



5 Reasons ZFS is Better Than Traditional RAID

1. Protects against bit-rot and I/O errors

Unlike traditional RAID which blindly trusts its stored data, ZFS stores a 256-bit checksum when writing data to disk and checksums the metadata, itself. ZFS will then verify the integrity of all data whenever reading data from disk, (including regularly scheduled data scrubs) and will automatically repair any corrupt data it detects.

2. Copy on Write

Data is always written to a new block and checksums calculated prior to changing any data pointers, ensuring that existing data remains fully in-tact until the new data is fully written to disk, preventing data loss in the event of an unexpected reboot or power loss.

3. Snapshot Performance

Traditional RAID performance suffers whenever snapshots are taken, but because of the Copy on Write method of writing data, ZFS snapshots can be taken without a performance impact, resulting in very fast snapshot creation.

4. Low Snapshot Space Impact

ZFS also results in very little space being required for many, many data snapshots, only requiring block level changes to be logged. This also results in near instantaneous snapshot recovery.

5. Seamless Compression

ZFS very efficiently compresses data on disk, increasing storage capacity without applications knowing any different.