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

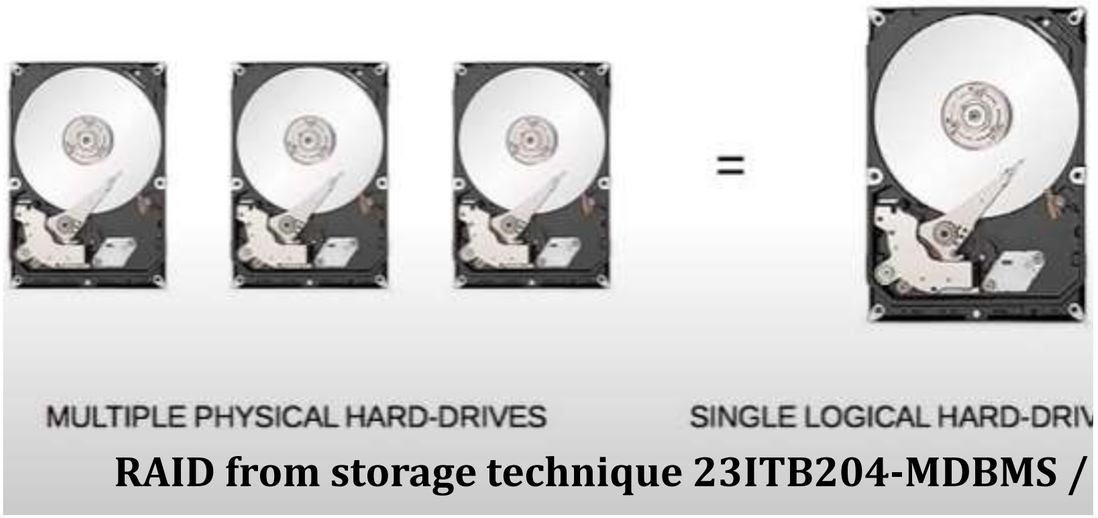

- Redundancy Array of the Independent Disk
- technology which is used to connect multiple secondary storage devices for increased performance, data redundancy or both.
- gives the ability to survive one or more drive failure depending upon the RAID level used.
- It consists of an array of disks in which 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 mod n) + 1

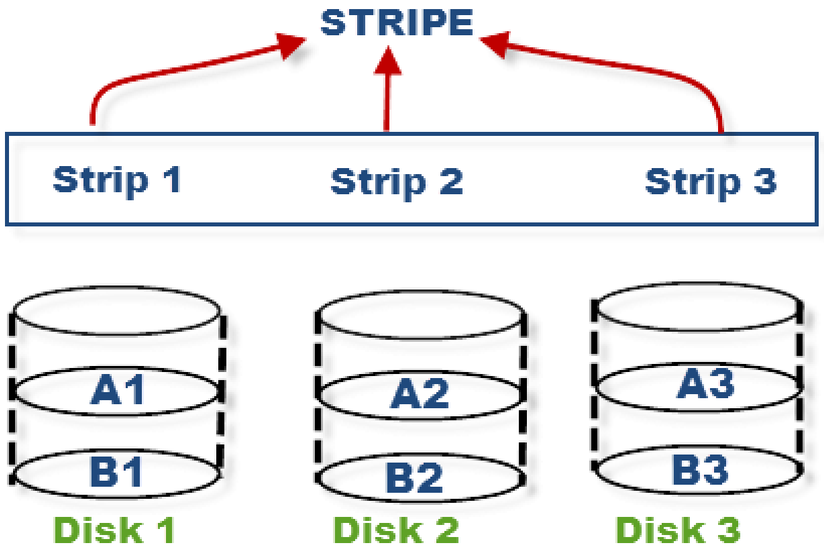
STRIPE
Strip 1
Strip 2
Strip 3
A1
B1
A2
B2
A3
B3
Disk 1
Disk 2
Disk 3

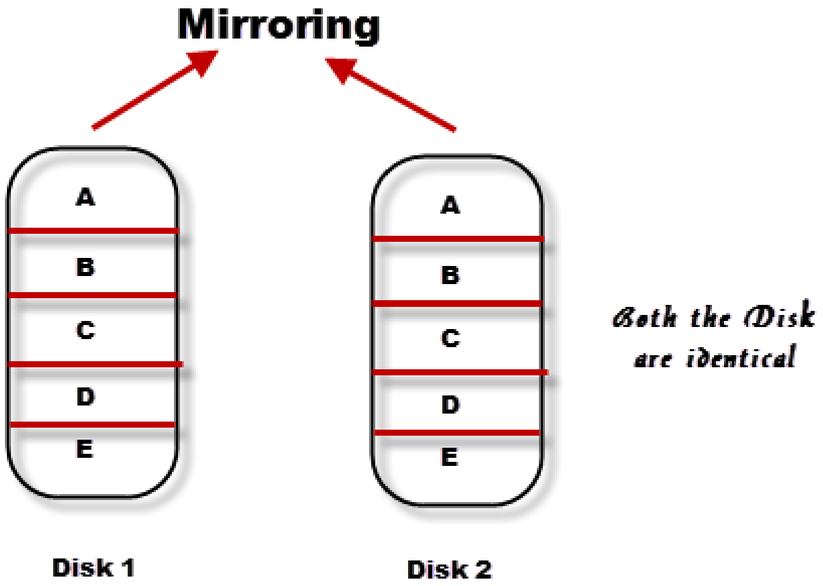
Mirroring
A
B
C
D
E
A
B
C
D
E
Both the Disk are identical
Disk 1
Disk 2

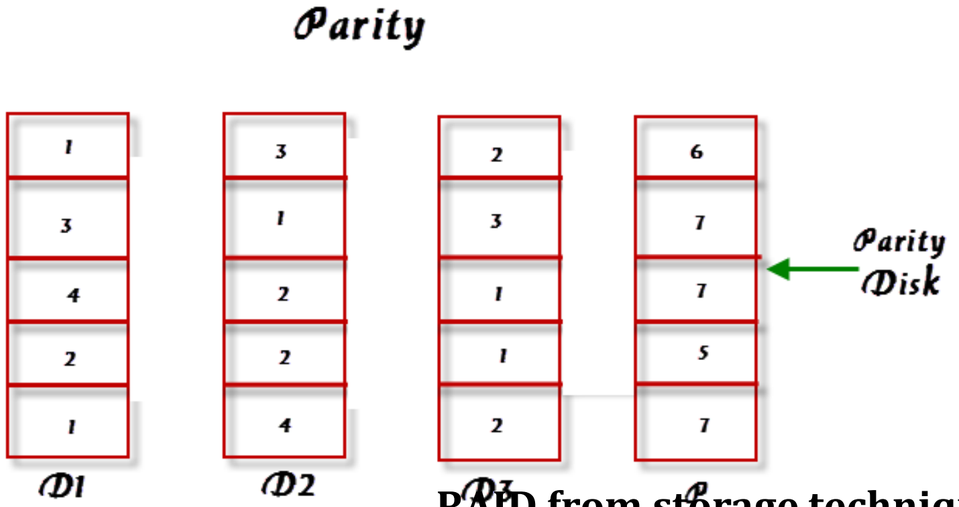
Parity
1
3
4
2
1
3
1
2
2
4
2
3
1
1
2
6
7
7
5
7
Parity Disk
D1
D2
D3
P

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

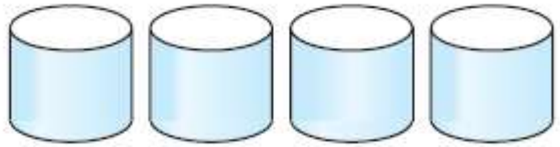

(a) RAID 0: nonredundant striping

- 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**

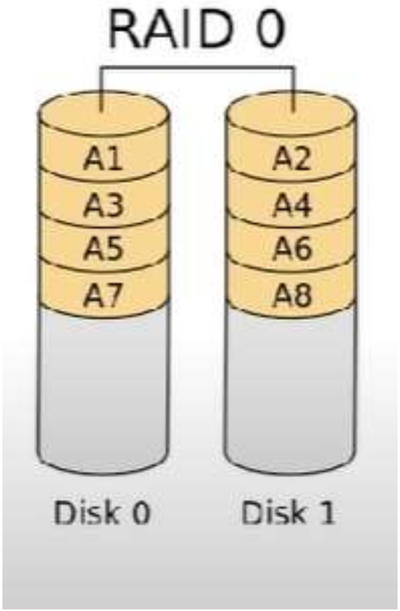

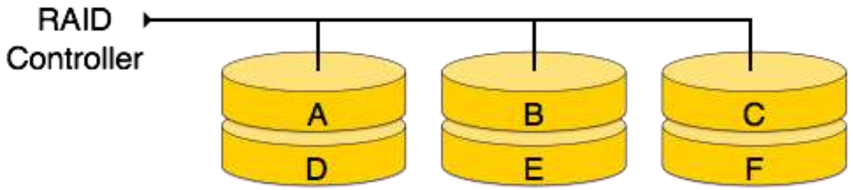

RAID LEVELS – RAID 0

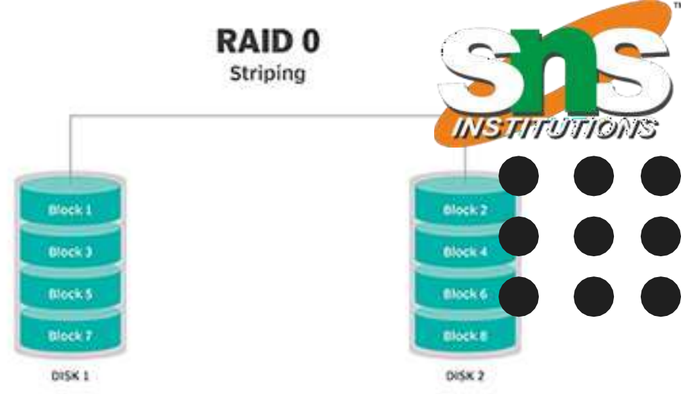

Disk 0	Disk 1	Disk 2
20	21	22
24	25	26
28	29	30
32	33	34

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 pı
- This level full utilizes the disk space and provides high performance.
- It requires minimum 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-tolerance.
- In this level, failure of either disk results in complete data loss in respe

RAID LEVELS – RAID 1

- This level is called ***mirroring*** of data
- **copies the data from drive 1 to drive 2**
- It provides **100% redundancy** in case of a failure

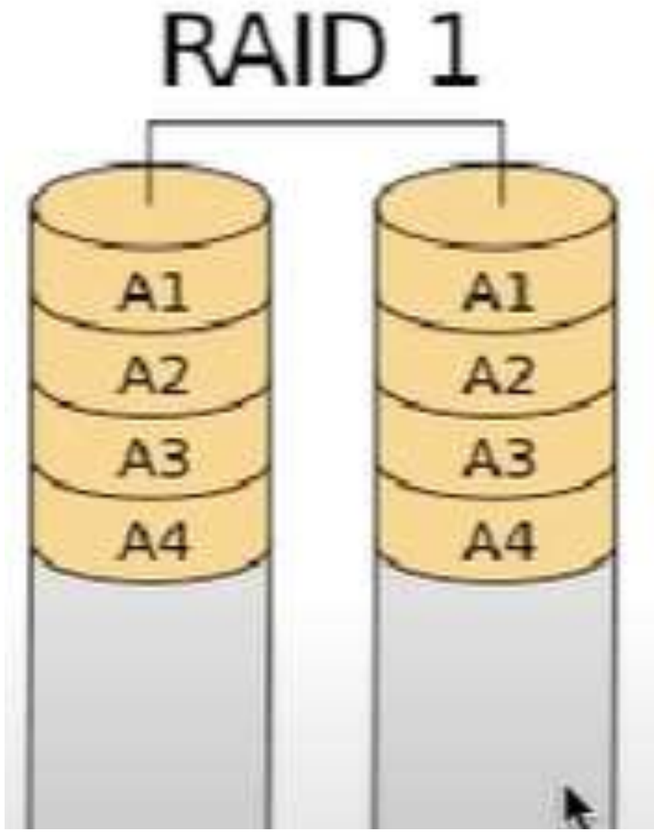

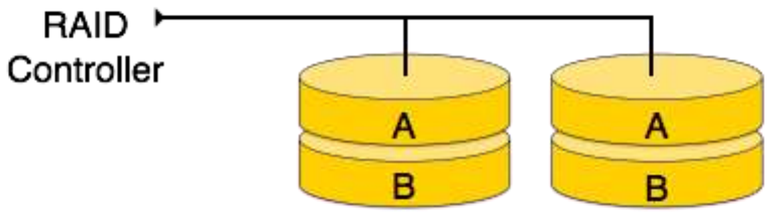

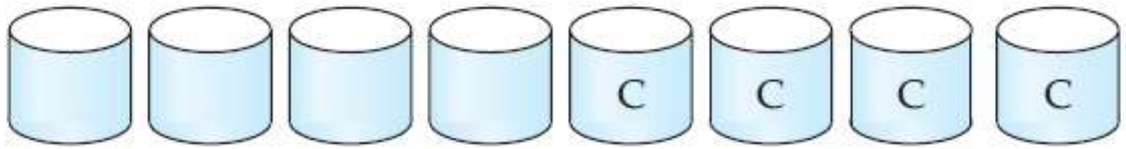

(b) RAID 1: mirrored disks

RAID LEVELS – RAID 1

Disk 0	Disk 1	Disk 2	Disk 3
A	A	B	B
C	C	D	D
E	E	F	F
G	G	H	H

Only half space of the drive is used to store the data.

The other half of drive is just a mirror to the already stored data.

RAID LEVELS – RAID 2

- RAID 2 **records Error Correction Code using Hamming distance** for its data, striped on different disks
- employs **parity bits**
- Each byte in a memory system may have a parity bit associated with it that records whether the numbers of bits in the byte that are set to 1 is even (parity = 0) or odd (parity = 1)
- If one of the bits in the byte gets damaged (either a 1 becomes a 0, or a 0 becomes a 1), the parity of the byte changes and thus will not match the stored parity

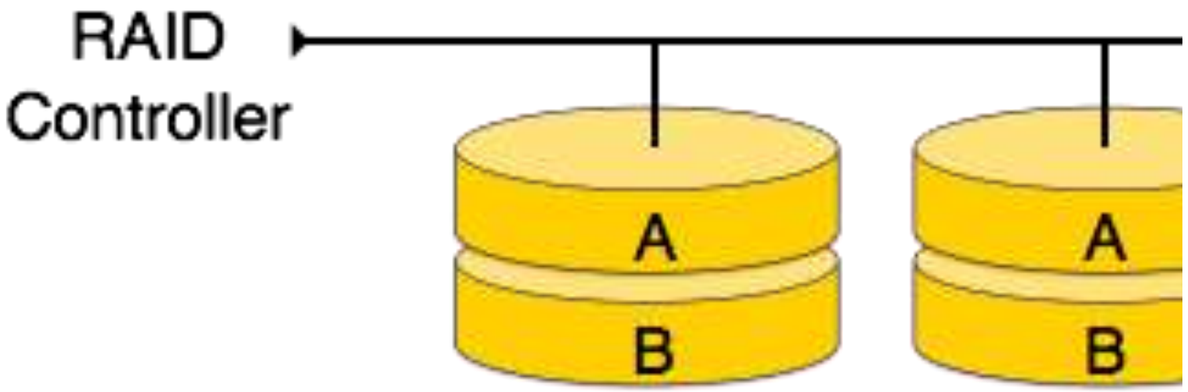

RAID LEVELS – RAID 2

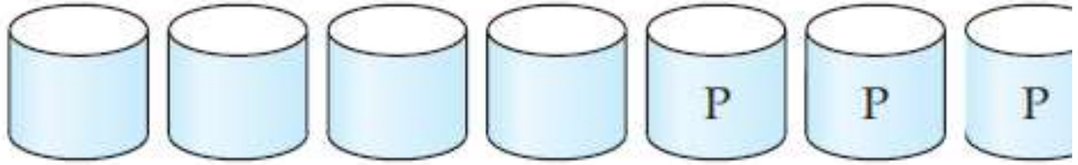

(c) RAID 2: memory-style error-correcting coc

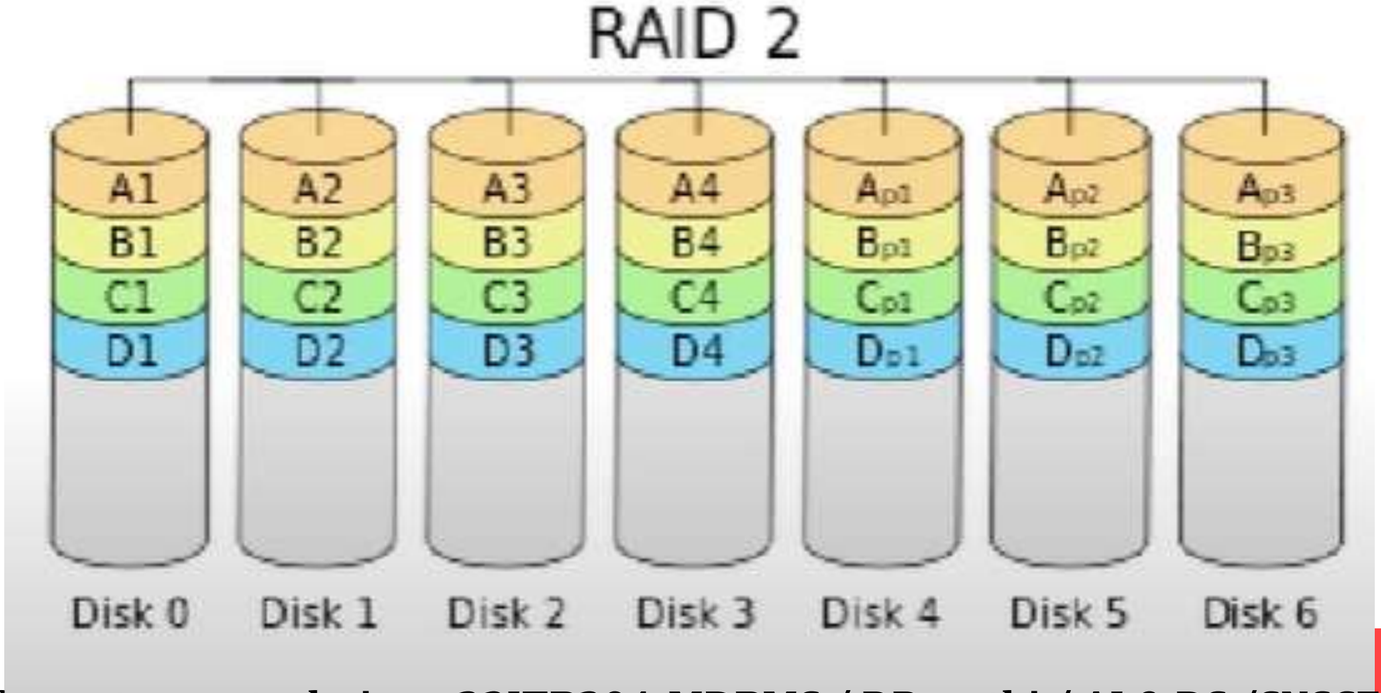

RAID LEVELS – RAID 3

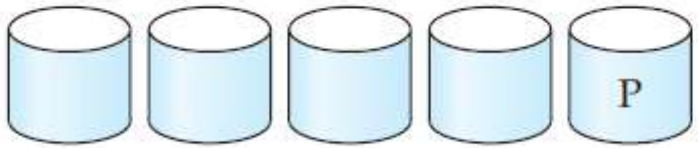

(d) RAID 3: bit-interleaved parity

- RAID 3 **stripes the data onto multiple disks**
- The parity bit generated for data word is stored on a different disk
- **In case of drive failure**, **the parity drive is accessed**, and data is reconstructed from the remaining devices.
- Once the failed drive is replaced, the missing data can be restored on the new drive.

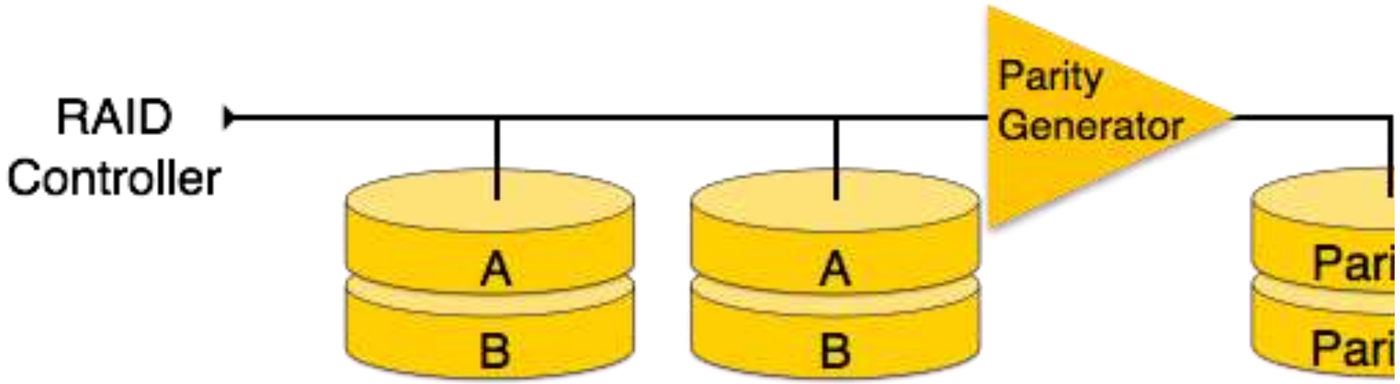

RAID LEVELS – RAID 3

Disk 0	Disk 1	Disk 2
A	B	C
D	E	F
G	H	I
J	K	L

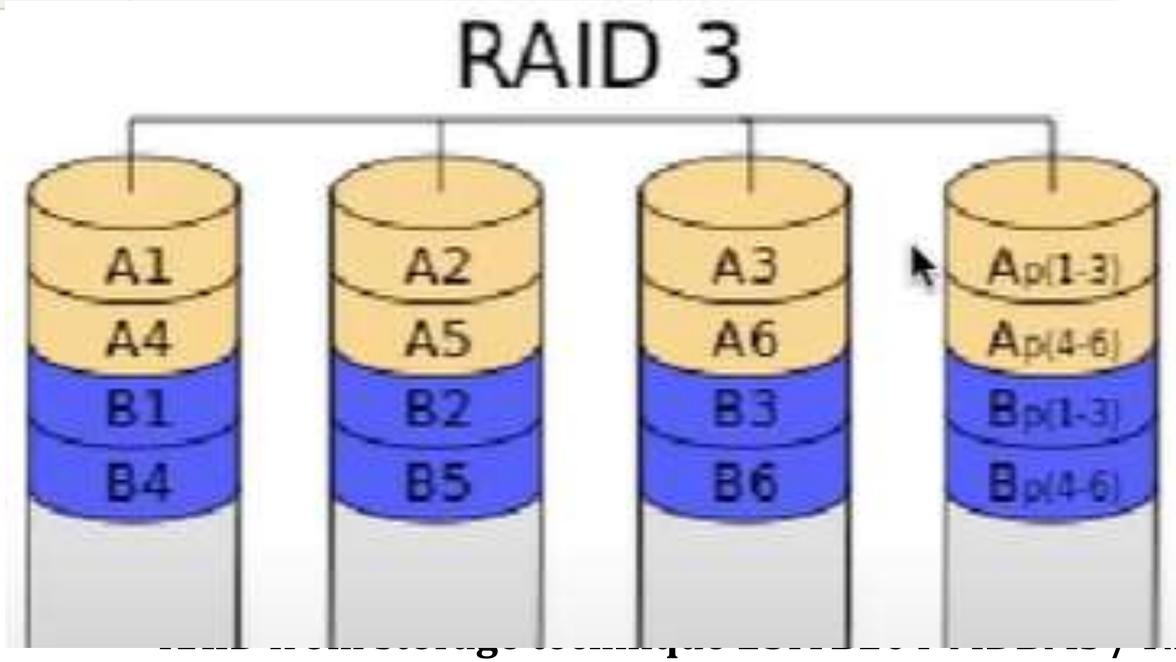

RAID LEVELS – RAID 4

- RAID 4 consists of **block-level stripping with a parity disk**
- This level **allows recovery of at most 1 disk failure due to the way parity works.**
- In this level, if more than one disk fails, then there is no way to recover the data

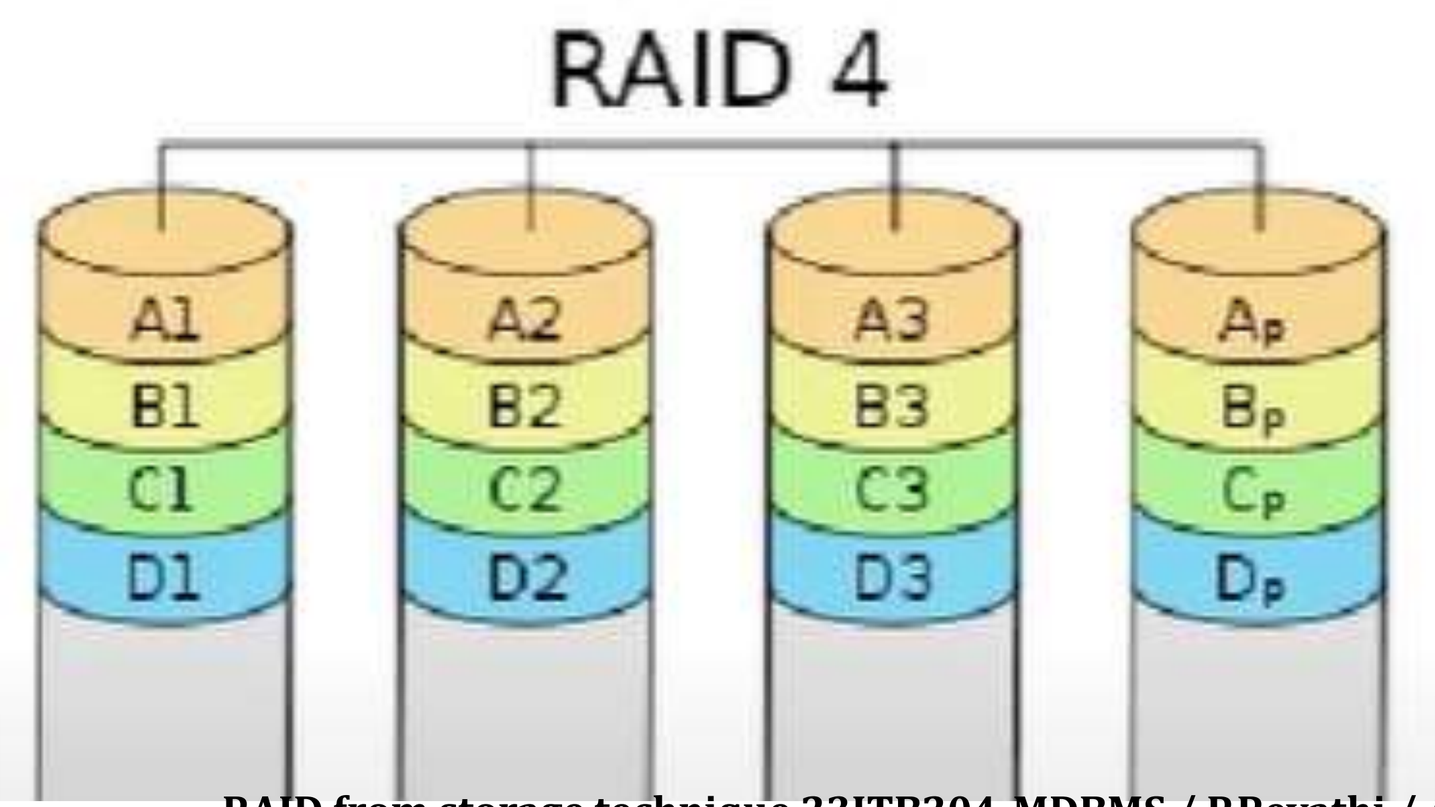

RAID LEVELS – RAID 5

- RAID 5 is a slight modification of the RAID 4 system.
- The only difference is that in RAID 5, the **parity rotates among the drives**
- It **consists of block-level striping with DISTRIBUTED parity**

Disk 0	Disk 1	Disk 2	Disk 3	Disk 4
0	1	2	3	P0
5	6	7	P1	4
10	11	P2	8	9
15	P3	12	13	14
P4	16	17	18	19

RAID LEVELS – RAID 5

- RAID 5 **writes whole data blocks onto different disks**, but the parity bits generated for data block stripe are distributed among all the data disks **rather than storing them on a different dedicated disk.**

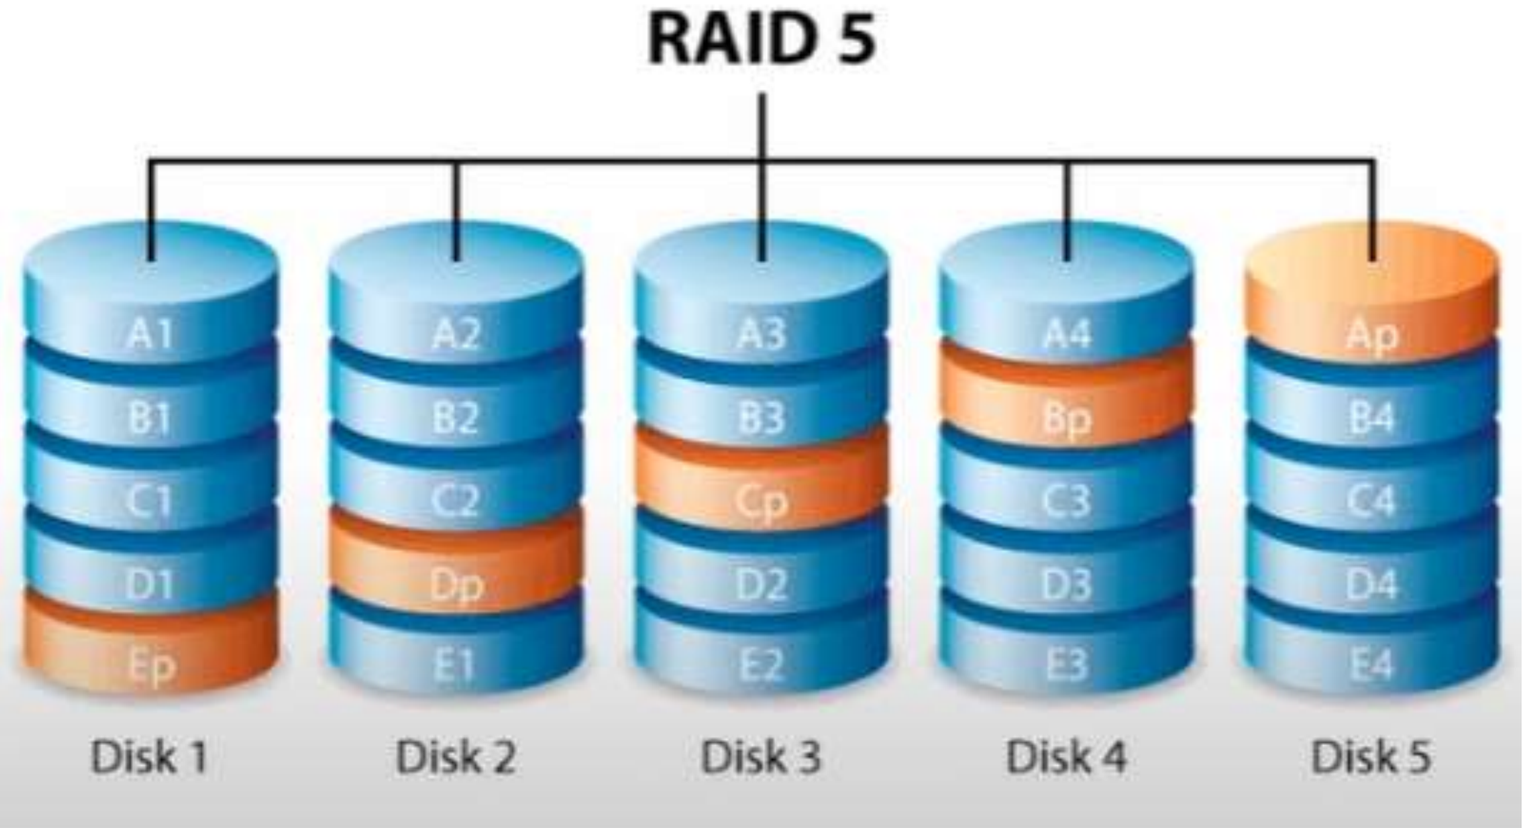

RAID LEVELS – RAID 6

- RAID 6 is an extension of level 5.
- In this level, **two independent parities** are generated and stored in distributed fashion among multiple disks.
- **Two parities provide additional fault tolerance**.
- This level **requires at least four disk drives to implement RAID**

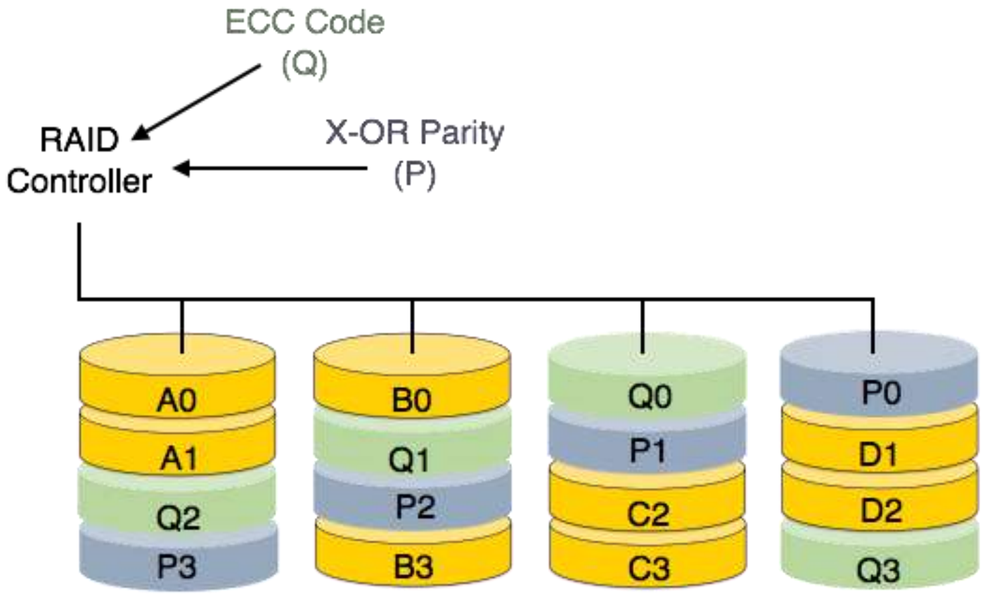

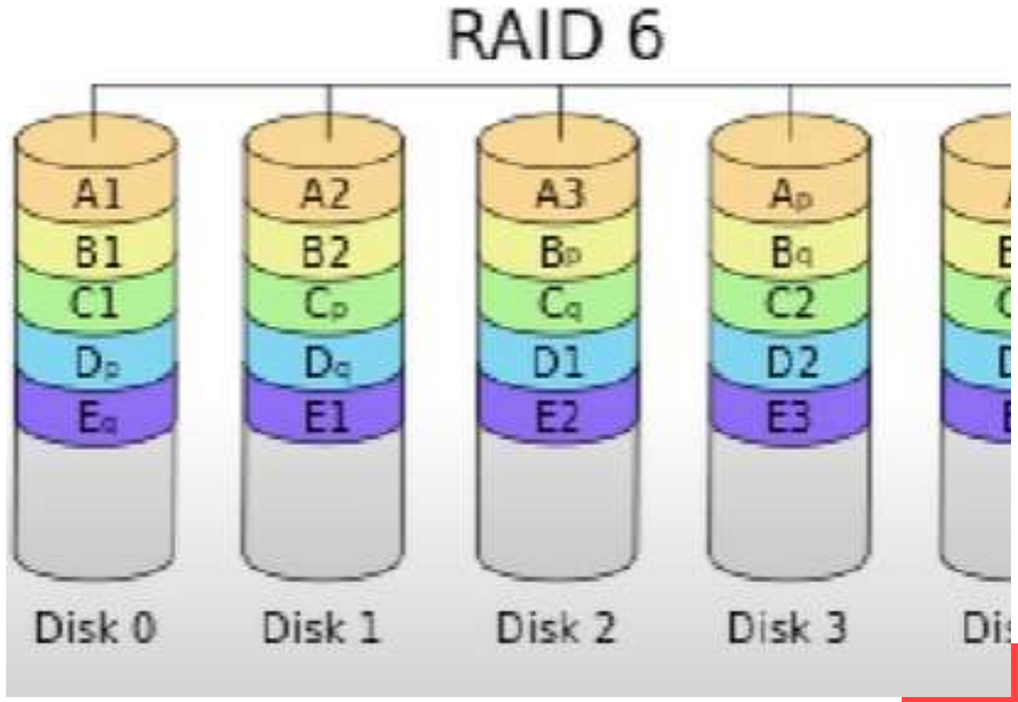

THANKYO
U