



SNS COLLEGE OF ENGINEERING

Kurumbapalayam(Po), Coimbatore - 641 107


Accredited by NAAC-UGC with 'A' Grade

Approved by AICTE, Recognized by UGC & Affiliated to Anna University, Chennai

Department of Artificial Intelligence and Data Science

**Course Name -23ITB204-Modern Database Management Systems
II Year / III Semester**

Topic- RAID from storage technique



REDUNDANT ARRAY OF INDEPENDENT DISKS

- way of storing the **same data in different places** on **multiple hard disks** or solid-state drives to protect data in the **case of a drive failure**
- connect multiple secondary storage devices for **increased performance, data redundancy or both**
- gives you the **ability to survive one or more drive failure** depending upon the RAID level used
- consists of an array of disks in which **multiple disks are connected** to achieve different goals



MULTIPLE PHYSICAL HARD-DRIVES

SINGLE LOGICAL HARD-DRIVE



- Redundancy Array of the Independent Disk
- technology which is used to connect multiple secondary storage devices for increased performance, data redundancy or both.
- gives the ability to survive one or more drive failure depending upon the RAID level used.
- It consists of an array of disks in which multiple disks are connected to achieve different goals
- RAID 0, RAID 1, RAID 2, RAID 3, RAID 4, RAID 5, RAID 6



- It contains a set of physical disk drives.
- In this technology, the operating system views these separate disks as a single logical disk.
- In this technology, data is distributed across the physical drives of the array.
- Redundancy disk capacity is used to store parity information.
- In case of disk failure, the parity information can be helped to recover the data.



WHY REDUNDANCY?



- although taking up **extra space**, adds to **disk reliability**
- in case of disk failure, if the **same data is also backed up onto another disk**, we can **retrieve the data** and go on with the operation
- if the data is spread across just multiple disks **without the RAID technique**, the **loss of a single disk can affect the entire data**.



MIRRORING



- approach to **introduce redundancy** is to duplicate every disk . This is called mirroring
 - A logical disk then consists of two physical disks, and every write is carried out on both disks.
- If one of the disks fails, the data can be read from the other.
- Data will be lost only if the second disk fails before the first failed disk is repaired



IMPROVEMENT IN PERFORMANCE VIA PARALLELISM



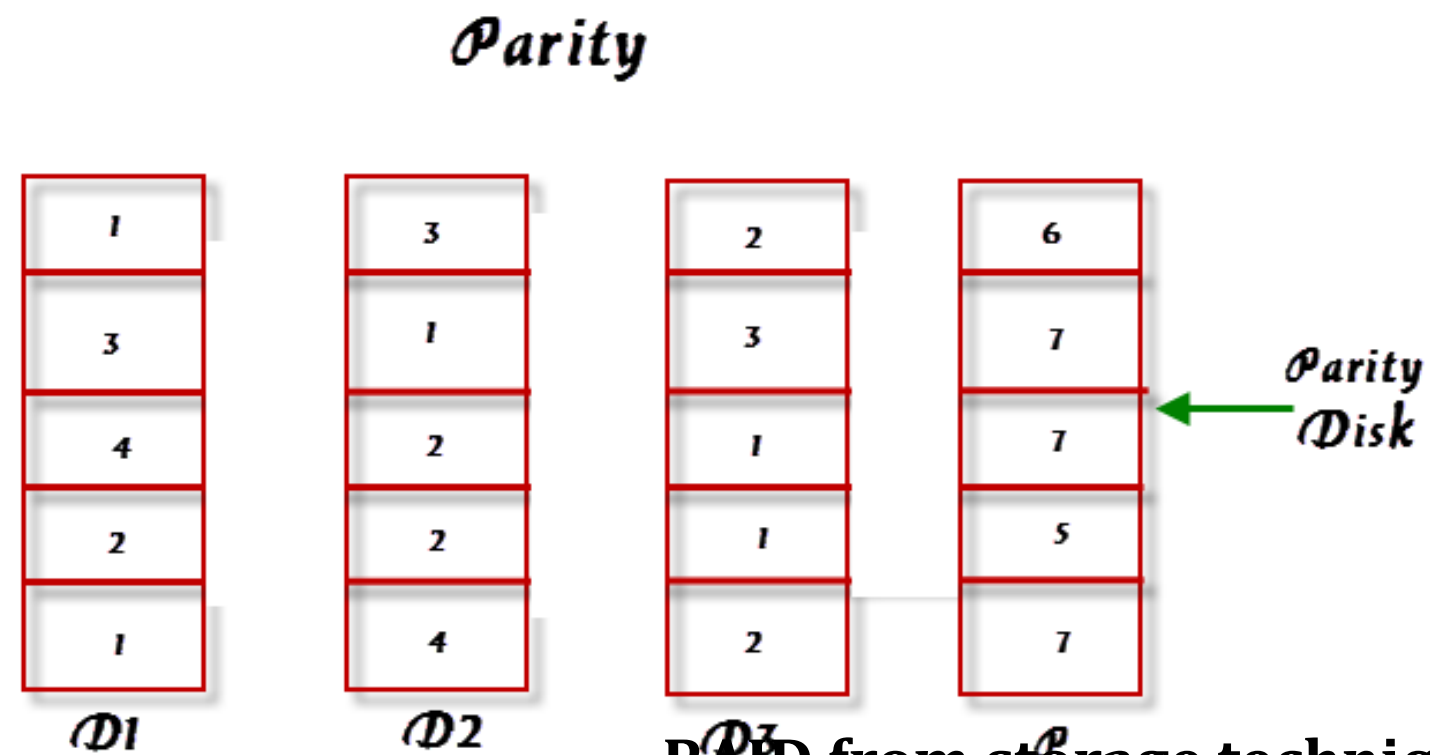
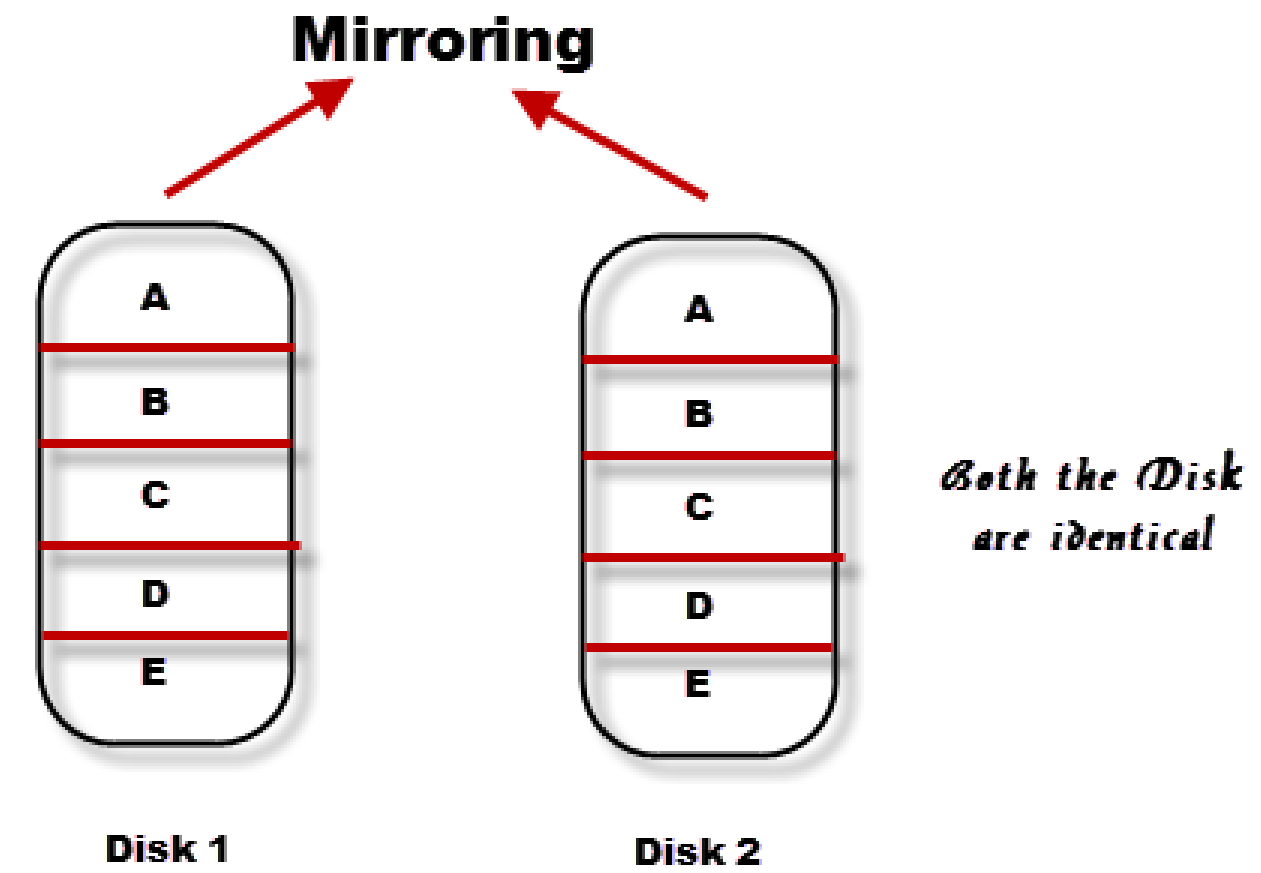

- with **Disk Mirroring** - rate at which read requests can be handled is doubled, since read requests can be sent to either disk
- we can improve the transfer rate as well (or instead) by **striping** data across multiple disks
- data striping consists of splitting the bits of each byte across multiple disks; such striping is called bit level striping.
- For e.g.,
 - if we have an array of eight disks, we write bit i of each byte to disk i
 - array of eight disks can be treated as a single disk - eight times the normal size - eight times the transfer rate



IMPROVEMENT IN PERFORMANCE VIA PARALLELISM



- Block-level striping stripes **blocks across multiple disks**
- treats the array of disks as a **single large disk**, and it gives **blocks logical numbers**
- array of **n disks**, block-level striping assigns **logical block i** of the disk array to disk **$(i \bmod n) + 1$**



RAID LEVELS

high reliability

Mirroring

expensive

Does not improve
high reliability

Striping


high
data-transfer rates

- 1. Provides redundancy
- 2. Lower cost
- 3. Disk striping with "parity" bits




RAID
Levels

RAID LEVELS – RAID 0

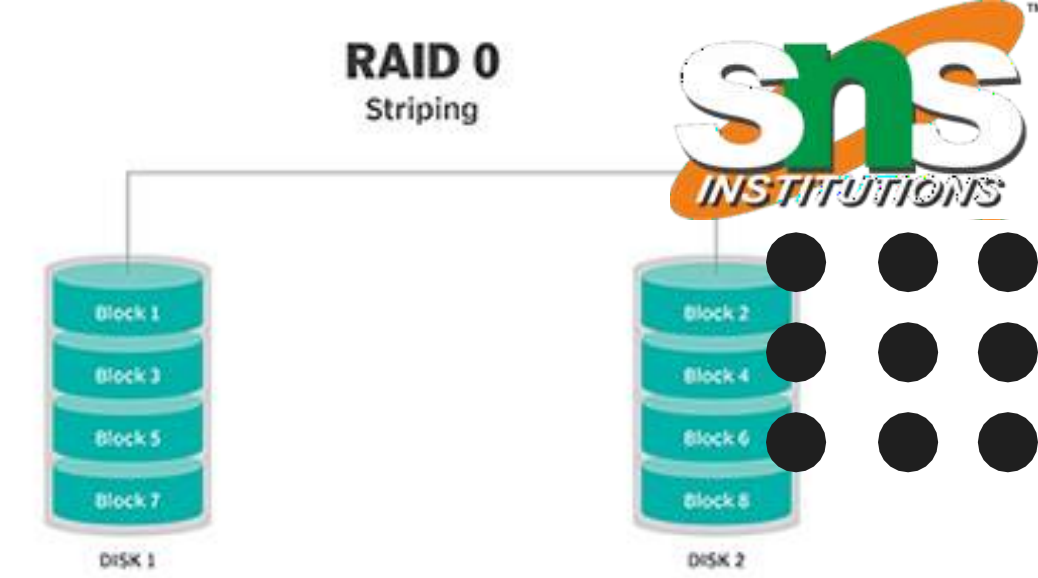


(a) RAID 0: nonredundant striping

- provides *data stripping* i.e., a *data can place across multiple disks*
- if **one disk fails then all data in the array is lost.**
- The *data is broken down into blocks* and the *blocks are distributed among disks*
- Each *disk receives a block of data to write/read in parallel*
- **Doesn't provide fault tolerance** but *increases the system performance*




RAID LEVELS – RAID 0



Disk 0	Disk 1	Disk 2
20	21	22
24	25	26
28	29	30
32	33	34

instead of placing just one block into a disk at a time, we can work with two or more blocks placed it into a disk before moving on to the next one

there is **no duplication of data**. Hence, **a block once lost cannot be recovered**.



Pros of RAID 0:

- In this level, throughput is increased because multiple data requests are processed in parallel.
- This level fully utilizes the disk space and provides high performance.
- It requires a minimum of 2 drives.

Cons of RAID 0:

- It doesn't contain any error detection mechanism.
- The RAID 0 is not a true RAID because it is not fault-tolerant.
- In this level, failure of either disk results in complete data loss in respect of the entire RAID.



THANKYOU