

RAID Levels



“Redundant Arrays of Independent Disks, or RAID, has revolutionized the way on-line data is stored in computers. Spanning the entire spectrum from personal computers to super computers, RAID technology offers significant improvements in reliability, availability and serviceability”. - Joe Molina - RAB Chairman

The Acronyms

The “RAID” acronym first appeared in 1988 in the earliest of the Berkeley Papers written by Patterson, Gibson & Katz of the University of California at Berkeley. The RAID Advisory Board has since substituted “Independent” for “Inexpensive”.

A series of papers written by the original three authors and others defined and categorized several data protection and mapping models for disk arrays. Some of the models described in these papers, such as mirroring, were known at the time, others were new.

The word levels used by the authors to differentiate the models from each other may suggest that a higher numbered RAID model is uniformly superior to a lower numbered one. This is not the case.

RAID Level 0 (Disk Striping)

RAID Level 0 is a performance oriented striped data mapping technique. Uniformly sized blocks of storage are assigned in regular sequence to all of an array's disks.

RAID Level 0 provides high I/O performance at low inherent cost. (No additional disks are required). The reliability of RAID Level 0, however is less than that of its member disks due to its lack of redundancy.

Despite the name, RAID Level 0 is not actually RAID, unless it is combined with other technologies to provide data and functional redundancy, regeneration and rebuilding.

RAID Level 1

RAID Level 1, also called mirroring, has been used longer than any other form of RAID. It remains popular because of its simplicity and high level of reliability and availability. Mirrored arrays consist of two or more disks. Each disk in a mirrored array holds an identical image of user data. A RAID Level 1 array may use parallel access for high transfer rate when reading. More commonly, RAID Level 1 array members operate independently and improve performance for read-intensive applications, but at relatively high inherent cost.

RAID Level 2

RAID Level 2 is one of two inherently parallel mapping and protection techniques defined in the Berkeley paper. It has not been widely deployed in industry largely because it requires special disk features. Since disk production volumes determine cost, it is more economical to use standard disks for RAID systems.



RAID Level 3

RAID Level 3 adds redundant information in the form of parity to a parallel access striped array, permitting regeneration and rebuilding in the event of a disk failure. One stripe of parity protects corresponding strip's of data on the remaining disks. RAID Level 3 provides for high transfer rate and high availability, at an inherently lower cost than mirroring. Its transaction performance is poor, however, because all RAID Level 3 array member disks operate in lockstep.

RAID Level 4

Like RAID Level 3, RAID Level 4 uses parity concentrated on a single disk to protect data. Unlike RAID Level 3, however, a RAID Level 4 array's member disks are independently accessible. Its performance is therefore more suited to transaction I/O than large file transfers. RAID Level 4 is seldom implemented without accompanying technology, such as write-back cache, because the dedicated parity disk represents an inherent write bottleneck.

RAID Level 5

By distributing parity across some or all of an array's member disks, RAID Level 5 reduces (but does not eliminate) the write bottleneck inherent in RAID Level 4. As with RAID Level 4, the result is asymmetrical performance, with reads substantially outperforming writes. To reduce or eliminate this intrinsic asymmetry, RAID level 5 is often augmented with techniques such as caching and parallel multiprocessors.