



Full credit is given to the above companies including the OS that this PDF file was generated!

Rocky Enterprise Linux 9.2 Manual Pages on command 'xfs_admin.8'

\$ man xfs_admin.8

xfs_admin(8) System Manager's Manual xfs_admin(8)

NAME

xfs_admin - change parameters of an XFS filesystem

SYNOPSIS

```
xfs_admin [ -eflpu ] [ -O featurelist ] [ -c 0|1 ] [ -L label ] [ -U
uuid ] [ -r rtdev ] device [ logdev ]
xfs_admin -V
```

DESCRIPTION

xfs_admin uses the xfs_db(8) command to modify various parameters of a filesystem.

Devices that are mounted cannot be modified. Administrators must unmount filesystems before xfs_admin or xfs_db(8) can convert parameters.

A number of parameters of a mounted filesystem can be examined and modified using the xfs_growfs(8) command.

The optional logdev parameter specifies the device special file where the filesystem's external log resides. This is required only for filesystems that use an external log. See the mkfs.xfs -l option, and refer to xfs(5) for a detailed description of the XFS log.

OPTIONS

- e Enables unwritten extent support on a filesystem that does not already have this enabled (for legacy filesystems, it can't be disabled anymore at mkfs time).

This option only applies to the deprecated V4 format.

- f Specifies that the filesystem image to be processed is stored in a regular file at device (see the mkfs.xfs -d file option).

- j Enables version 2 log format (journal format supporting larger log buffers).

This option only applies to the deprecated V4 format.

- l Print the current filesystem label.

- p Enable 32bit project identifier support (PROJID32BIT feature).

This option only applies to the deprecated V4 format.

- u Print the current filesystem UUID (Universally Unique Identifier).

- c 0|1 Enable (1) or disable (0) lazy-counters in the filesystem.

Lazy-counters may not be disabled on Version 5 superblock filesystems (i.e. those with metadata CRCs enabled).

In other words, this option only applies to the deprecated V4 format.

This operation may take quite a bit of time on large filesystems as the entire filesystem needs to be scanned when this option is changed.

With lazy-counters enabled, the superblock is not modified or logged on every change of the free-space and inode counters. Instead, enough information is kept in other parts of the filesystem to be able to maintain the counter values without needing to keep them in the superblock. This gives significant improvements in performance on some configurations and metadata intensive workloads.

- L label

Set the filesystem label to label. XFS filesystem labels can be at most 12 characters long; if label is longer than 12 characters?

ters, xfs_admin will truncate it and print a warning message.

The filesystem label can be cleared using the special "--" value for label.

-O feature1=status,feature2=status...

Add or remove features on an existing V5 filesystem. The features should be specified as a comma-separated list. status should be either 0 to disable the feature or 1 to enable the feature. Note, however, that most features cannot be disabled.

NOTE: Administrators must ensure the filesystem is clean by running xfs_repair -n to inspect the filesystem before performing the upgrade. If corruption is found, recovery procedures (e.g. reformat followed by restoration from backup; or running xfs_repair without the -n) must be followed to clean the filesystem.

Supported features are as follows:

inobtcount

Keep a count the number of blocks in each inode btree in the AGI. This reduces mount time by speeding up metadata space reservation calculations. The filesystem cannot be downgraded after this feature is enabled. Once enabled, the filesystem will not be writable by older kernels. This feature was added to Linux 5.10.

bigtime

Upgrade a filesystem to support larger timestamps up to the year 2486. The filesystem cannot be downgraded after this feature is enabled. Once enabled, the filesystem will not be mountable by older kernels. This feature was added to Linux 5.10.

-U uuid

Set the UUID of the filesystem to uuid. A sample UUID looks like this: "c1b9d5a2-f162-11cf-9ece-0020afc76f16". The uuid may also be nil, which will set the filesystem UUID to the null UUID. The uuid may also be generate, which will generate a new UUID for the filesystem. Note that on CRC-enabled filesystems,

this will set an incompatible flag such that older kernels will not be able to mount the filesystem. To remove this incompatible flag, use restore, which will restore the original UUID and remove the incompatible feature flag as needed.

-r rtdev

Specifies the device special file where the filesystem's realtime section resides. Only for those filesystems which use a realtime section.

-V Prints the version number and exits.

The mount(8) manual entry describes how to mount a filesystem using its label or UUID, rather than its block special device name.

SEE ALSO

mkfs.xfs(8), mount(8), xfs_db(8), xfs_growfs(8), xfs_repair(8), xfs(5).

 xfs_admin(8)