

ZFS metadata corruption

ZFS performs scrub tasks on a regular basis in which it reads each and every file, and compares that file to the checksum. If the checksum is different from the data on the disk, it will attempt to repair that data using the redundancy that has been setup with RAID. If it cannot do this, it will log a corruption error which can be seen using the `zpool status -v <poolname>` command. This error will remain until the file is removed and another scrub task is performed.

If the file is gone, but snapshots still reference that file, then an identifier will be used such as `<0x00xxx>:<0xXX00>`. The zpool will also report as being unhealthy.

If the actual ZFS pool metadata is corrupted, that is another story. If crucial metadata is corrupted then the pool will report as being corrupted, and may even fail to import upon boot. A pool with corrupted metadata should be backed up immediately, destroyed and restored from backup.

ZFS metadata corruption can report as being `<metadata:0xXXxx>`

Further reading [zfs pool metadata corrupt](#)

Revision #1

Created 10 April 2024 05:51:26 by ColtM

Updated 10 April 2024 06:12:33 by ColtM