Rescue System

How do I restore my Raid 1 after replacing a hard disk?

If a hard disk in a Raid network has failed, this can be remedied by means of the command:

cat /proc/mdstat

An example output in which the hard disk sdb (second hard disk in the system) has failed looks as follows:

cat /proc/mdstat Personalities : [raid1]

md125 : active raid1 sda3[0] sdb3 [1](F) 1073740664 blocks super 1.2 [2/1] [U_]

md126 : active raid1 sda5[0] sdb5[1](F) 524276 blocks super 1.2 [2/1] [U_]

md127 : active raid1 sda4[0] sdb4[1](F) 33553336 blocks super 1.2 [2/1] [U_]

As soon as the defective hard disk has been replaced by our support, the Raid must be restored. To do this, the first step is to copy the partitions from the non-defective hard disk (sda in the example).

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

You can use the following command to check whether the partitions have been created on sdb:

fdisk -l Disk /dev/sdb: 1,8 TiB, 2000398934016 bytes, 3907029168 sectors Units: sectors of 1 * 512 = 512 bytes Sector size (logical/physical): 512 bytes / 4096 bytes I/O size (minimum/optimal): 4096 bytes / 4096 bytes Disklabel type: gpt Disk identifier: A0FA7428-8E95-4683-AF1E-D6E8B75D1619

Device Start End Sectors Size Type /dev/sdb1 2048 4095 2048 1M BIOS boot /dev/sdb2 4096 1003519 999424 488M EFI System /dev/sdb3 1003520 17004543 16001024 7,6G Linux RAID /dev/sdb4 17004544 19005439 2000896 977M Linux RAID /dev/sdb5 19005440 3907028991 3888023552 1,8T Linux RAID

Disk /dev/sda: 1,8 TiB, 2000398934016 bytes, 3907029168 sectors Units: sectors of 1 * 512 = 512 bytes Sector size (logical/physical): 512 bytes / 4096 bytes I/O size (minimum/optimal): 4096 bytes / 4096 bytes Disklabel type: gpt Disk identifier: 319B1206-243C-4DE8-9D01-C96448249BD9

Device Start End Sectors Size Type /dev/sda1 2048 4095 2048 1M BIOS boot /dev/sda2 4096 1003519 999424 488M EFI System /dev/sda3 1003520 17004543 16001024 7,6G Linux RAID /dev/sda4 17004544 19005439 2000896 977M Linux RAID /dev/sda5 19005440 3907028991 3888023552 1,8T Linux RAID

> Page 1 / 3 © 2025 myLoc managed IT <faq@myloc.de> | 01.07.2025 05:32 URL: https://faq.myloc.de/index.php?action=faq&cat=17&id=353&artlang=en

Rescue System

Once the partitions have been created, the individual partitions can be added back to the raid group. This is done for each partition individually. In our example we have 3 partitions in the raid group.

mdadm --manage /dev/md125 --add /dev/sdb3 mdadm --manage /dev/md126 --add /dev/sdb5 mdadm --manage /dev/md127 --add /dev/sdb4

Now the raid is synchronised. You can monitor / view the sync process using the following command:

cat /proc/mdstat

An example output of a synchronisation looks like this:

Personalities : [linear] [multipath] [raid0] [raid1] [raid5] [raid4] [raid6] [raid10] md125 : active raid1 sda3[0] sdb3[1] 24418688 blocks [2/1] [U_] [=>.....] recovery = 9.9% (2423168/24418688) finish=2.8min speed=127535K/sec

Restore Swap Partition

During an installation via our installation system, the swap partition is created as Raid-0, even if Raid-1 was selected during the installation. If a hard disk fails, the Raid-0 is no longer repairable. In this case, Raid-0 must be deleted and then rebuilt.

Delete Raid

In the following example, the hard disk "/dev/sda" has failed. The raid is named "md0". Please note that the names may differ from your system and use the correct name from your system here.

To remove a raid, the corresponding hard disk/partition must not be mounted. This is unmounted in the first step.

umount -l /dev/sda2

Now the Raid-0 is stopped.

```
mdadm --stop /dev/md0
mdadm: stopped /dev/md0 -> Das Raid wurde erfolgreich gestoppt.
```

Now the superblock of the hard disk must be removed. The superblock defines the hard disks or partitions as a raid device.

mdadm --zero-superblock /dev/sda2

Create Raid

Finally, a new Raid-0 is created.

mdadm --create /dev/md0 --level=0 --raid-devices=2 /dev/sda2 /dev/sdb2

Reinstall boot loader (Grub2)

The first step is to set up a chroot environment. For this, the corresponding partitions must be mounted. Here, md126 is the largest raid, which most likely contains the data. This is now mounted to /mnt:

mount /dev/md126 /mnt

Page 2 / 3 © 2025 myLoc managed IT <faq@myloc.de> | 01.07.2025 05:32 URL: https://faq.myloc.de/index.php?action=faq&cat=17&id=353&artlang=en

Rescue System

Now the Proc and Dev file systems are mounted:

mount -t proc proc /mnt/proc mount -obind /dev /mnt/dev

Now the chroot environment can be started:

chroot /mnt /bin/bash

Now we first regenerate the device map at Grub2:

grub-mkdevicemap -n

Grub2 can then be repaired using the following command:

grub-install /dev/sdb

Unique solution ID: #1533 Author: Bettina Brauer Last update: 2021-04-22 06:41