CSC 4103 - Operating Systems
Spring 2008

LECTURE - XIX
MASS STORAGE AND I/O - II

Tevfik Koşar

Louisiana State University
April 10th, 2008

RAID Structure

- As disks get cheaper, adding multiple disks to the same system provides increased storage space, as well as increased reliability and performance.
- **RAID: Redundant Array of Inexpensive Disks**
  - multiple disk drives provides **reliability** via **redundancy**.
- RAID is arranged into six different levels.

RAID (cont)

- RAID schemes improve performance and improve the reliability of the storage system by storing redundant data.
  - *Mirroring (shadowing): duplicate each disk*
    - *Simplest but most expensive approach*
  - *Block interleaved parity* uses much less redundancy.
  - Data Striping: splitting each bit (or block) of a file across multiple disks.

RAID Levels

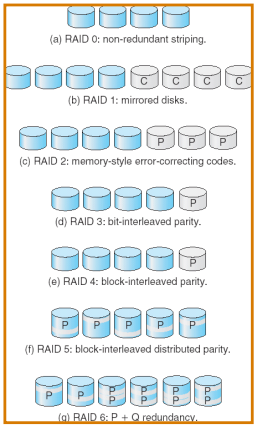

RAID Level 0

- Data is divided into blocks and is spread in a fixed order among all the disks in the array
- does not provide any fault tolerance
- also known as disk striping
- improves read and write performance via parallel access

RAID Level 0

Disk 0	Disk 1	Disk 2	Disk 3	Disk 4
Block 1	Block 2	Block 3	Block 4	Block 5
Block 6	Block 7	Block 8	Block 9	Block 10
Block 11	Block 12	Block 13	Block 14	Block 15
Block 16	Block 17	Block 18	Block 19	Block 20
Block 21	Block 22	Block 23	Block 24	Block 25

RAID Level 1

- All data written to the primary disk is written to the mirror disk
- provides a redundant, identical copy of all data
- provides fault tolerance
- also known as disk mirroring
- also generally improves read performance (but may degrade write performance).

RAID Level 1

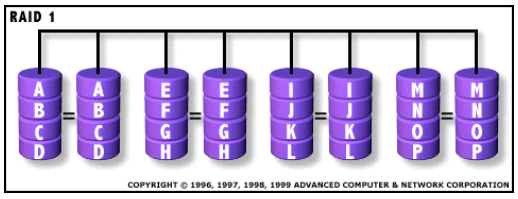

RAID Level 2

- uses error correcting algorithm that employs disk-striping strategy that breaks a file into bytes and spreads it across multiple disks
- The error-correction method requires several disks
- provides fault tolerance
- but is not as efficient as other RAID levels

RAID Level 2

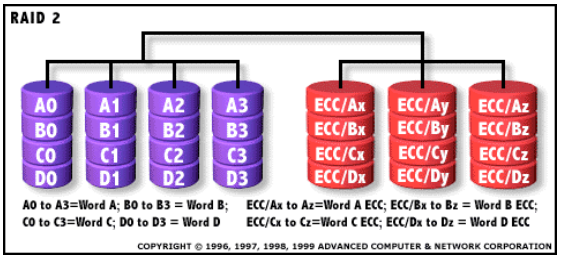

RAID Level 3

- similar to RAID level 2, but it requires only one disk for parity data
- suffers from a write bottleneck, because all parity data is written to a single drive
- but provides some read and write performance improvement.

RAID Level 3

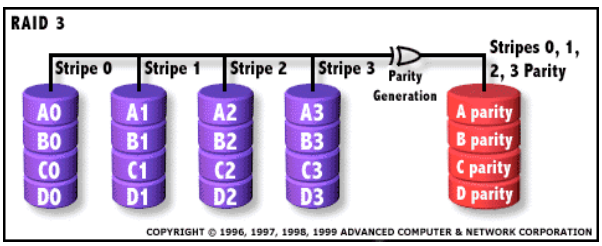

RAID Level 4

- similar to RAID level 3, but it employs striped data in much larger blocks or segments
- not as efficient as RAID level 5, because (as in RAID level 3) all parity data is written to a single drive
- so RAID level 4 suffers from a write bottleneck and is not generally used.

RAID Level 4

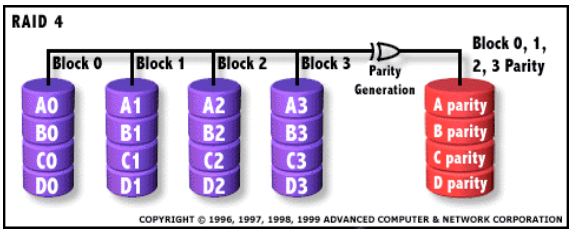

RAID Level 5

- known as striping with parity
- the most popular RAID level
- similar to level 4 in that it stripes the data in large blocks across all the disks in the array
- It differs in that it writes the parity across all the disks
- The data redundancy is provided by the parity information
- The data and parity information are arranged on the disk array so that the two are always on different disks

RAID Level 5

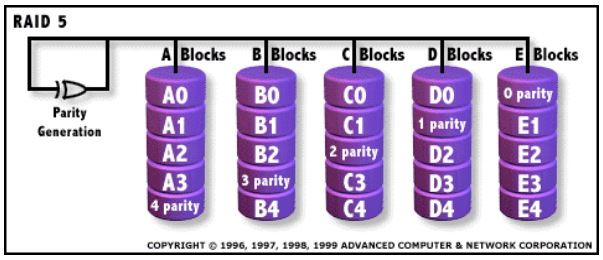

RAID Level 6

- Similar to RAID 5, but uses dual distributed parity
- more reliability versus less data space

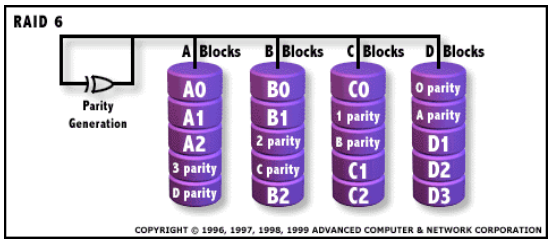

Hierarchical Storage Management (HSM)

- A hierarchical storage system extends the storage hierarchy beyond primary memory and secondary storage to incorporate tertiary storage — usually implemented as a jukebox of tapes or removable disks.
- Usually incorporate tertiary storage by extending the file system.
  - Small and frequently used files remain on disk.
  - Large, old, inactive files are archived to the jukebox.
- HSM is usually found in supercomputing centers and other large installations that have enormous volumes of data.

Hierarchical Storage Management