

Create ZFS Pool

Links

https://pve.proxmox.com/wiki/ZFS_on_Linux

Introduction

ZFS pool can be built with `wwn` or `partuuid`, to be able to run the pool even from an USB case if needed. For normal service, `disk ID` is sufficient.

Poolname should be customerid, to make deployment easier. [Poolname can be changed](#)

Create a ZFS pool with 2 HDD as mirror

Get WWN from HDD

```
ls -l /dev/disk/by-id/  
# look for the first target HDD, in our case ata-WDC_WD40EFRX-68WT0N0_WD-WCC4E4DF9NJD ->  
../../sdj  
# which means the first target HDD is also named sdj  
# now look for the WWN which belongs to sdj:  
# wwn-0x50014ee20c6324e6 -> ../../sdj  
# look for the second target HDD, in our case ata-WDC_WD40EFRX-68WT0N0_WD-WCC4E6FN513C ->  
../../sdk  
# which means the second target HDD is also named sdk  
# now look for the WWN which belongs to sdk:  
# wwn-0x50014ee20c629629 -> ../../sdk
```

Create ZFS Pool with WWN from above

```
zpool create -o ashift=12 -O compression=zstd poolname mirror /dev/disk/by-id/wwn-  
0x50014ee20c6324e6 /dev/disk/by-id/wwn-0x50014ee20c629629
```

Create ZFS Pool in Raidz1

Get WWN from HDD same as above for each HDD

Create ZFS Pool with WWN from above

```
zpool create -o ashift=12 -0 compression=lz4 poolname raidz1 /dev/disk/by-id/wwn-<1>  
/dev/disk/by-id/wwn-<2> /dev/disk/by-id/wwn-<3>
```

Create a ZFS pool on a single disk with no data redundancy

Find WWN of the HDD

```
ls -l /dev/disk/by-id/  
# look for the first target HDD, in our case ata-WDC_WD40EFRX-68WT0N0_WD-WCC4E4DF9NJD ->  
../../sdj  
# which means the first target HDD is also named sdj  
# now look for the WWN which belongs to sdj:  
# wwn-0x50014ee20c6324e6 -> ../../sdj
```

Create ZFS Pool with WWN from above

```
zpool create -o ashift=12 -0 compression=lz4 poolname /dev/disk/by-id/wwn-0x50014ee20c6324e6
```

Create a ZFS pool on a single disk with 5 partitions for data redundancy

Attention: Very slow performance

```
apt install parted  
parted /dev/disk/by-id/ata-WDC_WD40EFRX-68WT0N0_WD-WCC4E4DF9NJD mklabel gpt  
parted /dev/disk/by-id/ata-WDC_WD40EFRX-68WT0N0_WD-WCC4E4DF9NJD mkpart zfs 0% 20%  
parted /dev/disk/by-id/ata-WDC_WD40EFRX-68WT0N0_WD-WCC4E4DF9NJD mkpart zfs 20% 40%  
parted /dev/disk/by-id/ata-WDC_WD40EFRX-68WT0N0_WD-WCC4E4DF9NJD mkpart zfs 40% 60%
```

```
parted /dev/disk/by-id/ata-WDC_WD40EFRX-68WT0N0_WD-WCC4E4DF9NJD mkpart zfs 60% 80%
parted /dev/disk/by-id/ata-WDC_WD40EFRX-68WT0N0_WD-WCC4E4DF9NJD mkpart zfs 80% 100%
```

Find the letter of the disk

```
ls -l /dev/disk/by-id/ata-WDC_WD40EFRX-68WT0N0_WD-WCC4E4DF9NJD
```

Result

```
/dev/disk/by-id/ata-WDC_WD40EFRX-68WT0N0_WD-WCC4E4DF9NJD -> ../../sdh
```

Find all Partuuid of sdh

```
ls -l /dev/disk/by-partuuid/ | grep sdh
```

Result (just one displayed as example)

```
2c49c49f-4221-324e-afca-23bedbb06677 -> ../../sdh1 #2c49c49f-4221-324e-afca-23bedbb06677 is
the partuuid1
```

Create ZFS pool (use all 5 partuuid)

```
zpool create -o ashift=12 -0 compression=lz4 poolname raidz1 /dev/disk/by-partuuid/<partuuid1>
/dev/disk/by-partuuid/<partuuid2> ... /dev/disk/by-partuuid/<partuuid5>
```

Revision #11

Created 2024-04-18 22:19:01 UTC by Admin

Updated 2025-08-08 22:26:37 UTC by Admin