

Changing disk identifiers in ZFS zpool •

<https://plantroon.com/changing-disk-identifiers-in-zpool/>

With ZFS on Linux, it often happens that zpool is created using disk identifiers such as `/dev/sda`. While this is fine for most scenarios, the [recommended practice](#) is to use the more guaranteed disk identifiers such as the ones found in `/dev/disk/by-id`. This blog post describes 3 methods how to change the disk identifiers in such zpool **after** it has been created. All this without migrating data, adding disks or having physical access to the machine.

- [Export and import](#) - offline, any zpool
- [Offline and replace](#) - online, any zpool, buggy on Linux
- [Detach and attach](#) - online, mirror zpool only

Export and import

This is a well-known procedure which works if you can afford downtime. The zpool has to be unmounted for this to work.

Start by exporting the zpool (rpool0 in this example):

Import it right after with specifying the correct path where to look for the devices:

```
# zpool import -d /dev/disk/by-id/ rpool0
```

This might import the disks by their [world wide name](#) which is not a problem since these are guaranteed to be persistent for any given disk on Linux, no matter where you put it (different drive bay or a completely different machine).

Offline and replace

This is the official method for replacing disks as described in [Oracle docs](#). The command `zpool labelclear` refuses to do anything because it sees the disk as part of an active zpool. For some reason, none of the usual Linux tools like `wipefs`, `parted`, `gdisk`, `fdisk` could manage to properly clear ZFS metadata from the disk, so the only option is zeroing out the disk manually which takes a long time and unnecessarily wears out SSDs.

Detach and attach

This works for online renaming, for example, when your operating system's root is on the zpool. Sadly, this does not work for RaidZ pools.

Here is the zpool mirror used in the example:

```
pool: rpool0

state: ONLINE
  scan: scrub repaired 0 in 0h1m with 0 errors on Sun Apr 14 00:25:39 2019
config:

          NAME                                STATE        READ  WRITE CKSUM
          rpool0                              ONLINE         0      0      0
            mirror-0                          ONLINE         0      0      0
              sda                             ONLINE         0      0      0
              sdb                             ONLINE         0      0      0

errors: No known data errors
```

Start by removing the first disk in the mirror:

```
# zpool detach rpool0 sda
```

Wipe zpool information from the disk:

```
# zpool labelclear -f /dev/sda
```

If it fails, try:

Determine its name in /dev/disk/by-id:

```
# ls -al /dev/disk/by-id | grep sda
```

Attach it to the zpool mirror with the correct name:

```
# zpool attach rpool0 sdb /dev/disk/by-id/ata-disk-model-serial
```

Now watch `zpool status` and when the resilvering process completes, do the same for the other disk in the pool:

```
# zpool detach rpool0 sdb
# zpool labelclear -f /dev/sda
# wipefs -a /dev/sda
# ls -al /dev/disk/by-id | grep sdb
# zpool attach rpool0 sda /dev/disk/by-id/ata-disk-model-serial
```

Wait for the resilvering process to finish. Now you have a properly set up zpool mirror.