

NAME

xfs_scrub – check and repair the contents of a mounted XFS filesystem

SYNOPSIS

xfs_scrub [**-abCemTVx**] *mount-point*

xfs_scrub -V

DESCRIPTION

xfs_scrub attempts to check and repair all metadata in a mounted XFS filesystem.

WARNING! This program is **EXPERIMENTAL**, which means that its behavior and interface could change at any time!

xfs_scrub asks the kernel to scrub all metadata objects in the filesystem. Metadata records are scanned for obviously bad values and then cross-referenced against other metadata. The goal is to establish a reasonable confidence about the consistency of the overall filesystem by examining the consistency of individual metadata records against the other metadata in the filesystem. Damaged metadata can be rebuilt from other metadata if there exists redundant data structures which are intact.

Filesystem corruption and optimization opportunities will be logged to the standard error stream. Enabling verbose mode will increase the amount of status information sent to the output.

If the kernel scrub reports that metadata needs repairs or optimizations and the user does not pass **-n** on the command line, this program will ask the kernel to make the repairs and to perform the optimizations. See the sections about optimizations and repairs for a list of optimizations and repairs known to this program. The kernel may not support repairing or optimizing the filesystem. If this is the case, the filesystem must be unmounted and **xfs_repair(8)** run on the filesystem to fix the problems.

OPTIONS

-a errors

Abort if more than this many errors are found on the filesystem.

-b Run in background mode. If the option is specified once, only run a single scrubbing thread at a time. If given more than once, an artificial delay of 100us is added to each scrub call to reduce CPU overhead even further.

-C fd This option causes **xfs_scrub** to write progress information to the specified file description so that the progress of the filesystem check can be monitored. If the file description is a tty, a fancy progress bar is rendered. Otherwise, a simple numeric status dump compatible with the **fsck -C** format is output.

-e Specifies what happens when errors are detected. If *shutdown* is given, the filesystem will be taken offline if errors are found. If *continue* is given, no action is taken if errors are found; this is the default behavior.

-k Do not call TRIM on the free space.

-m file Search this file for mounted filesystems instead of */etc/mtab*.

-n Only check filesystem metadata. Do not repair or optimize anything.

-T Print timing and memory usage information for each phase.

-v Enable verbose mode, which prints periodic status updates.

-V Prints the version number and exits.

-x Read all file data extents to look for disk errors. **xfs_scrub** will issue **O_DIRECT** reads to the block device directly. If the block device is a SCSI disk, it will instead issue **READ VERIFY** commands directly to the disk. If media errors are found, the error report will include the disk offset, in bytes. If the media errors affect a file, the report will also include the inode number and file offset, in bytes. These actions will confirm that all file data blocks can be read from storage.

OPTIMIZATIONS

Optimizations supported by this program include, but are not limited to:

- Instructing the underlying storage to discard unused extents via the **TRIM** ioctl.
- Updating secondary superblocks to match the primary superblock.
- Turning off shared block write checks for files that no longer share blocks.

REPAIRS

Repairs are performed by calling into the kernel. This limits the scope of repair activities to rebuilding primary data structures from secondary data structures, or secondary structures from primary structures. The existence of secondary data structures may require features that can only be turned on from **mkfs.xfs(8)**. If errors cannot be repaired, the filesystem must be unmounted and **xfs_repair(8)** run. Repairs supported by the kernel include, but are not limited to:

- Reconstructing extent allocation data.
- Rebuilding free space information.
- Rebuilding inode indexes.
- Fixing minor corruptions of inode records.
- Recalculating reference count information.
- Reconstructing reverse mapping data from primary extent allocation data.
- Scheduling a quotacheck for the next mount.

If corrupt metadata is successfully repaired, this program will log that a repair has succeeded instead of a corruption report.

EXIT CODE

The exit code returned by **xfs_scrub** is the sum of the following conditions:

- | | |
|---|---------------------------------------|
| 0 | – No errors |
| 1 | – File system errors left uncorrected |
| 2 | – File system optimizations possible |
| 4 | – Operational error |
| 8 | – Usage or syntax error |

CAVEATS

xfs_scrub is an immature utility! Do not run this program unless you have backups of your data! This program takes advantage of in-kernel scrubbing to verify a given data structure with locks held and can keep the filesystem busy for a long time. The kernel must be new enough to support the **SCRUB_METADATA** ioctl.

If errors are found and cannot be repaired, the filesystem must be unmounted and repaired.

SEE ALSO

xfs_repair(8).