

RAID on GnuLinux - Mdadm Reference

1 Overview

There are a few options for software RAID on Gnu Linux. Among them is Btrfs and ZFS, however today I will focus on using mdadm. This is historically the oldest software raid, therefore should be better vetted, although its performance may be less of that of the first two mentioned - for simple servers, mdadm might be the most stable choice.

2 Details

I've worked with this in setting up some Core 2 Duo PCs, with 2 to 4 Sata HDDs. This will be a reference. Let's begin.

2.1 Creation of RAID:

Will not be covered here (yet). You must create the partition tables. Create the raid with mdadm. mkfs.ext4 on the raid partition. Add mdadm to grub config. Reinstall grub. Details may be provided later.

2.2 Details of RAID:

```
root@advacoONE:/dev# sudo mdadm -D /dev/md127
/dev/md127:
    Version : 1.2
  Creation Time : Fri Feb  1 01:00:25 2019
    Raid Level : raid1
    Array Size : 57638912 (54.97 GiB 59.02 GB)
  Used Dev Size : 57638912 (54.97 GiB 59.02 GB)
    Raid Devices : 3
  Total Devices : 2
 Persistence : Superblock is persistent

    Update Time : Fri Feb  1 02:40:44 2019
      State : clean, degraded
 Active Devices : 2
Working Devices : 2
```

```
Failed Devices : 0
Spare Devices : 0
```

```
Name : devuan:root
UUID : 83a8dc03:802a4129:26322116:c2cfe1d4
Events : 82
```

```
Number Major Minor RaidDevice State
-      0      0      0      removed
1       8     17      1      active sync  /dev/sdb1
2       8     33      2      active sync  /dev/sdc1
root@advacoONE:/dev#--
```

so you can see, one was removed (it auto removes, when unplugged)

2.3 Add Drive to RAID:

```
sudo mdadm --add /dev/md127 /dev/sda1
```

NOTE2: If you setup 2 hdds, in a raid, and want to add a third, if you just `--add`, it will show up as a spare... if you do `mdadm --grow /dev/md127 --raid-devices=3` then the third might be active sync (what we want) note that the `--grow`, seems to allow for parameter changes after you have already created the raid. you can also specify the exact same command, `raid-devices=3` in the setup of the raid (see install doc). Note that if you lose a drive, you can simply add it.

NOTE: don't worry about `mkfs.ext4` on the raid members, after initial setup. The RAID will manage that.

NOTE: if you have a new drive and need to copy the hdd partition tables: <https://unix.stackexchange.com/questions/12986/how-to-copy-the-partition-layout-of-a-whole-disk-using-standard-tools> or aka

(FOR MBR ONLY)

Save:

```
sfdisk -d /dev/sda > part_table
```

```

Restore:
sfdisk /dev/NEWHDD < part_table

(FOR GPT:)
# Save MBR disks
sgdisk --backup=/partitions-backup-$(basename $source).sgdisk $source
sgdisk --backup=/partitions-backup-$(basename $dest).sgdisk $dest

# Copy $source layout to $dest and regenerate GUIDs
sgdisk --replicate=$dest $source
sgdisk -G $dest

root@advacoONE:/dev# mdadm --add /dev/md127 /dev/sda1
mdadm: added /dev/sda1
root@advacoONE:/dev# sudo mdadm -D /dev/md127
/dev/md127:
    Version : 1.2
  Creation Time : Fri Feb  1 01:00:25 2019
    Raid Level : raid1
    Array Size : 57638912 (54.97 GiB 59.02 GB)
  Used Dev Size : 57638912 (54.97 GiB 59.02 GB)
    Raid Devices : 3
  Total Devices : 3
  Persistence : Superblock is persistent

    Update Time : Fri Feb  1 02:41:43 2019
      State : clean, degraded, recovering
  Active Devices : 2
Working Devices : 3
  Failed Devices : 0
   Spare Devices : 1

Rebuild Status : 0% complete

    Name : devuan:root
    UUID : 83a8dc03:802a4129:26322116:c2cfe1d4
    Events : 92

  Number   Major   Minor   RaidDevice State
     3         8         1         0     spare rebuilding   /dev/sda1

```

```
      1      8      17      1      active sync  /dev/sdb1
      2      8      33      2      active sync  /dev/sdc1
root@advacoONE:/dev#
```

Looks good.

Rebuild Status : 6% complete

```
Name : devuan:root
UUID : 83a8dc03:802a4129:26322116:c2cfe1d4
Events : 103
```

```
Number  Major  Minor  RaidDevice State
      3      8      1      0      spare rebuilding  /dev/sda1
      1      8      17      1      active sync  /dev/sdb1
      2      8      33      2      active sync  /dev/sdc1
```

as it progresses, you see the RAID rebuilding.

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

```
Every 1.0s: cat /proc/mdstat
advacoONE: Fri Feb  1 02:43:24 2019
```

```
Personalities : [raid1] [linear] [multipath] [raid0] [raid6] [raid5] [raid4] [raid10]
md127 : active raid1 sda1[3] sdb1[1] sdc1[2]
      57638912 blocks super 1.2 [3/2] [_UU]
      [==>.....] recovery = 11.2% (6471936/57638912) finish=13.2min spe
```

```
unused devices: <none>
```

WARNING: Reinstall grub on the new drive again as well afterwards.

2.4 Email Notifications on mdadm

Test emails on mdadm.. first configure email however you prefer (i currently use ssmtp, see this link: https://wiki.zoneminder.com/How_to_get_ssmtp_working_with_Zoneminder

—

then edit /etc/mdadm/mdadm.conf to have your email in mailaddr then

```
sudo mdadm --monitor --scan --test --oneshot
```

should send an email

<https://ubuntuforums.org/showthread.php?t=1185134> for more details
on email sending

3 References

The section about degraded disks

<https://help.ubuntu.com/lts/serverguide/advanced-installation.html.en>

General partition tips.

<https://github.com/zfsonlinux/zfs/wiki/Debian-Stretch-Root-on-ZFS>

SSMTP email setup:

https://wiki.zoneminder.com/How_to_get_ssmtp_working_with_Zoneminder
wiki.zoneminder.com/SMS_Notifications