

NAME

zpool-scrub - begin or resume scrub of ZFS storage pools

SYNOPSIS

zpool scrub [-s|-p] [-w] [-e] *pool*<?>

DESCRIPTION

Begins a scrub or resumes a paused scrub. The scrub examines all data in the specified pools to verify that it checksums correctly. For replicated (mirror, raidz, or draid) devices, ZFS automatically repairs any damage discovered during the scrub. The **zpool status** command reports the progress of the scrub and summarizes the results of the scrub upon completion.

Scrubbing and resilvering are very similar operations. The difference is that resilvering only examines data that ZFS knows to be out of date (for example, when attaching a new device to a mirror or replacing an existing device), whereas scrubbing examines all data to discover silent errors due to hardware faults or disk failure.

When scrubbing a pool with encrypted filesystems the keys do not need to be loaded. However, if the keys are not loaded and an unrepairable checksum error is detected the file name cannot be included in the **zpool status -v** verbose error report.

Because scrubbing and resilvering are I/O-intensive operations, ZFS only allows one at a time.

A scrub is split into two parts: metadata scanning and block scrubbing. The metadata scanning sorts blocks into large sequential ranges which can then be read much more efficiently from disk when issuing the scrub I/O.

If a scrub is paused, the **zpool scrub** resumes it. If a resilver is in progress, ZFS does not allow a scrub to be started until the resilver completes.

Note that, due to changes in pool data on a live system, it is possible for scrubs to progress slightly beyond 100% completion. During this period, no completion time estimate will be provided.

OPTIONS

-s Stop scrubbing.

-p

Pause scrubbing. Scrub pause state and progress are periodically synced to disk. If the system is restarted or pool is exported during a paused scrub, even after import, scrub will remain paused until it is resumed. Once resumed the scrub will pick up from the place where it was last checkpointed to

disk. To resume a paused scrub issue **zpool scrub** or **zpool scrub -e** again.

-w

Wait until scrub has completed before returning.

-e

Only scrub files with known data errors as reported by **zpool status -v**. The pool must have been scrubbed at least once with the **head_errlog** feature enabled to use this option. Error scrubbing cannot be run simultaneously with regular scrubbing or resilvering, nor can it be run when a regular scrub is paused.

EXAMPLES**Example 1**

Status of pool with ongoing scrub:

zpool status

...

scan: scrub in progress since Sun Jul 25 16:07:49 2021

403M / 405M scanned at 100M/s, 68.4M / 405M issued at 10.0M/s

0B repaired, 16.91% done, 00:00:04 to go

...

Where metadata which references 403M of file data has been scanned at 100M/s, and 68.4M of that file data has been scrubbed sequentially at 10.0M/s.

PERIODIC SCRUB

On machines using systemd, scrub timers can be enabled on per-pool basis. **weekly** and **monthly** timer units are provided.

```
systemctl enable zfs-scrub-weekly@rpool.timer --now
```

```
systemctl enable zfs-scrub-monthly@otherpool.timer --now
```

SEE ALSO

systemd.timer(5), zpool-iostat(8), zpool-resilver(8), zpool-status(8)