This is the fourth lecture of Chapter 7

Chapter 7

Input/Output Systems (D)

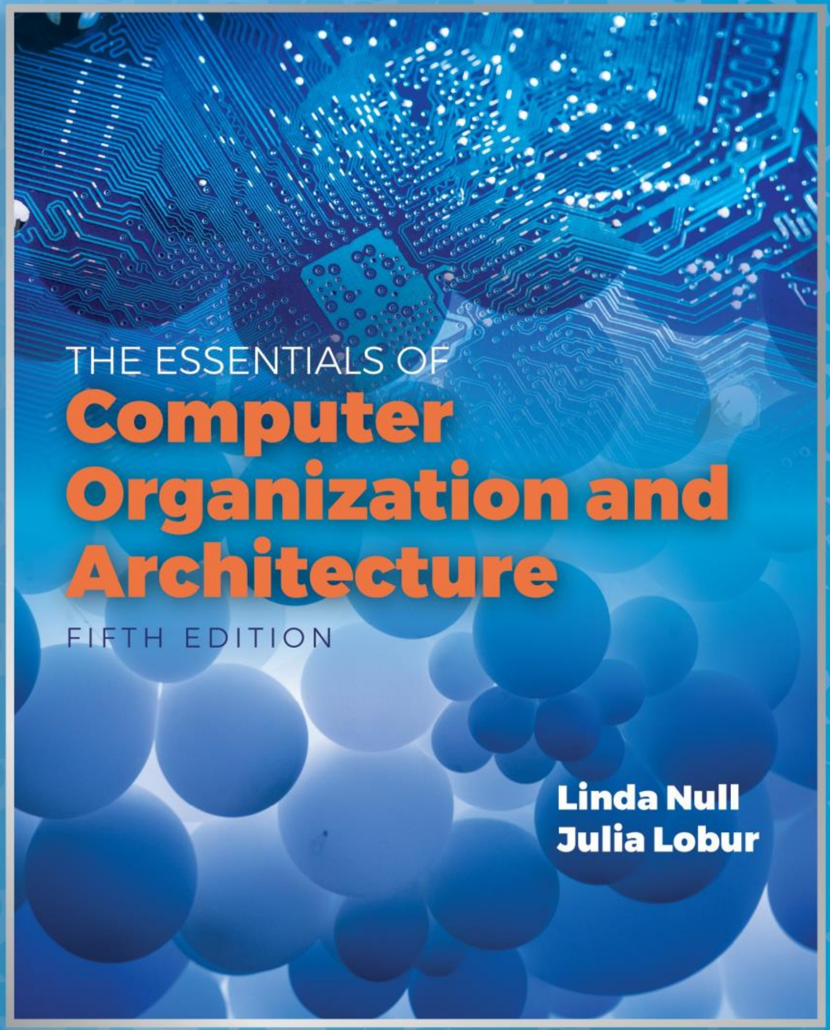

© Nicemonkey/Shutterstock. Copyright © 2019 by Jones & Bartlett Learning, LLC an Ascend Learning Company
www.jblearning.com

Quick review of last lecture

- Example: Disk Capacity and Performance
- Master Control Block (Disk Directory)
- Advantage and limitations of hard disks
- Solid State Drives
  - Store data in non-volatile *flash* memory circuits
  - Access time and transfer rates are *typically* 100 times faster than magnetic disk
  - Flash is block-addressable (like disk drives)
- Optical Disks (Reading)
  - Data is recorded in a single spiral track, starting from the center of the disk and spanning outward

© Nicemonkey/Shutterstock. Copyright © 2019 by Jones & Bartlett Learning, LLC an Ascend Learning Company
www.jblearning.com

7.8 Magnetic Tape (1 of 6)

- First-generation magnetic tape was not much more than wide analog recording tape, having capacities under 11MB.
- Data was usually written in nine vertical tracks:

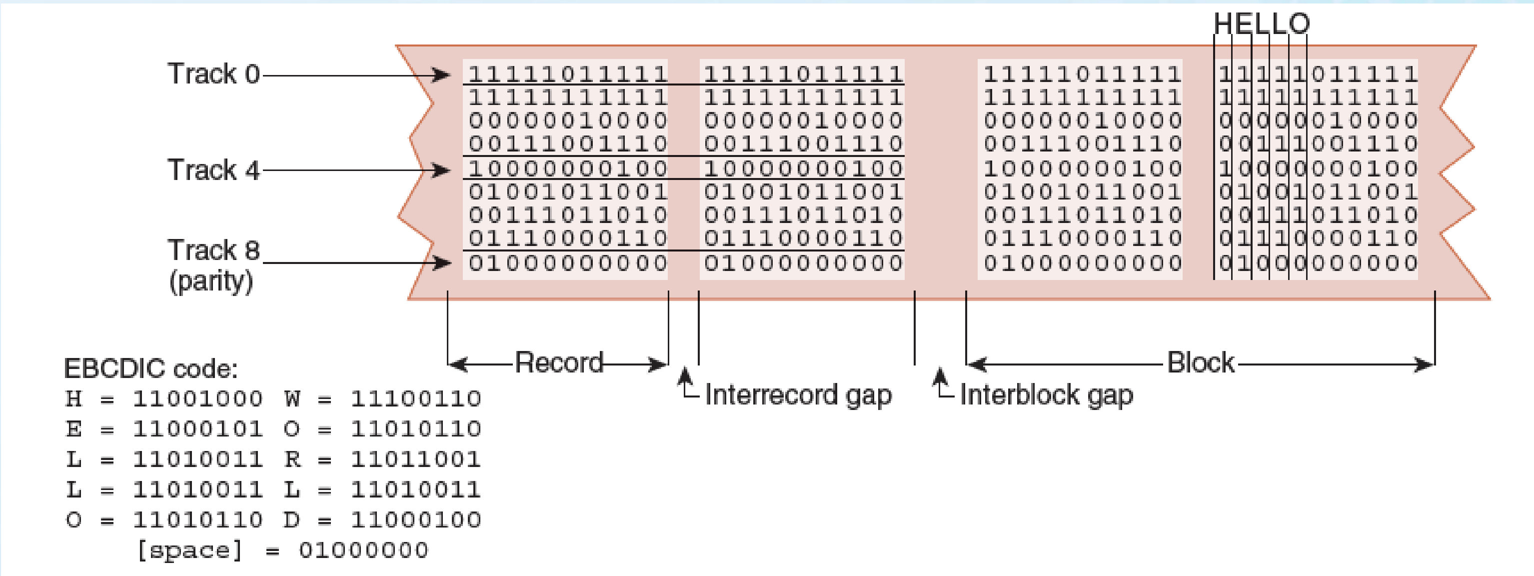

© Nicemonkey/Shutterstock. Copyright © 2019 by Jones & Bartlett Learning, LLC an Ascend Learning Company
www.jblearning.com

7.8 Magnetic Tape (2 of 6)

- Today's tapes are digital, and provide multiple gigabytes of data storage.
- Two dominant recording methods are *serpentine* and *helical scan*, which are distinguished by how the read-write head passes over the recording medium.
- Serpentine recording is used in *digital linear tape* (DLT) and *quarter inch cartridge* (QIC) tape systems.
- *Digital audio tape* (DAT) systems employ helical scan recording.
- These two recording methods are shown on the next slide.

© Nicemonkey/Shutterstock. Copyright © 2019 by Jones & Bartlett Learning, LLC an Ascend Learning Company
www.jblearning.com

7.8 Magnetic Tape (3 of 6)

Serpentine

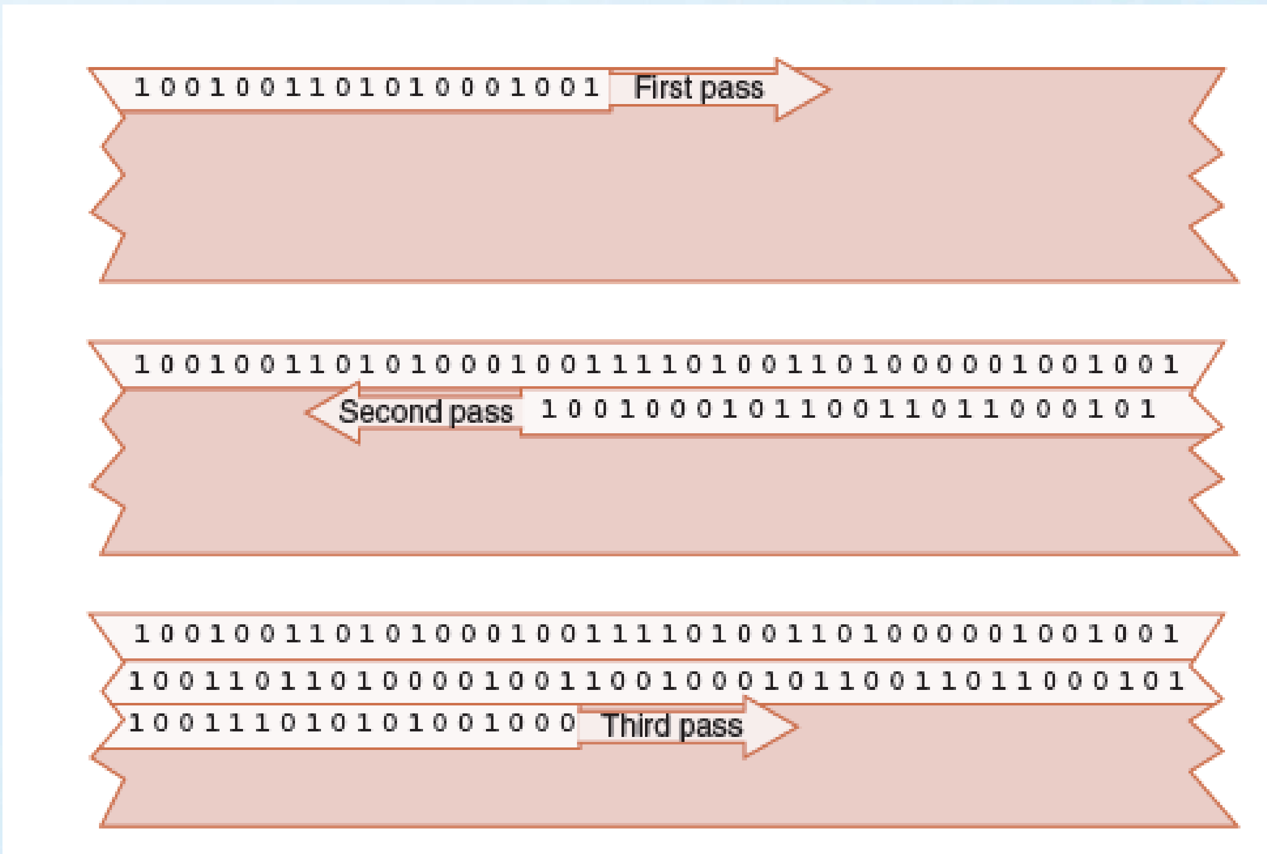

© Nicemonkey/Shutterstock. Copyright © 2019 by Jones & Bartlett Learning, LLC an Ascend Learning Company
www.jblearning.com

7.8 Magnetic Tape (4 of 6)

Helical Scan

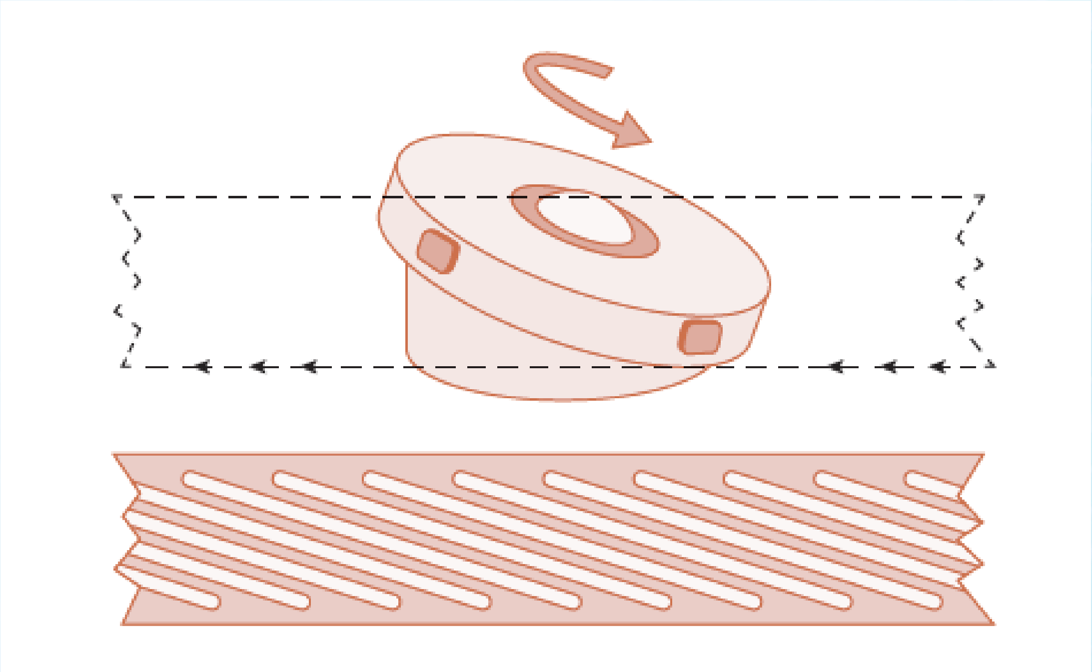

- More complicated to implement
- More wear on the tape and is generally less reliable
- Slower

© Nicemonkey/Shutterstock. Copyright © 2019 by Jones & Bartlett Learning, LLC an Ascend Learning Company
www.jblearning.com

7.8 Magnetic Tape (5 of 6)

- Numerous incompatible tape formats emerged over the years.
  - Sometimes even different models of the same manufacturer's tape drives were incompatible!
- Finally, in 1997, HP, IBM, and Seagate collaboratively invented a best-of-breed tape standard.
- They called this new tape format *Linear Tape Open* (LTO) because the specification is openly available.

© Nicemonkey/Shutterstock. Copyright © 2019 by Jones & Bartlett Learning, LLC an Ascend Learning Company
www.jblearning.com

7.8 Magnetic Tape (6 of 6)

- LTO, as the name implies, is a linear digital tape format.
- The specification allowed for the refinement of the technology through four "generations."
- Generation 5 was released in 2010.
  - Without compression, the tapes support a transfer rate of 208MB per second and each tape can hold up to 1.4TB.
- LTO supports several levels of error correction, providing superb reliability.
  - Tape has a reputation for being an error-prone medium.

© Nicemonkey/Shutterstock. Copyright © 2019 by Jones & Bartlett Learning, LLC an Ascend Learning Company
www.jblearning.com

7.9 RAID (1 of 11)

- 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.

© Nicemonkey/Shutterstock. Copyright © 2019 by Jones & Bartlett Learning, LLC an Ascend Learning Company
www.jblearning.com

7.9 RAID (2 of 11)

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

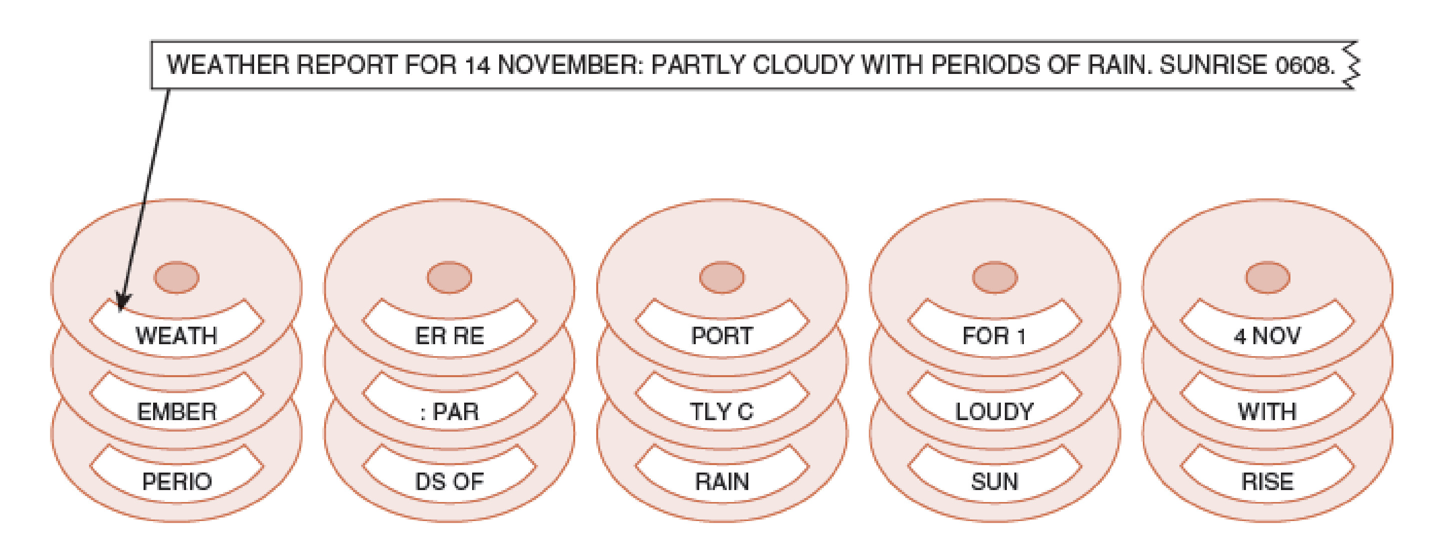

  - The disadvantage of RAID 0 is in its low reliability.

© Nicemonkey/Shutterstock. Copyright © 2019 by Jones & Bartlett Learning, LLC an Ascend Learning Company
www.jblearning.com

7.9 RAID (3 of 11)

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

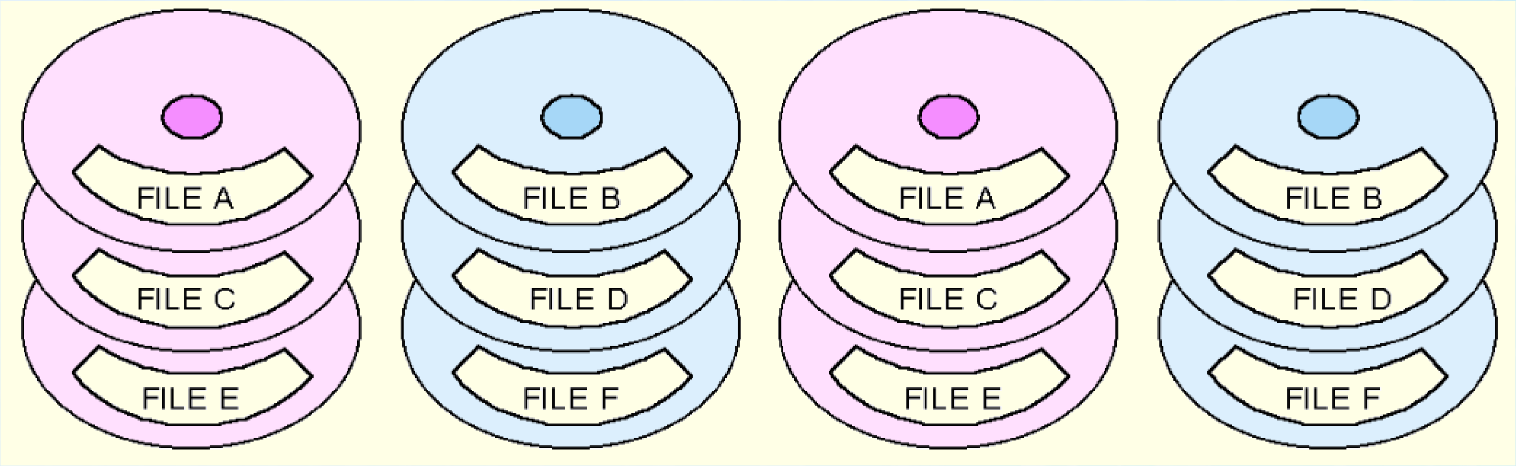

  - The disadvantage of RAID 1 is cost.

© Nicemonkey/Shutterstock. Copyright © 2019 by Jones & Bartlett Learning, LLC an Ascend Learning Company
www.jblearning.com

7.9 RAID (4 of 11)

- 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.

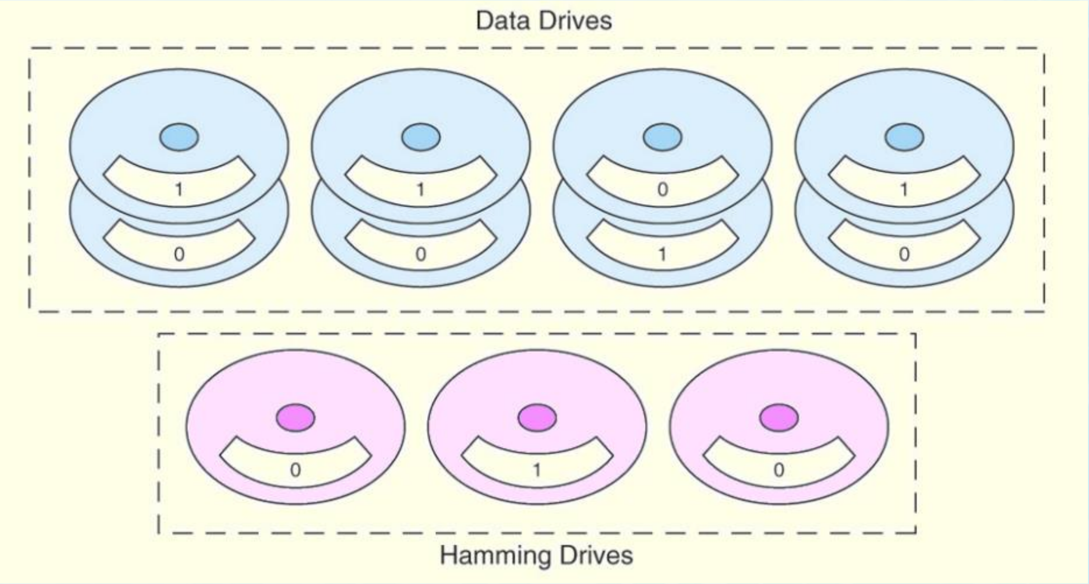

  - RAID 2 performance is poor and the cost is relatively high.

© Nicemonkey/Shutterstock. Copyright © 2019 by Jones & Bartlett Learning, LLC an Ascend Learning Company
www.jblearning.com

7.9 RAID (5 of 11)

- 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.

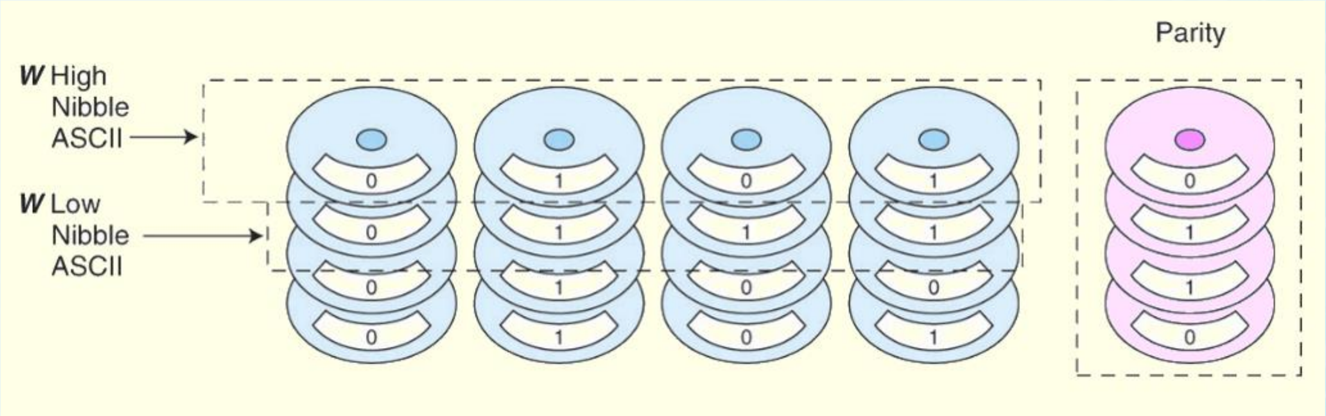

  - RAID 3 is not suitable for commercial applications, but is good for personal systems.

© Nicemonkey/Shutterstock. Copyright © 2019 by Jones & Bartlett Learning, LLC an Ascend Learning Company
www.jblearning.com

7.9 RAID (6 of 11)

- 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.

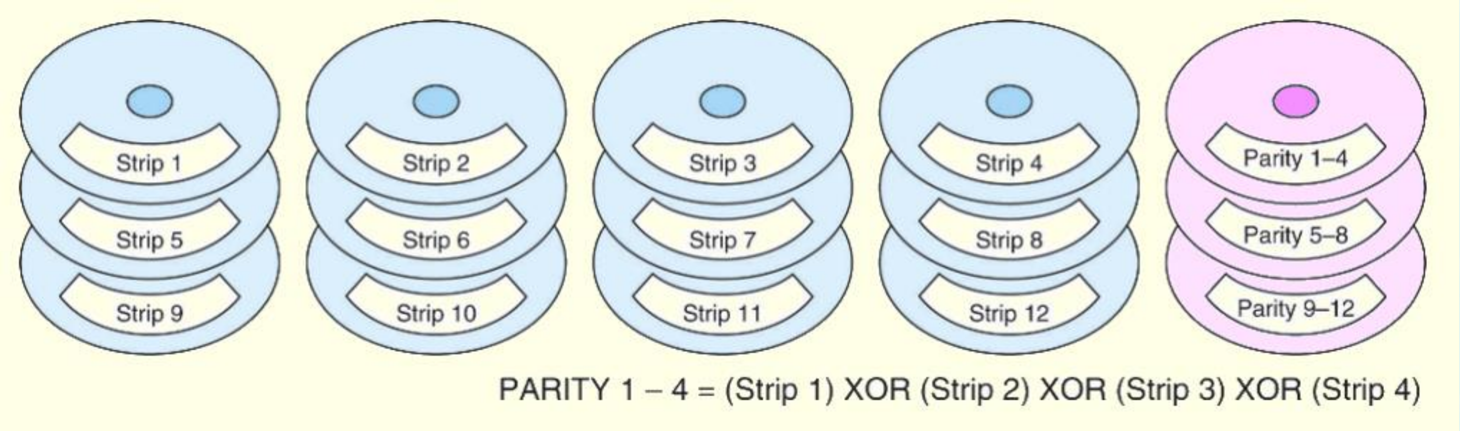

  - RAID 4 would be feasible if all record blocks were the same size.

© Nicemonkey/Shutterstock. Copyright © 2019 by Jones & Bartlett Learning, LLC an Ascend Learning Company
www.jblearning.com

7.9 RAID (7 of 11)

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

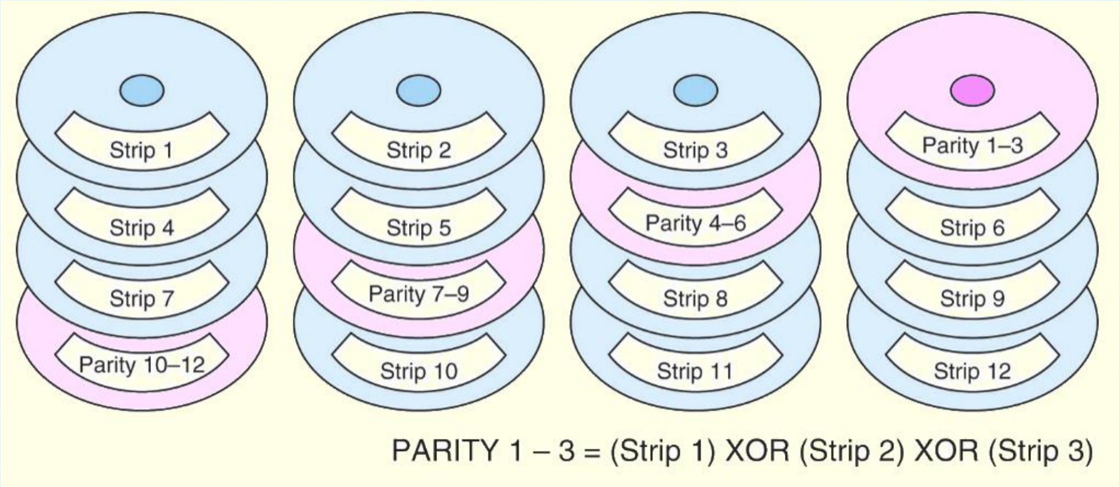

  - RAID 5 is used in many commercial systems.

© Nicemonkey/Shutterstock. Copyright © 2019 by Jones & Bartlett Learning, LLC an Ascend Learning Company
www.jblearning.com

7.9 RAID (8 of 11)

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

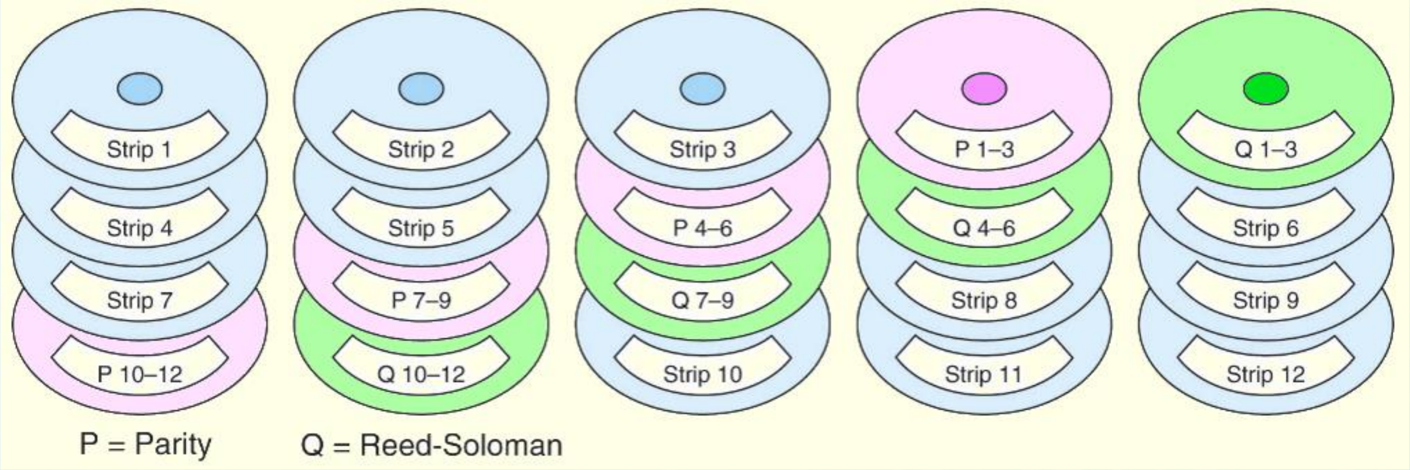

  - RAID 6 is write-intensive, but highly fault-tolerant.

© Nicemonkey/Shutterstock. Copyright © 2019 by Jones & Bartlett Learning, LLC an Ascend Learning Company
www.jblearning.com

7.9 RAID (9 of 11)

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

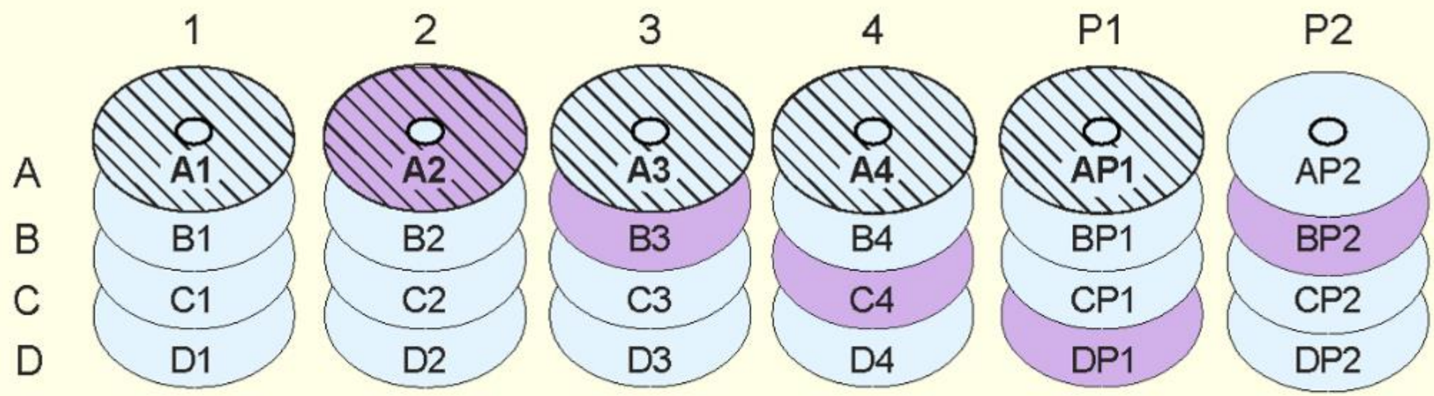

$$AP1 = A1 \oplus A2 \oplus A3 \oplus A4$$
$$BP1 = B1 \oplus B2 \oplus B3 \oplus B4$$
$$CP1 = C1 \oplus C2 \oplus C3 \oplus C4$$
$$DP1 = D1 \oplus D2 \oplus D3 \oplus D4$$

$$AP2 = A1 \oplus B2 \oplus C3 \oplus D4$$
$$BP2 = A2 \oplus B3 \oplus C4 \oplus DP1$$
$$CP2 = A3 \oplus B4 \oplus CP1 \oplus D1$$
$$DP2 = A4 \oplus BP1 \oplus C1 \oplus D2$$

© Nicemonkey/Shutterstock. Copyright © 2019 by Jones & Bartlett Learning, LLC an Ascend Learning Company
www.jblearning.com

7.9 RAID (10 of 11)

- 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.

© Nicemonkey/Shutterstock. Copyright © 2019 by Jones & Bartlett Learning, LLC an Ascend Learning Company
www.jblearning.com

7.9 RAID (11 of 11)

- 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.

© Nicemonkey/Shutterstock. Copyright © 2019 by Jones & Bartlett Learning, LLC an Ascend Learning Company
www.jblearning.com