DATABASE MANAGEMENT SYSTEMS

Associate Professor Dr. Raed Ibraheem Hamed

University of Human Development, College of Science and Technology
Computer Science Department

2015 – 2016

Points to Cover

- Storing Data in a DBMS
- Primary Storage
- Secondary Storage
- Tertiary Storage
- Redundant Array of Independent Disks
  - RAID 0
  - RAID 1
  - RAID 2
  - RAID 3
  - RAID 4
  - RAID 5
  - RAID 6

Storing Data in a DBMS

Databases are stored in file formats, which contain records. At physical level, the actual data is stored in electromagnetic format on some device. These storage devices can be broadly categorized into three types −

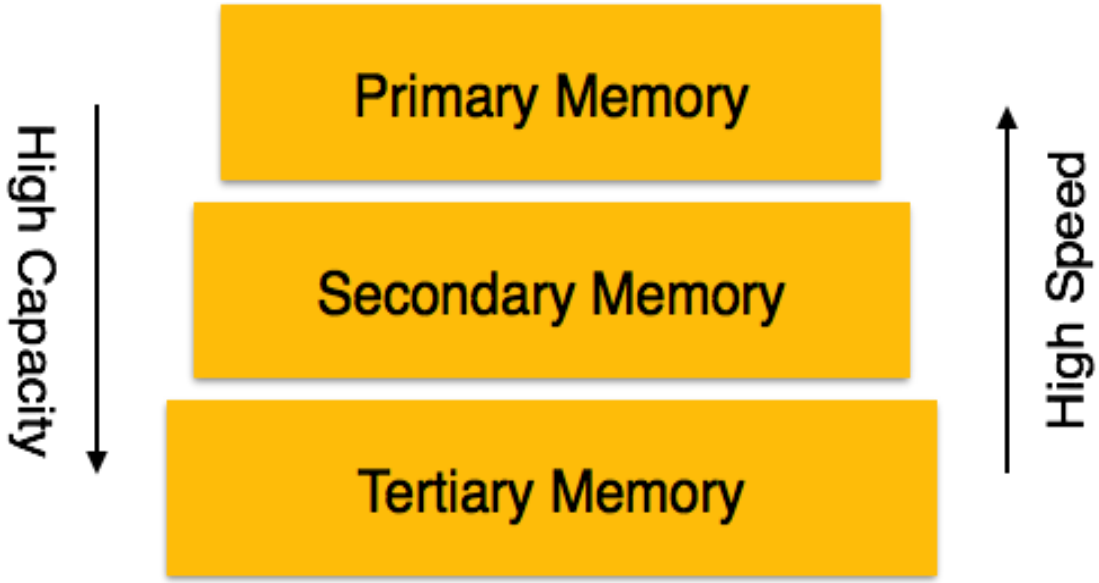

Primary Storage

Primary Storage, also known as main storage or memory, is the main area in a computer in which data is stored for **quick access** by the **computer's processor**. Information must be transferred to primary storage. On today's smaller computers, especially personal computers, the term random access memory (**RAM**) - or just memory - is used instead of primary, main storage, core memory, or immediate access storage.

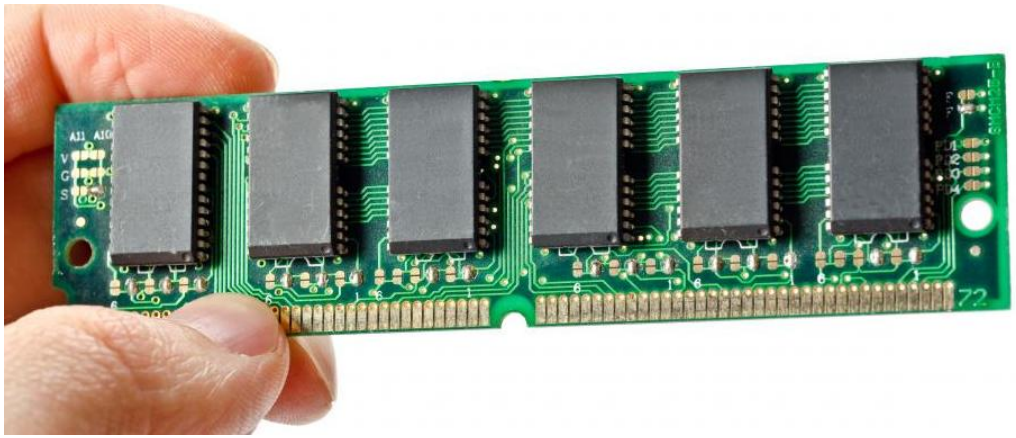

Secondary Storage

Secondary Storage – Secondary storage devices are used to store data for future use or as backup. Secondary storage includes memory devices that are not a part of the CPU for example, magnetic disks, optical disks (DVD, CD, etc.), hard disks, flash drives, and magnetic tapes.

Tertiary Storage

Tertiary Storage – is used to store huge volumes of data. Since such storage devices are external to the computer system, they are the slowest in speed. **Tape libraries** and **Optical jukeboxes** are widely used as tertiary storage. Typically it involves a **robotic mechanism** which will insert and removable mass storage media into a storage device according to the system's demands.

optical jukeboxes

tape library

Redundant Array of Independent Disks

- **RAID** or **R**edundant **A**rray of **I**ndependent **D**isks, is a technology to connect multiple secondary storage devices and use them as a single storage media.
- **RAID** consists of an array of disks in which multiple disks are connected together to achieve different goals.

Redundant Array of Independent Disks

Storage servers with 24 hard disk drives and built-in hardware RAID controllers supporting various RAID levels.

What are the different operating methods RAID uses?

RAID uses the following operating methods : -

1. **Striping** : Dividing a body of data into blocks and spreading the data blocks across several partitions on numerous drives. Striping is primarily used to increase performance.
2. **Mirroring** : Replicating data onto two or more disks to provide good fault tolerance. Data mirroring can be performed without **data loss** or **system stoppages** if a fault occurs.

Data Block

A **Data Block** is the smallest unit of I/O used by a database. Oracle stores data in ***data blocks*** .

One data block corresponds to a specific number of bytes of physical database space on disk. You set the data block size for every Oracle database when you create the database.

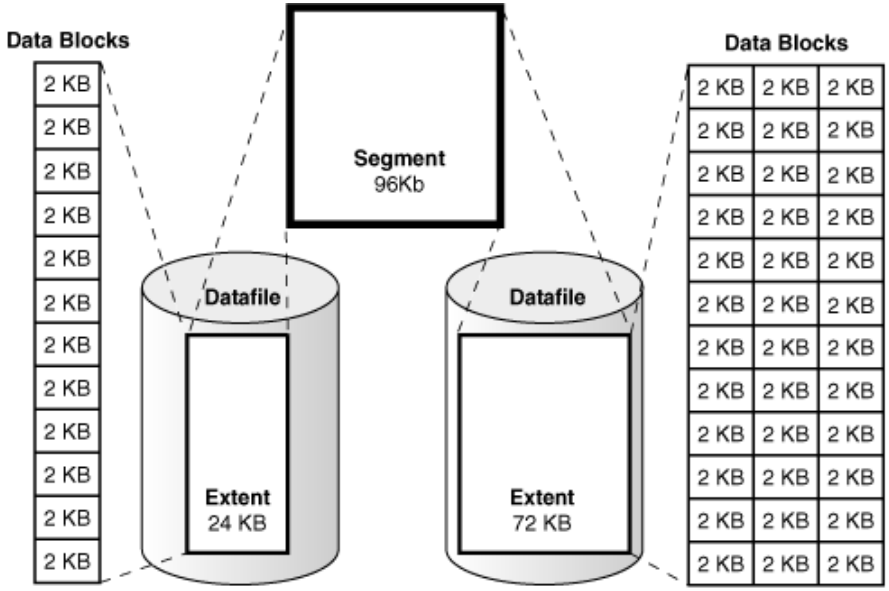

Redundant Array of Independent Disks

RAID 0

RAID 0 writes data to two or more drives alternately to provide the best performance. RAID 0 has a simple design, is easier to implement and has no overheads for parity. However, data is stored on only one disk. So, if one disk fails, data stored within those disks is **lost** since it does not use parity. For this reason, RAID 0 has very limited usage,

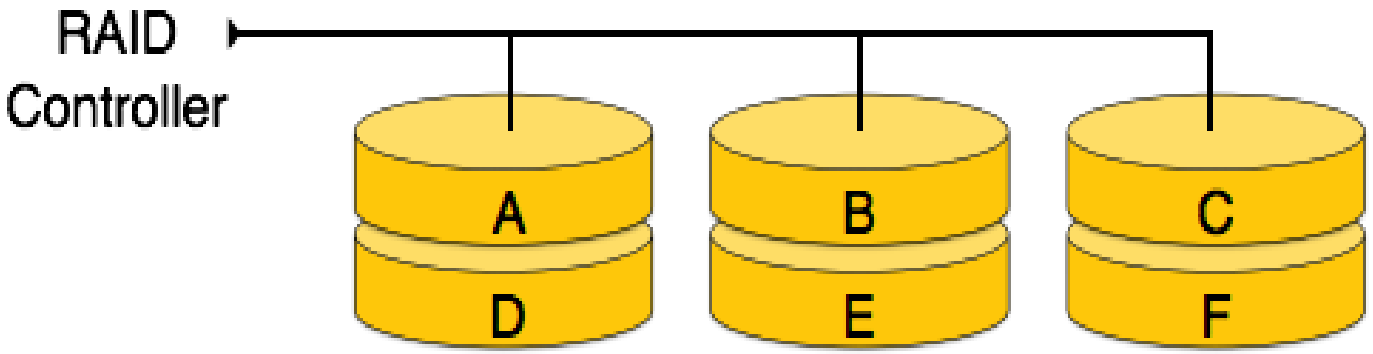

Parity Bit

A **Parity Bit**, or **check bit** is a bit added to the end of a string of binary code that indicates whether the number of bits in the string with the value **one** is **even or odd**. Parity bits are used as the simplest form of **error detecting code**.

7 bits of data	(count of 1 bits)	8 bits including parity	
		even	odd
0000000	0	00000000	00000001
1010001	3	10100011	10100010
1101001	4	11010010	11010011
1111111	7	11111111	11111110

Redundant Array of Independent Disks

RAID 1

RAID 1 provides a **redundant**, identical copy of a disk, yielding good fault tolerance. Replacement of a faulty disk does not affect any data loss or system stoppages. However, **usable capacity** is limited by the smallest capacity and the **write performance** is lower than RAID 0 since it **writes data twice**. Also, RAID 1 has a large storage overhead and a high cost/capacity ratio. RAID level 1 is also called **mirroring** .

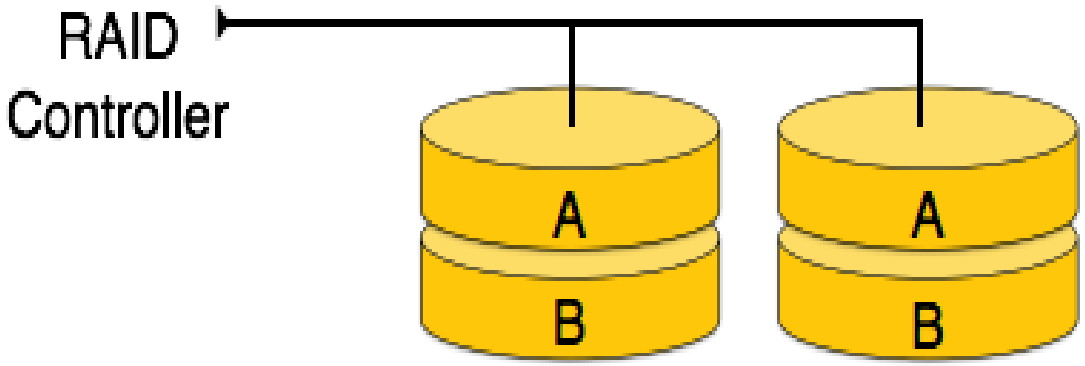

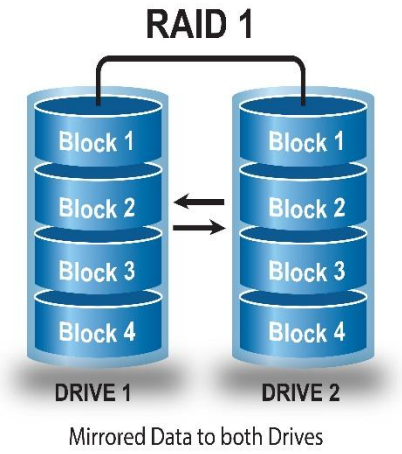

Redundant Array of Independent Disks

RAID 2

Like level 1, with **Error-correcting code** memory (ECC memory). ECC is a type of computer data storage that can detect and correct the most common kinds of internal data corruption. Due to its **complex structure** and **high cost**, RAID 2 is not commercially available.

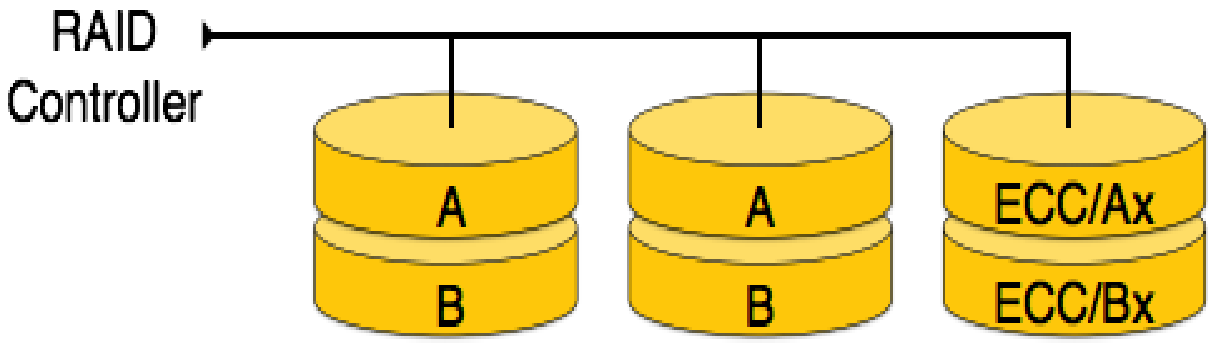

Redundant Array of Independent Disks

RAID 3

RAID 3 stripes the data onto multiple disks. The parity bit generated for data word is stored on a different disk. This technique makes it to overcome single disk failures. In this diagram A and B are bytes. This is not commonly used.

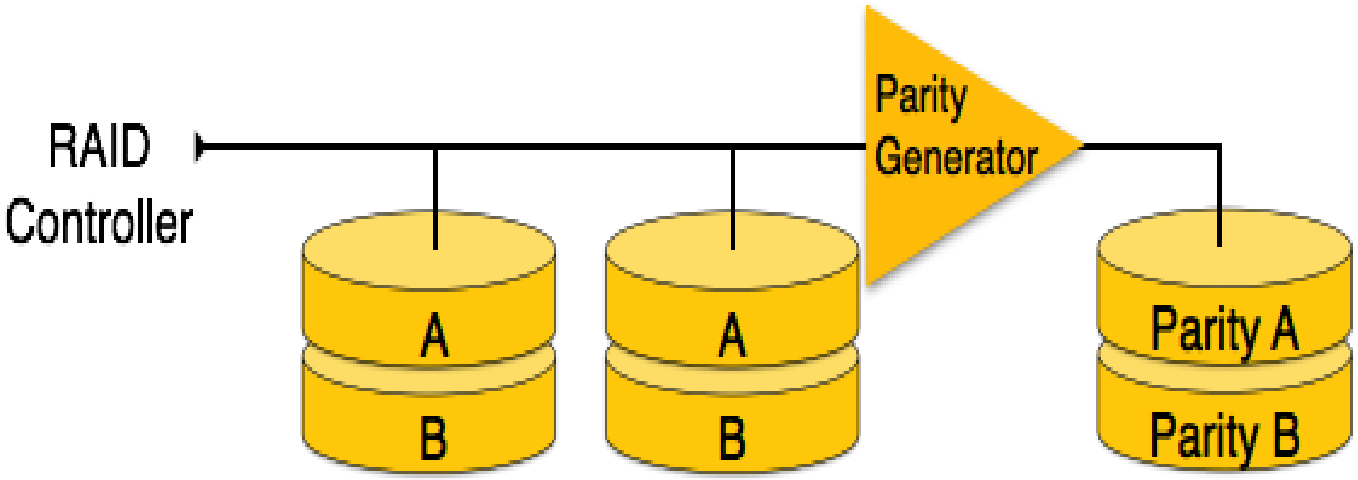

Redundant Array of Independent Disks

RAID 4

In this level, an entire block of data is written onto data disks and then the parity is generated and stored on a different disk. In this diagram A and B are Blocks. This is not commonly used.

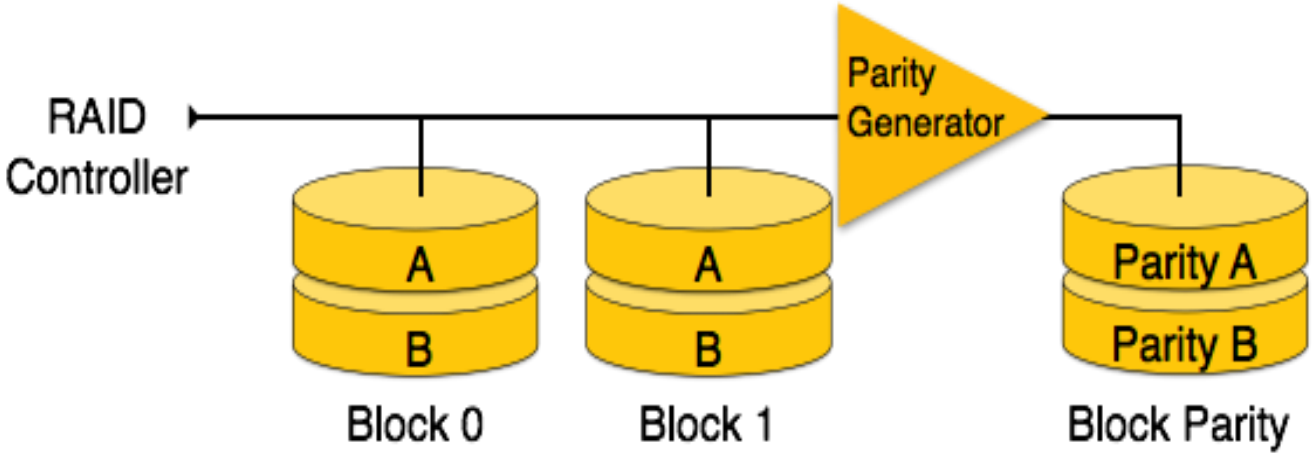

Redundant Array of Independent Disks

RAID 5

RAID 5 uses disk striping with parity across three or more disks. Data redundancy is provided by the parity information. Since **data** and **parity** information are arranged on the disk array, two types of information are always on different disks. If one disk fails, just replace it with a new disk and the array rebuilds itself. RAID 5 has a **higher read rate and makes good use of capacity**. The drawbacks of RAID 5 are slower write rates and slower rebuild times.

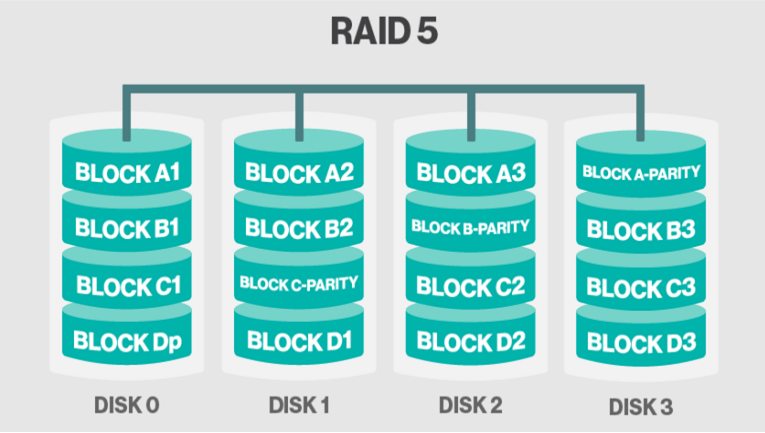

Redundant Array of Independent Disks

RAID 6

Just like RAID 5, this does block level striping. However, it uses dual parity. In this level, two independent parities are generated and stored in distributed fashion among multiple disks. This level requires at least four disk drives to implement RAID. Can handle two disk failure.

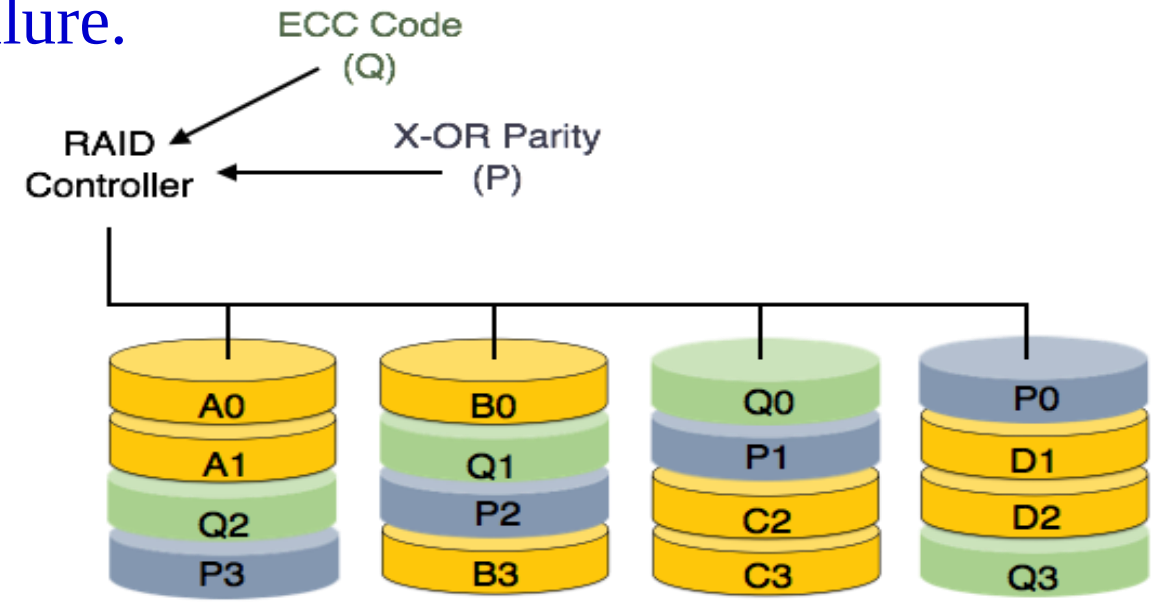

Thank you

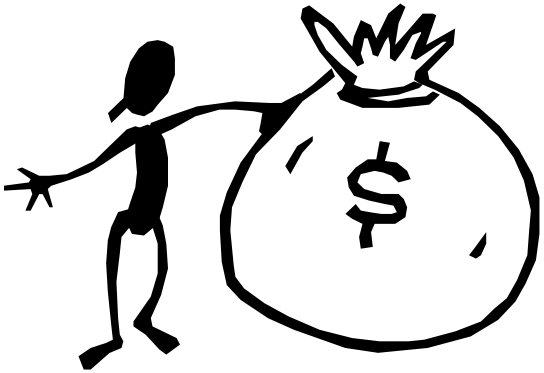

???