



Replacing a drive in MD array

20.10.23

Pegasi Knowledge

<https://ghost.pegasi.fi/wiki/>

Table of Contents

Replacing a drive in MD array	3
The checking part	3
Detach old drive	3
Activate spare drive and add the new disk	3

Replacing a drive in MD array

We assume we have 3 drives: sda, sdb and hot spare drive sdc. The first drive sda is failing due to temperature problems and we want to take it offline and put sdc online.

The checking part

Check the configuration and drives

```
mdadm -D /dev/md0
mdadm -D /dev/md1
```

Check the condition of drives and be sure that you have the right one to blame

```
smartctl --all /dev/sda
smartctl --all /dev/sdb
smartctl --all /dev/sdc
```

If the broken drive is still alive but adding latency and CPU load, you might want to see it for yourself before doing harsh decisions. So add an io monitor first

```
iostat -x 1
```

and do a stress test with something like this (remember to do it inside the mount you want it to be in)

```
dd if=/dev/zero of=/tmp/koe bs=1024k count=2000
```

Look at the monitor and see which drive is receiving the penalty. That would be the one which we will get rid of.

And remember to have [other drives bootable](#)

Detach old drive

Remove the faulty drive from array

```
mdadm /dev/md0 --fail /dev/sda1 --remove /dev/sda1
mdadm /dev/md1 --fail /dev/sda2 --remove /dev/sda2
```

Activate spare drive and add the new disk

This would be the part where you stuff like this

```
mdadm /dev/md0 --add /dev/sdc1
mdadm /dev/md0 --add /dev/sdc2
```

But since my spare /dev/sdc was already there waiting the md subsystem automatically started utilizing it without the need for intervention.

Look at /proc/mdstat and wait until /dev/sdc is in sync.

Just to be on the safe side you should do the following steps:

- Try booting from sdb with BIOS boot override from second disk
- Disconnect the first drive cable, boot and confirm that sda really is the one disconnected (check /etc/fstab /dev/disk/by-uuid)
- Connect the new drive as first drive, boot and do following

```
sfdisk -d /dev/sdb > sdb.txt
sfdisk --force /dev/sda < sdb.txt
mdadm /dev/md0 --add /dev/sda1
mdadm /dev/md0 --add /dev/sda2
```

Now look at the arrays with mdadm -D /dev/mdX and confirm that the new drive partitions are as spares. They will activate as soon as one of actives fail or is marked as failed.

Or if you have 2 drives raid /proc/mdstat shows you it is syncing /dev/sda and after it's completion all is done.

Oh, and remember again to [make the new drive bootable](#)

Voilà!