

Hardware

How do I repair the RAID1 of a server after replacing a defective hard drive ?

In the following example we have a RAID1 with the drives `/dev/sda` and `/dev/sdb`.

The hard drives are divided into three primary partitions:

```
/dev/sdX1    /boot  
/dev/sdX2    swap  
/dev/sdX3    /
```

So we have 3 RAID Arrays

```
/dev/md0    /boot  
/dev/md1    swap  
/dev/md2    /
```

In our example the second hard drive - `/dev/sdb` - is defective.

To install a replacement, we take out the faulty drive from the current RAID Array first :

```
mdadm /dev/md0 -r /dev/sdb1  
mdadm /dev/md1 -r /dev/sdb2  
mdadm /dev/md2 -r /dev/sdb3
```

Then we can change the faulty drive (possibly the system has to be switched off). After the replacement we need to set up the partition layout on the replacement drive.

This can be done via `sfdisk`:

```
sfdisk -d /dev/sda | sfdisk /dev/sdb
```

Hereby the partition is transferred 1:1 from `sda` to `sdb`.

There's also an alternative way possible:

Hardware

It is quite simple in our example, because we have no extended partitions:

```
dd if=/dev/sda of=/dev/sdb bs=512 count=1
```

Hereby we copy the MBR (Size = 512 Byte) to the new hard drive. If the hard drive scheme is a construct with extended partitions, they must be created additionally.

Let's have a look at the partition scheme with fdisk:

```
fdisk -ul /dev/sda
```

We have to note the default values of the extended partitions and hand them over to dd's parameters skip and seek:

```
dd if=/dev/sda of=/dev/sdb count=1 skip=STARTVALUE seek=STARTVALUE
```

After we have completely acquired the partition layout, we tell this the kernel:

```
blockdev --rereadpt /dev/sdb
```

Now we can begin with rebuilding the RAID:

```
mdadm /dev/md0 -a /dev/sdb1  
mdadm /dev/md1 -a /dev/sdb2  
mdadm /dev/md2 -a /dev/sdb3
```

The status can be read anytime from the file /proc/mdstat,

Hardware

e.g. with

```
watch -n1 cat /proc/mdstat
```

Now we create the Swap space on the new hard drive:

The rebuilding of the partitions takes a while, depending on the size of the partitions and RAID mode.

On completion of the sync process our RAID is completely ready again.

Unique solution ID: #1372

Author: EUserV Support

Last update: 2012-07-18 14:37