional resources

- The **stratis(8)** man page

## 8.3.3. Related information

- The *Stratis Storage* website: <https://stratis-storage.github.io/>

## 8.4. USING SNAPSHOTS ON STRATIS FILE SYSTEMS

You can use snapshots on Stratis file systems to capture file system state at arbitrary times and restore it in the future.

### 8.4.1. Characteristics of Stratis snapshots

This section describes the properties and limitations of file system snapshots on Stratis.

In Stratis, a snapshot is a regular Stratis file system created as a copy of another Stratis file system. The snapshot initially contains the same file content as the original file system, but can change as the snapshot is modified. Whatever changes you make to the snapshot will not be reflected in the original file system.

The current snapshot implementation in Stratis is characterized by the following:

- A snapshot of a file system is another file system.
- A snapshot and its origin are not linked in lifetime. A snapshotted file system can live longer than the file system it was created from.
- A file system does not have to be mounted to create a snapshot from it.
- Each snapshot uses around half a gigabyte of actual backing storage, which is needed for the XFS log.

### 8.4.2. Creating a Stratis snapshot

This procedure creates a Stratis file system as a snapshot of an existing Stratis file system.

#### Prerequisites

- Stratis is installed. See [Section 8.1.4, “Installing Stratis”](#).
- The **stratisd** service is running.
- You have created a Stratis file system. See [Section 8.1.6, “Creating a Stratis file system”](#).

#### Procedure

- To create a Stratis snapshot, use:

```
# stratis fs snapshot my-pool my-fs my-fs-snapshot
```

#### Additional resources

- The **stratis(8)** man page

### 8.4.3. Accessing the content of a Stratis snapshot

This procedure mounts a snapshot of a Stratis file system to make it accessible for read and write operations.

## Prerequisites

- Stratis is installed. See [Section 8.1.4, “Installing Stratis”](#).
- The **stratisd** service is running.
- You have created a Stratis snapshot. See [Section 8.4.2, “Creating a Stratis snapshot”](#).

## Procedure

- To access the snapshot, mount it as a regular file system from the `/stratis/my-pool/` directory:

```
# mount /stratis/my-pool/my-fs-snapshot mount-point
```

## Additional resources

- [Section 8.1.7, “Mounting a Stratis file system”](#)
- The **mount(8)** man page

## 8.4.4. Reverting a Stratis file system to a previous snapshot

This procedure reverts the content of a Stratis file system to the state captured in a Stratis snapshot.

### Prerequisites

- Stratis is installed. See [Section 8.1.4, “Installing Stratis”](#).
- The **stratisd** service is running.
- You have created a Stratis snapshot. See [Section 8.4.2, “Creating a Stratis snapshot”](#).

### Procedure

1. Optionally, back up the current state of the file system to be able to access it later:

```
# stratis filesystem snapshot my-pool my-fs my-fs-backup
```

2. Unmount and remove the original file system:

```
# umount /stratis/my-pool/my-fs
# stratis filesystem destroy my-pool my-fs
```

3. Create a copy of the snapshot under the name of the original file system:

```
# stratis filesystem snapshot my-pool my-fs-snapshot my-fs
```

4. Mount the snapshot, which is now accessible with the same name as the original file system:

```
# mount /stratis/my-pool/my-fs mount-point
```

The content of the file system named `my-fs` is now identical to the snapshot `my-fs-snapshot`.

## Additional resources

- The **stratis(8)** man page

### 8.4.5. Removing a Stratis snapshot

This procedure removes a Stratis snapshot from a pool. Data on the snapshot are lost.

#### Prerequisites

- Stratis is installed. See [Section 8.1.4, “Installing Stratis”](#).
- The **stratisd** service is running.
- You have created a Stratis snapshot. See [Section 8.4.2, “Creating a Stratis snapshot”](#).

#### Procedure

1. Unmount the snapshot:

```
# umount /stratis/my-pool/my-fs-snapshot
```

2. Destroy the snapshot:

```
# stratis filesystem destroy my-pool my-fs-snapshot
```

#### Additional resources

- The **stratis(8)** man page

### 8.4.6. Related information

- The *Stratis Storage* website: <https://stratis-storage.github.io/>

## 8.5. REMOVING STRATIS FILE SYSTEMS

You can remove an existing Stratis file system or a Stratis pool, destroying data on them.

### 8.5.1. Components of a Stratis volume

Externally, Stratis presents the following volume components in the command-line interface and the API:

#### **blockdev**

Block devices, such as a disk or a disk partition.

#### **pool**

Composed of one or more block devices.

A pool has a fixed total size, equal to the size of the block devices.

The pool contains most Stratis layers, such as the non-volatile data cache using the **dm-cache** target.

Stratis creates a **/stratis/my-pool/** directory for each pool. This directory contains links to devices that represent Stratis file systems in the pool.

## filesystem

Each pool can contain one or more file systems, which store files.

File systems are thinly provisioned and do not have a fixed total size. The actual size of a file system grows with the data stored on it. If the size of the data approaches the virtual size of the file system, Stratis grows the thin volume and the file system automatically.

The file systems are formatted with XFS.



### IMPORTANT

Stratis tracks information about file systems created using Stratis that XFS is not aware of, and changes made using XFS do not automatically create updates in Stratis. Users must not reformat or reconfigure XFS file systems that are managed by Stratis.

Stratis creates links to file systems at the `/stratis/my-pool/my-fs` path.



### NOTE

Stratis uses many Device Mapper devices, which show up in `dmsetup` listings and the `/proc/partitions` file. Similarly, the `lsblk` command output reflects the internal workings and layers of Stratis.

## 8.5.2. Removing a Stratis file system

This procedure removes an existing Stratis file system. Data stored on it are lost.

### Prerequisites

- Stratis is installed. See [Section 8.1.4, “Installing Stratis”](#).
- The `stratisd` service is running.
- You have created a Stratis file system. See [Section 8.1.6, “Creating a Stratis file system”](#).

### Procedure

1. Unmount the file system:

```
# umount /stratis/my-pool/my-fs
```

2. Destroy the file system:

```
# stratis filesystem destroy my-pool my-fs
```

3. Verify that the file system no longer exists:

```
# stratis filesystem list my-pool
```

### Additional resources

- The `stratis(8)` man page

### 8.5.3. Removing a Stratis pool

This procedure removes an existing Stratis pool. Data stored on it are lost.

#### Prerequisites

- Stratis is installed. See [Section 8.1.4, “Installing Stratis”](#).
- The `stratisd` service is running.
- You have created a Stratis pool. See [Section 8.1.5, “Creating a Stratis pool”](#).

#### Procedure

1. List file systems on the pool:

```
# stratis filesystem list my-pool
```

2. Unmount all file systems on the pool:

```
# umount /stratis/my-pool/my-fs-1 \  
          /stratis/my-pool/my-fs-2 \  
          /stratis/my-pool/my-fs-n
```

3. Destroy the file systems:

```
# stratis filesystem destroy my-pool my-fs-1 my-fs-2
```

4. Destroy the pool:

```
# stratis pool destroy my-pool
```

5. Verify that the pool no longer exists:

```
# stratis pool list
```

#### Additional resources

- The `stratis(8)` man page

### 8.5.4. Related information

- The *Stratis Storage* website: <https://stratis-storage.github.io/>