7.9 RAID

- RAID, an acronym for *Redundant Array of Independent Disks* was invented to address problems of disk reliability, cost, and performance.
- In RAID, data is stored across many disks, with extra disks added to the array to provide error correction (redundancy).
- The inventors of RAID, David Patterson, Garth Gibson, and Randy Katz, provided a RAID taxonomy that has persisted for a quarter of a century, despite many efforts to redefine it.

7.9.1 RAID Level 0

- RAID Level 0, also known as *drive spanning*, provides improved performance, but no redundancy.
  - Data is written in blocks across the entire array

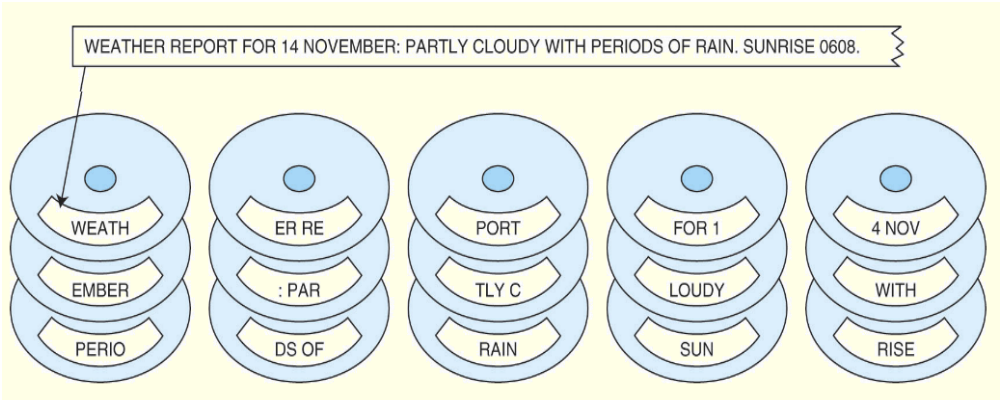

Figure 7.26 A Record Written Using RAID-0, Block Interleave Data Striping with No Redundancy

  - The disadvantage of RAID 0 is in its low reliability.

7.9.2 RAID Level 1

- RAID Level 1, also known as *disk mirroring*, provides 100% redundancy, and good performance.
  - Two matched sets of disks contain the same data.

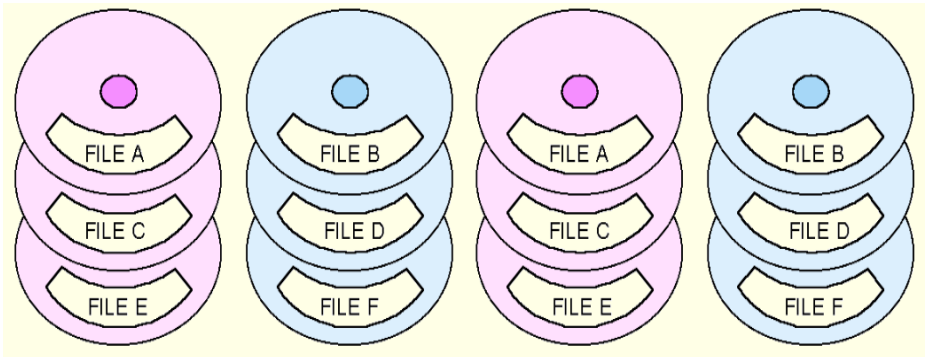

Figure 7.27 RAID-1, Disk Mirroring

  - The disadvantage of RAID 1 is cost.

7.9.3 RAID Level 2

- A RAID Level 2 configuration consists of a set of data drives, and a set of Hamming code drives.
  - Hamming code drives provide error correction for the data drives.

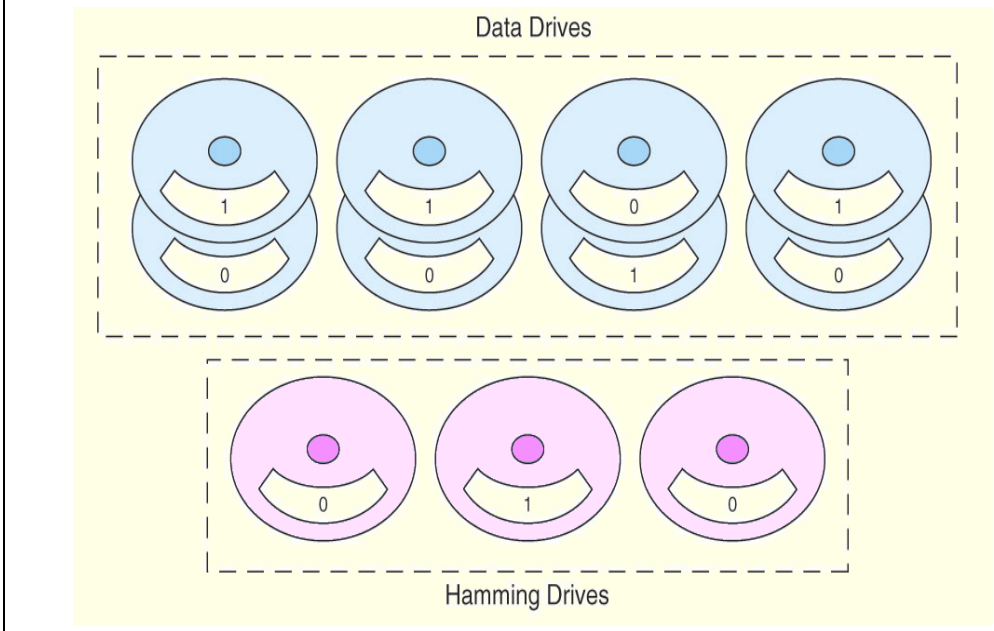

Figure 7.28 RAID-2,Bit Interleave Data Striping with a Hamming Code

  - RAID 2 performance is poor and the cost is relatively high.
  - Forms the theoretical bridge between RAID-1 and RAID-3.

7.9.4 RAID Level 3

- RAID Level 3 stripes bits across a set of data drives and provides a separate disk for parity.
  - Parity is the XOR of the data bits.
  - Data on the other seven data drives and the parity drive can be used to reconstruct the data on the damaged drive.

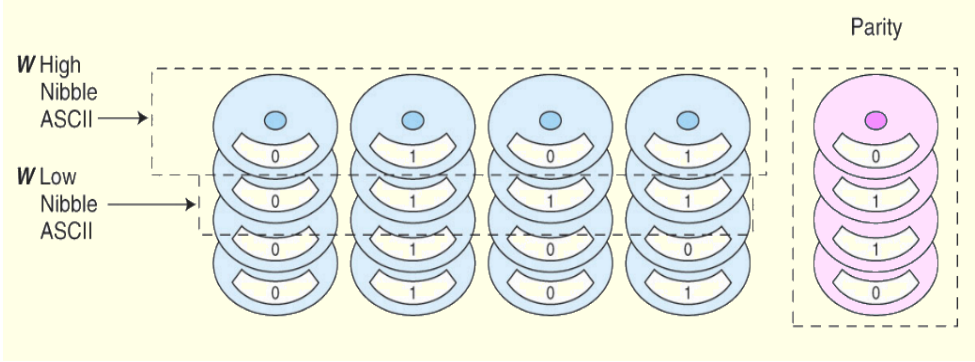

Figure 7.29 RAID-3,Bit Interleave Data Striping with Parity Disk

  - RAID 3 is not suitable for commercial applications, but is good for personal systems.
  - RAID 3 is not well suited for transaction-oriented applications

7.9.5 RAID Level 4

- RAID Level 4 is like adding parity disks to RAID 0.
  - Data is written in blocks across the data disks, and a parity block is written to the redundant drive.

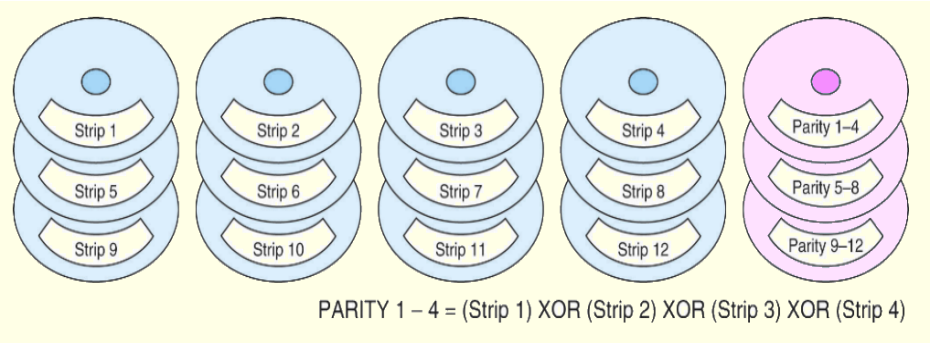

Figure 7.30 RAID-4,Block Interleave Data Striping with One Parity Disk

  - RAID 4 would be feasible if all record blocks were the same size.

7.9.6 RAID Level 5

- RAID Level 5 is RAID 4 with distributed parity.
  - With distributed parity, some accesses can be serviced concurrently, giving good performance and high reliability.

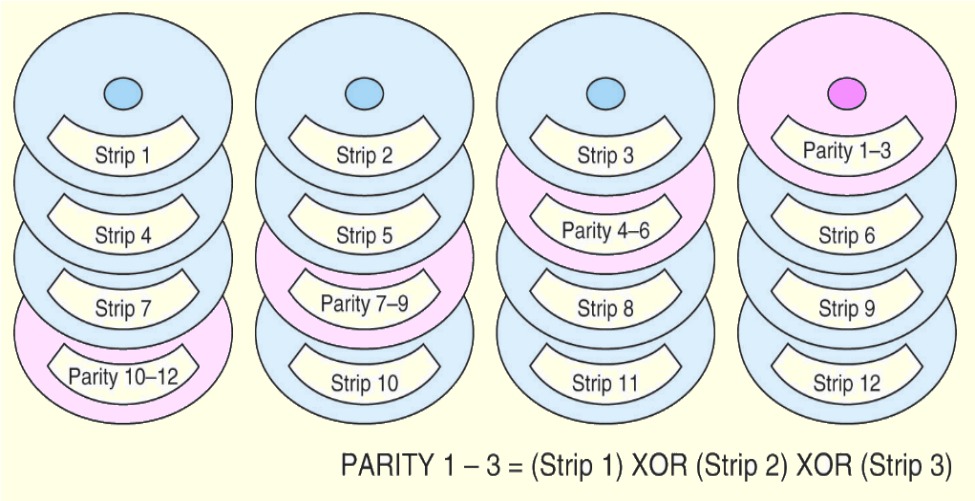

Figure 7.31 RAID-5,Block Interleave Data Striping with Distributed Parity

  - RAID 5 is used in many commercial systems.

7.9.7 RAID Level 6

- RAID Level 6 carries two levels of error protection over striped data: Reed-Soloman and parity.
  - It can tolerate the loss of two disks.

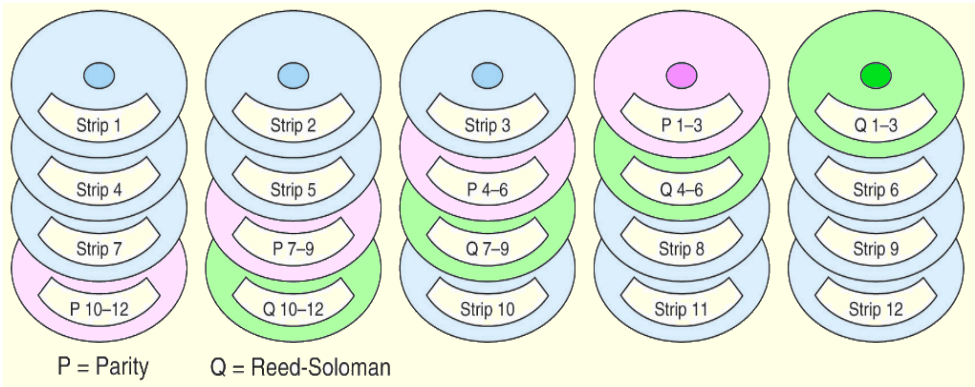

Figure 7.32 RAID-5,Block Interleave Data Striping with Dual Error Protection

  - RAID 6 is write-intensive, but highly fault-tolerant.

7.9.8 RAID DP

- Double parity RAID (RAID DP) employs pairs of overlapping parity blocks that provide linearly independent parity functions.

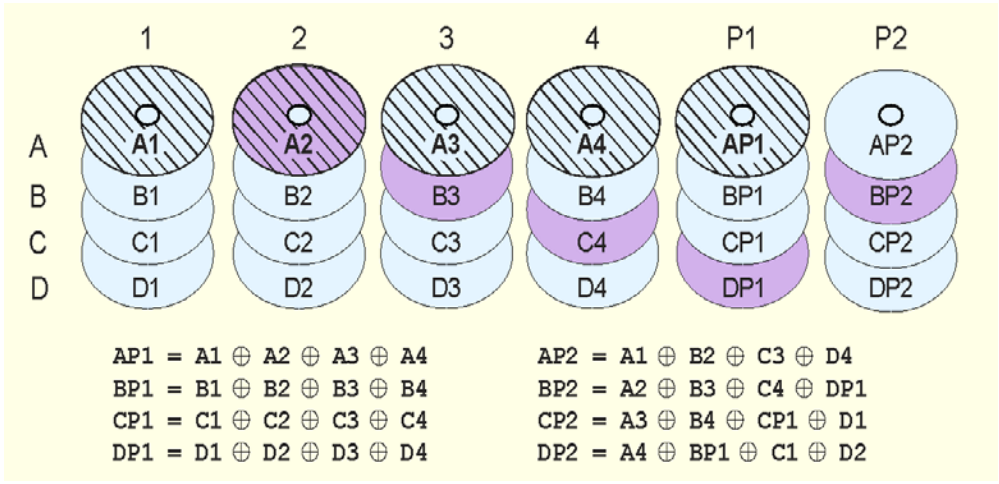

Figure 7.33 Error Recovery Pattern for RAID DP

The recovery of A2 is provided by the overlap of equations AP1 and BP2

- Like RAID 6, RAID DP can tolerate the loss of two disks.
- The use of simple parity functions provides RAID DP with better performance than RAID 6.
- Of course, because two parity functions are involved, RAID DP's performance is somewhat degraded from that of RAID 5.
  - RAID DP is also known as EVENODD, diagonal parity RAID, RAID 5DP, advanced data guarding RAID (RAID ADG) and-- erroneously-- RAID 6.

- Large systems consisting of many drive arrays may employ various RAID levels, depending on the criticality of the data on the drives.
  - A disk array that provides program workspace (say for file sorting) does not require high fault tolerance.
- Critical, high-throughput files can benefit from combining RAID 0 with RAID 1, called RAID 10.
- RAID 50 combines striping and distributed parity. For good fault tolerance and high capacity.
  - Note: Higher RAID levels do not necessarily mean "better" RAID levels. It all depends upon the needs of the applications that use the disks.