

RAID - mdadm

RAID IN LINUX BASICS

RAID 0 – STRIPING

- RAID 0 is designed for speed. It splits or "stripes" data across two or more drives, with each block of data going to a different drive. There is no duplication or parity involved. You need at least two drives, and 100% of the combined storage is usable. For example, with two 1TB drives, you get 2TB of usable space.
- Offers the best performance among all RAID types due to simultaneous read/write operations. It utilizes full disk capacity without overhead.
- There is no redundancy. If even one drive fails, all data is lost.

```
$ mdadm --create --verbose /dev/md0 --level=0 --raid-devices=2 /dev/sd{ab}1
```



RAID 1 – MIRRORING

- RAID 1 focuses on reliability. It mirrors the same data onto two drives, so if one fails, the other still holds an intact copy. It requires a minimum of two drives, but you only get 50% of the total space as usable storage. With two 1TB drives, for instance, you get just 1TB of usable space.
- Excellent fault tolerance. The system can continue to operate even if one drive fails.
- Storage efficiency is low—half of your total storage is used for redundancy.


```
$ mdadm --create --verbose /dev/md0 --level=1 --raid-devices=2 /dev/sd{ab}
```



RAID 3 / RAID 4 – STRIPING + DEDICATED PARITY

- RAID 3 and RAID 4 both use a dedicated parity disk. Data is striped across the other drives. They require at least three drives, and the usable capacity is (N-1) x smallest drive size. For instance, with 4x1TB, 3TB is usable. RAID 3 works at the byte level, RAID 4 at the block level.
- Offers fault tolerance and improved performance compared to RAID 1.
- The dedicated parity disk can become a bottleneck during writes, making it less efficient for heavy I/O workloads.

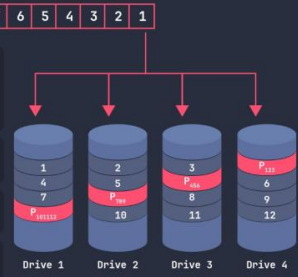
```
$ mdadm --create --verbose /dev/md0 --level=4 --raid-devices=4 /dev/sd{abcd}1
```




RAID 5 – STRIPING WITH DISTRIBUTED PARITY

- RAID 5 offers a balance between performance and redundancy. Data and parity are distributed across four drives. One drive's worth of space is used for parity, so the usable storage is (N-1) x smallest drive size. For example, with 4x1TB drives, you get 3TB usable.
- Good performance with fault tolerance. A single drive can fail without data loss.
- Rebuilding a failed drive can be time-consuming and slightly risky during the rebuild process.

```
$ mdadm --create --verbose /dev/md0 --level=5 --raid-devices=4 /dev/sd{abcd}1
```



RAID 10 – (NESTED RAID 1+0)



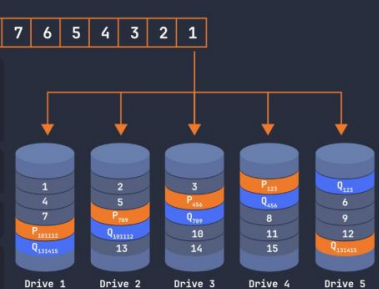
- RAID 10 combines mirroring and striping to deliver high performance and reliability. It mirrors pairs of drives and then stripes across those pairs. This setup requires at least four drives, and only 50% of total storage is usable. For example, 4x1TB drives give you 2TB of usable space.
- You get the speed benefits of striping and the redundancy of mirroring. It's one of the fastest and safest options.
- It's expensive in terms of drive count and offers only 50% storage efficiency.

```
$ mdadm --create --verbose /dev/md0 --level=10 --raid-devices=4 /dev/sd{abcd}1
```

RAID 6 – STRIPING WITH DOUBLE DISTRIBUTED PARITY


- RAID 6 takes RAID 5 further by adding a second parity block, allowing for two simultaneous drive failures. It requires at least four drives, and the usable storage is (N-2) x smallest drive size. For example, with 5x1TB drives, you get 3TB usable.
- High fault tolerance. Two drives can fail at once without losing data.
- Slightly slower write performance due to dual parity, and even less usable space compared to RAID 5.

```
$ mdadm --create --verbose /dev/md0 --level=6 --raid-devices=5 /dev/sd{abcde}1
```



JBOD – Just a Bunch Of Disks (not RAID!)

- JBOD isn't technically RAID but is supported by most RAID controllers. It allows you to combine multiple disks of different sizes into a single logical volume. There's no minimum number of drives, and 100% of the combined capacity is usable, minus any filesystem overhead.
- Simple and cost-effective for expanding storage using mismatched drives.
- No redundancy or performance gains. If one disk fails, you lose only the data on that disk, but the overall array may become unusable depending on how the volume is spanned.



RAID ⇒ R/A/I/D
(Redundant / Array of / Independent / Disk)

- Use identical drives (same size, speed, and model) to avoid imbalances or wasted space.
- Avoid "green" or low-power drives, as they may conflict with RAID due to power-down features.
- RAID is not a backup — always maintain separate, tested backups.
- Enable mdadm monitoring to get email alerts on disk or array issues.

```
# Create Raid 0  
$ mdadm --create --verbose /dev/md0 --level=0 --raid-devices=2 /dev/sd{ab}1  
# Format an Array  
$ mkfs.ext4 /dev/md0  
# Create a mount point and mount an array  
$ mkdir /mnt/raid0  
$ mount /dev/md0 /mnt/raid
```

