
RAID

Article URL

[RAID](#)

Author

SecurityHome.eu

Published: 12 September 2006

Last updated on 31 May 2012.

Â

This technology makes it possible to recover the data if one of the hard drives crashes.

There are different levels of RAID
RAID0

Disk Striping
RAID1

Disk Mirroring
RAID2

Disk Striping without error-correction code (ECC)
RAID3

Disk Striping with ECC as parity
RAID4

Disk Striping big parts; parity on 1 disk
RAID5

Disk Striping with parity on more disks
RAID6

Disk Striping with 2x parity on more disks
RAID0+1

Combination of RAID0 and RAID1
RAID1+0

Combination of RAID1 and RAID0
RAID5+0

Combination of RAID5 and RAID0
RAID5+1

Combination of RAID5 and RAID1
X-RAID2

is loosely equivalent to RAID5

RAID0 - Disk Stripping

Here all the data is distributed in blocks of 64K over all the disk in the stripe set.

- + Fast access.
- If one disk fails all the data is lost.

RAID1 - Disk Mirroring

With disk mirroring there is a exact copy on another disk. Doesn't have to be an even number of disks.

- + If 1 disk fails the data is still there on the other disk.
- You can use half of the total diskspace.
- With uneven disk, if the middle one fails the data is still lost.
(The middle one has parts of both mirrors).

RAID1 - Disk Duplicating

This is the same as disk Mirroring but has always an even number of disks.

- + If 1 disk fails the data is still there on the other disk.
- You can use half of the total diskspace.

Raid 2 Redundancy through Hamming code

Disks are synchronized and striped in very small stripes, often in single bytes/words. Hamming codes error correction is calculated across corresponding bits on disks, and is stored on multiple parity disks

- + If 1 disk fails, the data can still be restored.
- + Fast access.

RAID3 - Parallel transfer disks with parity

RAID level 3 stripes data across multiple drives, with an additional drive dedicated to parity, for error correction/recovery.

- + If 1 disk fails, the data can still be restored.

-
- + Fast access.
 - On 1 drive all space is "useless".

RAID 4 Block level parity

Identical to RAID 3 but does block-level striping instead of byte-level striping. The size of each stripe is large (Block). Now a single file can be stored in a block, each disk operate independently and so many different I/O request can be satisfied in parallel. But data transfer speed for each I/O request will be less. The error detection is done by block level parity and is stored on a separate disk.

- + If 1 disk fails, the data can still be restored.
- + Fast access.
- On 1 drive all space is "useless".

RAID5 - Disk Striping with parity on more disks

Here a stripe-set must contain at least 3 disks. All the data is distributed over the disks. On each disk there is space reserved for parity. The parity contains information to restore the data on another disk.

Example: Disk 2 and 3 can restore the data on disk 1, disk 1 and 3 can restore the data on disk 2, ...

- + Fast access.
- + If 1 disk fails no problem.

RAID 6 Striped set with dual parity

Provides fault tolerance from two drive failures; array continues to operate with up to two failed drives. This makes larger RAID groups more practical, especially for high availability systems. This becomes increasingly important because large-capacity drives lengthen the time needed to recover from the failure of a single drive. Single parity RAID levels are vulnerable to data loss until the failed drive is rebuilt: the larger the drive, the longer the rebuild will take. Dual parity gives time to rebuild the array without the data being at risk if one drive, but no more, fails before the rebuild is complete.

RAID 0+1: striped sets in a mirrored set

(minimum four disks; even number of disks) provides fault tolerance and improved performance but increases complexity. The key difference from RAID 1+0 is that RAID 0+1 creates a second striped set to mirror a primary striped set. The array continues to operate with one or more drives failed in the same mirror set, but if drives fail on both sides of the mirror the data on the RAID system is lost.

RAID 1+0: mirrored sets in a striped set

(minimum four disks; even number of disks) provides fault tolerance and improved performance but increases complexity. The key difference from RAID 0+1 is that RAID 1+0 creates a striped set from a series of mirrored drives. In a failed disk situation RAID 1+0 performs better because all the remaining disks continue to be used. The array can sustain multiple drive losses so long as no mirror loses both its drives.

RAID 5+0: stripe across distributed parity RAID systems

RAID 5+1: mirror striped set with distributed parity

(some manufacturers label this as RAID 53).

X-RAID2

X-RAID2 uses stripe like RAID5, but does not deliver a fixed-size logical unit.
So, you can keep adding drives and it will automatically expand.
And the drives don't have to be of the same size.
Here is a cool presentation

.

Some of these can be applied true software but for others you need hardware.
On hardware raids you mostly need drive of the same size (or even the same type/manufacture performance)
For software RAID it is not that important (But still recommended).

Some Operating systems allow you do combine partitions of drives in raid.
For example you can have 2 disks, with a part in mirroring (Raid 0) and another part as striped (raid 1).