

Software RAID 10

Jem Camba - 2018-12-10 - in Storage

This how-to will demonstrate a Software RAID configuration for a Packet [m1.xlarge](#) server, “[The Virtualizer](#)”, running Ubuntu 16.04.

Setup

SSH into your server as **root** and run **cat /proc/mdstat** to view the current status of the server’s software RAIDed disks.

```
cat /proc/mdstat
Personalities : [raid1] [linear] [multipath] [raid0] [raid6] [raid5] [raid4] [raid10]
md126 : active (auto-read-only) raid1 sdb2[1] sda2[0]
        1995776 blocks super 1.2 [2/2] [UU]

md127 : active raid1 sdb3[1] sda3[0]
        466718720 blocks super 1.2 [2/2] [UU]
        bitmap: 0/4 pages [0KB], 65536KB chunk

unused devices: <none>
```

To get a list of hard drives available to you for configuration run **lsblk**.

```
root@raid:~# lsblk
NAME        MAJ:MIN RM  SIZE RO TYPE  MOUNTPOINT
sda         8:0    0 447.1G 0 disk
├─sda1      8:1    0    2M 0 part
├─sda2      8:2    0   1.9G 0 part
├─┬md127   9:127  0   1.9G 0 raid1 [SWAP]
└─sda3      8:3    0 445.2G 0 part
```

```

└─md126  9:126  0 445.1G  0 raid1 /
sdb      8:16   0 447.1G  0 disk
├─sdb1   8:17   0   2M  0 part
├─sdb2   8:18   0  1.9G  0 part
└─┬─md127 9:127  0  1.9G  0 raid1 [SWAP]
   └─sdb3   8:19   0 445.2G  0 part
└─md126  9:126  0 445.1G  0 raid1 /
sdc      8:32   0 447.1G  0 disk
sdd      8:48   0 447.1G  0 disk
sde      8:64   0 447.1G  0 disk
sdf      8:80   0 447.1G  0 disk

```

As you can see, there are four disks available, (sdc, sdd, sde, sdf), which we can use in a RAID.

Step 2: Preparing the Disks

Now that we know what drives to use, we can prepare them by using **fdisk** or **parted**, to create a partition table and change the partition type for our drives.

Note: You'll need to repeat this for the four available drives (sdc, sdd, sde, sdf)

Bring up the disk utility with **fdisk /dev/sdc**

The command will bring up the command line interface. It works by typing the code of the option you'd like to select (left side) and confirming the selection by hitting the return/enter key.

```
fdisk /dev/sdc
```

```
Welcome to fdisk (util-linux 2.27.1).
```

```
Changes will remain in memory only, until you decide to write them.
```

```
Be careful before using the write command.
```

```
Device does not contain a recognized partition table.
```

```
Created a new DOS disklabel with disk identifier 0x6eac01f2.
```

```
Command (m for help): m
```

```
Help:
```

```
DOS (MBR)
```

- a toggle a bootable flag
- b edit nested BSD disklabel
- c toggle the dos compatibility flag

Generic

- d delete a partition
- F list free unpartitioned space
- l list known partition types
- n add a new partition
- p print the partition table
- t change a partition type
- v verify the partition table
- i print information about a partition

Misc

- m print this menu
- u change display/entry units
- x extra functionality (experts only)

Script

- I load disk layout from sfdisk script file
- O dump disk layout to sfdisk script file

Save & Exit

- w write table to disk and exit
- q quit without saving changes

Create a new label

- g create a new empty GPT partition table
- G create a new empty SGI (IRIX) partition table
- o create a new empty DOS partition table
- s create a new empty Sun partition table

Command (m for help):

Select **n**, "add a new partition" to create a new partition. We use the default recommendations here but feel free to configure as you see fit!

Command (m for help): n

Partition type

- p primary (0 primary, 0 extended, 4 free)
- e extended (container for logical partitions)

Select (default p): p

Partition number (1-4, default 1):

First sector (2048-937703087, default 2048):

Last sector, +sectors or +size{K,M,G,T,P} (2048-937703087, default 937703087):

Created a new partition 1 of type 'Linux' and of size 447.1 GiB.

Select **t**, "change the partition type" to change the drives type. Typing L will show you all the the available types. We're going to select Linux raid.

Command (m for help): t

Selected partition 1

Partition type (type L to list all types): L

0	Empty	24	NEC DOS	81	Minix / old Lin	bf	Solaris
1	FAT12	27	Hidden NTFS Win	82	Linux swap / So	c1	DRDOS/sec (FAT-
2	XENIX root	39	Plan 9	83	Linux	c4	DRDOS/sec (FAT-
3	XENIX usr	3c	PartitionMagic	84	OS/2 hidden or	c6	DRDOS/sec (FAT-
4	FAT16 <32M	40	Venix 80286	85	Linux extended	c7	Syrinx
5	Extended	41	PPC PReP Boot	86	NTFS volume set	da	Non-FS data
6	FAT16	42	SFS	87	NTFS volume set	db	CP/M / CTOS / .
7	HPFS/NTFS/exFAT	4d	QNX4.x	88	Linux plaintext	de	Dell Utility
8	AIX	4e	QNX4.x 2nd part	8e	Linux LVM	df	BootIt
9	AIX bootable	4f	QNX4.x 3rd part	93	Amoeba	e1	DOS access
a	OS/2 Boot Manag	50	OnTrack DM	94	Amoeba BBT	e3	DOS R/O
b	W95 FAT32	51	OnTrack DM6 Aux	9f	BSD/OS	e4	SpeedStor
c	W95 FAT32 (LBA)	52	CP/M	a0	IBM Thinkpad	hi	ea Rufus alignment
e	W95 FAT16 (LBA)	53	OnTrack DM6 Aux	a5	FreeBSD	eb	BeOS fs
f	W95 Ext'd (LBA)	54	OnTrackDM6	a6	OpenBSD	ee	GPT
10	OPUS	55	EZ-Drive	a7	NeXTSTEP	ef	EFI (FAT-12/16/
11	Hidden FAT12	56	Golden Bow	a8	Darwin UFS	f0	Linux/PA-RISC b
12	Compaq diagnost	5c	Priam Edisk	a9	NetBSD	f1	SpeedStor
14	Hidden FAT16 <3	61	SpeedStor	ab	Darwin boot	f4	SpeedStor
16	Hidden FAT16	63	GNU HURD or Sys	af	HFS / HFS+	f2	DOS secondary
17	Hidden HPFS/NTF	64	Novell Netware	b7	BSDI fs	fb	VMware VMFS
18	AST SmartSleep	65	Novell Netware	b8	BSDI swap	fc	VMware VMKCORE
1b	Hidden W95 FAT3	70	DiskSecure Mult	bb	Boot Wizard hid	fd	Linux raid auto
1c	Hidden W95 FAT3	75	PC/IX	bc	Acronis FAT32 L	fe	LANstep
1e	Hidden W95 FAT1	80	Old Minix	be	Solaris boot	ff	BBT

Don't forget to write/save the changes you made!

Partition type (type L to list all types): fd

Changed type of partition 'Linux' to 'Linux raid autodetect'.

Command (m for help): w

The partition table has been altered.

Calling ioctl() to re-read partition table.

Syncing disks.

Do the same changes with the remaining 3 drives.

Step 3: Configuring the Software RAID 10

Confirm the partitions! Run lsblk to see your newly created partitions.

```
sdc      8:32  0 447.1G  0 disk
└─sdc1   8:33  0 447.1G  0 part
sdd      8:48  0 447.1G  0 disk
└─sdd1   8:49  0 447.1G  0 part
sde      8:64  0 447.1G  0 disk
└─sde1   8:65  0 447.1G  0 part
sdf      8:80  0 447.1G  0 disk
└─sdf1   8:81  0 447.1G  0 part
```

The command we'll be using to create RAID arrays is **mdadm**:

```
mdadm --create --verbose dev/name-of-your-drive --level=the-raid-configuration --raid-
devices=the-number-of-drives dev/drive-1-name dev/drive-2-name
```

To configure RAID 10 run this command:

```
mdadm --create --verbose /dev/md0 --level=10 --raid-devices=4 /dev/sdc1 /dev/sdd1
/dev/sde1 /dev/sdf1
```

You can check your RAID configurations by viewing the **mdstat** file with **cat /proc/mdstat**

```
cat /proc/mdstat
Personalities : [raid1] [linear] [multipath] [raid0] [raid6] [raid5] [raid4] [raid10]
md0 : active raid10 sdf1[3] sde1[2] sdd1[1] sdc1[0]
      937438208 blocks super 1.2 512K chunks 2 near-copies [4/4] [UUUU]
      [>.....] resync = 0.1% (1232768/937438208) finish=75.9min
      speed=205461K/sec
      bitmap: 7/7 pages [28KB], 65536KB chunk
```

Now that we have a Raid 10 with our 4 drives, it's time to make a filesystem, and mount it.

Step 4: Creating and Mount the filesystem

Create a filesystem for the new **md0** Raid

```
mkfs.ext4 -F /dev/md0
```

Next, create a mount point for it:

```
mkdir /mnt/raid
```

Mount the filesystem:

```
mount -t ext4 /dev/md0 /mnt/raid
```

And you can check it with the following:

```
root@raid:~# df -h -x devtmpfs -x tmpfs
Filesystem      Size  Used Avail Use% Mounted on
/dev/md126      438G  856M  415G   1% /
/dev/md0        880G   72M  836G   1% /mnt/raid
```

Finally, we have to add the new filesystem mount options to the `/etc/fstab` file for automatic mounting at boot.

Here we will use the UUID rather than the `/dev/md0` name.

```
blkid | grep md0
/dev/md0: UUID="d920df8a-d89e-4de6-8ce3-24507b17bc29" TYPE="ext4"
```

And we add it on the `/etc/fstab` file:

```
echo 'UUID="d920df8a-d89e-4de6-8ce3-24507b17bc29" /mnt/raid ext4 defaults 0 0' |
tee -a /etc/fstab
```

If you want to test the **fstab**, you can simply unmount the raid with **umount /mnt/raid** and them use **mount -a** which will mount everything from the **fstab** file.

Tags

RAID10