

RAID Levels

RAID is an acronym first defined by David A. Patterson, Garth A. Gibson, and Randy Katz at the University of California, Berkeley in 1987 to describe a redundant array of inexpensive disks,[1] a technology that allowed computer users to achieve high levels of storage reliability from low-cost and less reliable PC-class disk-drive components, via the technique of arranging the devices into arrays for redundancy. Marketers representing industry RAID manufacturers later reinvented the term to describe a redundant array of independent disks as a means of dissociating a "low cost" expectation from RAID technology. In this article we will give you an idea about the different RAID levels. Feel free to comment the article on the bottom of it. Be aware that you need to register in order to leave your comment.

RAID level 0 is optimal for performance but suboptimal for fault tolerance. A RAID 0 array consists of one or more logical volumes cut across two or more physical volumes. Theoretically, this will improve logical disk I/O rates and decrease fault tolerance by a proportion equal to the number of physical volumes. For example, if the array consists of four disks, then it will be possible to read and write all them concurrently; a given amount of data can be transferred to the logical volume in a quarter of the time it would take if the logical volume were on only one disk. But if any of the four disks is damaged, the logical volume will be affected, so it is four times more likely to fail than if it were on only one disk.

RAID level 1 is optimal for fault tolerance. There may be performance gains, but that is not why you use it. Where RAID 1 is definitely suboptimal is cost. A RAID 1 array consists of one or more logical volumes mirrored across two or more disks: whenever data is written to the logical volume, copies of it will be written concurrently to two or more physical volumes. If any one physical volume is lost, the logical volume will survive because all data on the lost physical volume is available on another physical volume. There may be a performance improvement for read operations if it is possible to read different data from the mirrors concurrently; this will depend on the capabilities of the LVM. The cost problem is simply that you will require double the disk capacity—more than double, if you want a higher degree of mirroring. In the four-disk example, the logical volume will be equivalent in size to only two of the physical disks, but you can lose any one disk, and possibly two disks, before the logical volume is damaged.

RAID level 5 is a compromise between the performance of RAID 0 and the fault tolerance of RAID 1. The logical volume is cut across multiple physical volumes (so concurrent read and writes are possible), and a checksumming algorithm writes out enough information to allow reconstruction of the data on any one physical volume, if it gets damaged. Thus you do not get all the performance gain of RAID 0, because of the checksumming overhead; in particular, write operations can be slow because each write operation needs to calculate the checksum before the write can take place. You do not get all the fault tolerance of RAID 1, because you can survive the loss of only one disk. In the four-disk example, the logical volume will be equivalent in size to three physical volumes, and if any one of the four is lost, the logical volume will survive.

RAID 0+1 is optimal for both fault tolerance and performance: you mirror your striped disks. In the four-disk example, your system administrators would create one logical volume striped across two of the physical disks and mirror this to the other two disks. This should result in double the performance and double the safety (and double the price) of mapping one logical volume directly onto one physical volume.