



SNS COLLEGE OF TECHNOLOGY, COIMBATORE-35

(AN AUTONOMOUS INSTITUTION)



DEPARTMENT OF COMPUTER SCIENCE AND ENGINEERING

19CST202-DATABASE MANAGEMENT SYSTEM

UNIT-V

PHYSICAL STORAGE AND MONGODB

Topic: RAID

RAID:

RAID (redundant array of independent disks) is a way of storing the same data in different places on multiple hard disks or solid-state drives (SSDs) to protect data in the case of a drive failure. There are different RAID levels, however, and not all have the goal of providing redundancy.

RAID works

RAID works by placing data on multiple disks and allowing input/output (I/O) operations to overlap in a balanced way, improving performance. Because using multiple disks increases the [mean time between failures](#), storing data redundantly also increases fault tolerance.

RAID arrays appear to the operating system (OS) as a single logical drive.

RAID employs the techniques of disk mirroring or disk striping. Mirroring will copy identical data onto more than one drive. Striping [partitions](#) help spread data over multiple disk drives. Each drive's storage space is divided into units ranging from a [sector](#) of 512 bytes up to several megabytes. The stripes of all the disks are interleaved and addressed in order. Disk mirroring and disk striping can also be combined in a RAID array.

In a single-user system where large records are stored, the stripes are typically set up to be small (512 bytes, for example) so that a single record spans all the disks and can be accessed quickly by reading all the disks at the same time.


In a multiuser system, better performance requires a stripe wide enough to hold the typical or maximum size record, enabling overlapped disk I/O across drives.

RAID levels


RAID devices use different versions, called levels. The original paper that coined the term and developed the RAID setup concept defined six levels of RAID -- 0 through 5. This numbered system enabled those in IT to differentiate RAID versions. The number of levels has since expanded and has been broken into three categories: standard, nested and nonstandard RAID levels.

Standard RAID levels


RAID 0. This configuration has striping but no redundancy of data. It offers the best performance, but it does not provide fault tolerance.




RAID 1. Also known as *disk mirroring*, this configuration consists of at least two drives that duplicate the storage of data. There is no striping. Read performance is improved, since either disk can be read at the same time. Write performance is the same as for single disk storage.




RAID 2. This configuration uses striping across disks, with some disks storing error checking and correcting (ECC) information. RAID 2 also uses a dedicated Hamming code parity, a linear form of ECC. RAID 2 has no advantage over RAID 3 and is no longer used.



RAID 3. This technique uses striping and dedicates one drive to storing parity information. The embedded ECC information is used to detect errors. Data recovery is accomplished by calculating the exclusive information recorded on the other drives. Because an I/O operation addresses all the drives at the same time, RAID 3 cannot overlap I/O. For this reason, RAID 3 is best for single-user systems with long record applications.



RAID 4. This level uses large stripes, which means a user can read records from any single drive. Overlapped I/O can then be used for read operations. Because all write operations are required to update the parity drive, no I/O overlapping is possible.



RAID 5. This level is based on parity block-level striping. The parity information is striped across each drive, enabling the array to function, even if one drive were to fail. The array's architecture enables read and write operations to span multiple drives. This results in performance better than that of a single drive, but not as high as a RAID 0 array. RAID 5 requires at least three disks, but it is often recommended to use at least five disks for performance reasons.

RAID 5 arrays are generally considered to be a poor choice for use on write-intensive systems because of the performance impact associated with writing parity data. When a disk fails, it can take a long time to rebuild a RAID 5 array.

Diagram illustrating RAID 5 configuration with four disks (Disk 0, Disk 1, Disk 2, Disk 3). The data is distributed across the disks in a striped manner, with parity information stored on the disks. The data blocks are labeled as follows:

- Disk 0: Block A1, Block B1, Block C1, Block Dp
- Disk 1: Block A2, Block B2, Block C-Parity, Block D1
- Disk 2: Block A3, Block B-Parity, Block C2, Block D2
- Disk 3: Block A-Parity, Block B3, Block C3, Block D3

RAID 6. This technique is similar to RAID 5, but it includes a second parity scheme distributed across the drives in the array. The use of additional parity enables the array to continue functioning, even if two disks fail simultaneously. However, this extra protection comes at a cost. RAID 6 arrays often have slower write performance than RAID 5 arrays.

Diagram illustrating a RAID 5 configuration with 5 disks (DISK 0 to DISK 4). Each disk contains 6 blocks, organized into 6 data groups (A-E). The blocks are distributed such that each disk contains one block from each data group, ensuring data redundancy across the array.

Block Group	DISK 0	DISK 1	DISK 2	DISK 3	DISK 4
Block A	Block A1	Block A2	Block A3	Block Ap	Block Aq
Block B	Block B1	Block B2	Block Bp	Block Bq	Block B3
Block C	Block C1	Block Cp	Block Cq	Block C2	Block C3
Block D	Block Dp	Block Dq	Block D1	Block D2	Block D3
Block E	Block E1	Block E2	Block E3	Block E2	Block Ep