ZFS

- Replace bad disk on ZFS pool

Replace bad disk on ZFS pool

Check the status of the pool

```
root@pve:~# zpool status
  pool: Mirror1
 state: DEGRADED
status: One or more devices could not be used because the label is missing or
        invalid.  Sufficient replicas exist for the pool to continue
        functioning in a degraded state.
action: Replace the device using 'zpool replace'.
   see: https://openzfs.github.io/openzfs-docs/msg/ZFS-8000-4J
  scan: resilvered 3.16G in 00:14:53 with 0 errors on Sat Jun  1 04:46:50 2024
config:

        NAME                                    STATE     READ WRITE CKSUM
        Mirror1                                 DEGRADED     0     0     0
          mirror-0                              DEGRADED     0     0     0
            55ee9c9d-1a2e-412d-827e-1875474cdd90  ONLINE       0     0     0
            12500327251565222082                UNAVAIL      0     0     0  was /dev/disk/by-partuuid/52151b98-6ccc-
4bfd-8318-4e72ced09a5f
```

52151b98-6ccc-4bfd-8318-4e72ced09a5f

Put offline the bad drive

```
root@pve:~# zpool offline Mirror1 12500327251565222082
root@pve:~# zpool status
  pool: Mirror1
 state: DEGRADED
```

```
status: One or more devices has been taken offline by the administrator.
        Sufficient replicas exist for the pool to continue functioning in a
        degraded state.
action: Online the device using 'zpool online' or replace the device with
        'zpool replace'.
  scan: resilvered 2.50G in 00:11:31 with 0 errors on Tue Jun  4 18:33:06 2024
config:

        NAME                                    STATE     READ WRITE CKSUM
        Mirror1                                 DEGRADED     0     0     0
          mirror-0                              DEGRADED     0     0     0
            55ee9c9d-1a2e-412d-827e-1875474cdd90  ONLINE       0     0     0
            52151b98-6ccc-4bfd-8318-4e72ced09a5f  OFFLINE      0     0     0
```

Remove drive and insert new one, this case its device /dev/sdc

```
root@pve:~# lsblk | grep sd
sda                         8:0    0 111.8G  0 disk
├─sda1                      8:1    0  1007K  0 part
├─sda2                      8:2    0     1G  0 part /boot/efi
└─sda3                      8:3    0 110.8G  0 part
sdb                         8:16   0 111.8G  0 disk
└─sdb1                      8:17   0 111.8G  0 part /mnt/data
sdc                         8:32   0   1.8T  0 disk
sdd                         8:48   0   7.3T  0 disk
├─sdd1                      8:49   0   7.3T  0 part
└─sdd9                      8:57   0     8M  0 part
sde                         8:64   0   1.8T  0 disk
├─sde1                      8:65   0     2G  0 part
└─sde2                      8:66   0   1.8T  0 part
sdf                         8:80   0   7.3T  0 disk
├─sdf1                      8:81   0   7.3T  0 part
└─sdf9                      8:89   0     8M  0 part
```

Retrieve disk ID, our drive is number 9

```
root@pve:~# ls -1 /dev/disk/by-id/ | grep ata
ata-ADATA_SU650_2M1629QJKNU1
ata-ADATA_SU650_2M1629QJKNU1-part1
ata-ADATA_SU650_4M10277WTRC7
ata-ADATA_SU650_4M10277WTRC7-part1
ata-ADATA_SU650_4M10277WTRC7-part2
ata-ADATA_SU650_4M10277WTRC7-part3
ata-ST2000DM001-9YN164_Z1E1B9AK
ata-ST2000DM006-2DM164_Z4Z806BP
ata-ST2000DM006-2DM164_Z4Z806BP-part1
ata-ST2000DM006-2DM164_Z4Z806BP-part2
ata-ST8000DM004-2CX188_ZR14EZCZ
ata-ST8000DM004-2CX188_ZR14EZCZ-part1
ata-ST8000DM004-2CX188_ZR14EZCZ-part9
ata-ST8000DM004-2U9188_ZR12SA93
ata-ST8000DM004-2U9188_ZR12SA93-part1
ata-ST8000DM004-2U9188_ZR12SA93-part9
```

replace mirrir old device with new device by id and wait for resilver (could take long)

```
root@pve:~# zpool replace Mirror1 52151b98-6ccc-4bfd-8318-4e72ced09a5f ata-ST2000DM001-
9YN164_Z1E1B9AK
root@pve:~# zpool status Mirror1
  pool: Mirror1
 state: DEGRADED
status: One or more devices is currently being resilvered.  The pool will
        continue to function, possibly in a degraded state.
action: Wait for the resilver to complete.
  scan: resilver in progress since Tue Jun  4 18:47:00 2024
        246G / 246G scanned, 211M / 246G issued at 106M/s
        195M resilvered, 0.08% done, 00:39:46 to go
```

config:

```
        NAME                                    STATE     READ WRITE CKSUM
        Mirror1                                 DEGRADED     0     0     0
          mirror-0                              DEGRADED     0     0     0
            55ee9c9d-1a2e-412d-827e-1875474cdd90    ONLINE       0     0     0
            replacing-1                         DEGRADED     0     0     0
              52151b98-6ccc-4bfd-8318-4e72ced09a5f  OFFLINE      0     0     0
              ata-ST2000DM001-9YN164_Z1E1B9AK       ONLINE       0     0     0  (resilvering)
```