ZFS Storage

I use ZFS as filesystem for my backend storage, due to its seamless integration with Snapshots and the Proxmox Backup Server. It is decently performant compared to the likes of ext, but offers great resilience in RAID-like solutions, is flexible and easy to use. I recommend reading through the basic ZFS introduction to clarify terms and usecases before following this guide, as this only reflects a good usecase for me.

1. Create Pool
    1. `Zpool create -o ashift=12 mirror POOL /dev/sdc /dev/sdd mirror /dev/sde /dev/sdf`
    2. Hardening against reboot issues with /dev/sdX nomenclature
        1. `Zpool export POOL`
        2. `Zpool import -d /dev/disk/dev-by-id`
        3. `Zpool import -c /etc/zfs/zpool.cache`
        4. `Zpool status`
2. Set compression: `zfs set compression=lz4 POOL`
3. Create datasets (+tuning)
    1. `zfs create pool/dataset-name`
    2. `sudo zfs set recordsize=[size] data/media/series`
    3. No-snap für datenbanksets d. Set mountpoints for datasets e. ZFS Scrub
4. Set mountpoints for datasets
5. ZFS Scrub
6.

https://docs.oracle.com/cd/E24841_01/html/820-2313/gbbwa.html

https://www.usenix.org/system/files/login/articles/login_winter16_09_jude.pdf

https://shatteredsilicon.net/blog/2020/06/05/mysql-mariadb-innodb-on-zfs/

https://docs.oracle.com/cd/E24841_01/html/820-2313/gbbwa.html

https://unix.stackexchange.com/questions/288599/forcing-zpool-to-use-dev-disk-by-id-in-ubuntu-xenial https://docs.oracle.com/cd/E19253-01/819-5461/gaynr/index.html

---

Revision #5
Created 3 September 2022 20:53:46 by TJN
Updated 3 May 2025 12:33:09 by TJN