

# Mass Storage

Jonathan Geisler

April 5, 2006



# Disk management

i.e., Low-level details

- Formatting
  - Physical
  - Logical
- Partitioning
- Boot block(s)
- Bad blocks
  - Forwarding (sector sparing)
  - Sector slipping



# Swap-space management

- Suffers from fragmentation
- Location
  - File inside file system
  - Separate partition
- Allocation
  - At process creation
  - As needed



- More disks = more failures
- Redundancy allows
  - Problem fixes without data loss
  - Increased throughput or transfer rate
- Accomplished via
  - Mirroring
  - Error correcting codes (e.g., Hamming)
  - Parity bits

- Traditional
  - ① None. Striping

- Traditional
  - ① None. Striping
  - ① Mirror

- Traditional
  - ① None. Striping
  - ① Mirror
  - ② ECC

- Traditional
  - ① None. Striping
  - ② Mirror
  - ③ ECC
  - ④ Bit-interleaved

- Traditional
  - ① None. Striping
  - ① Mirror
  - ② ECC
  - ③ Bit-interleaved
  - ④ Block-interleaved

- Traditional
  - ① None. Striping
  - ② Mirror
  - ③ ECC
  - ④ Bit-interleaved
  - ⑤ Block-interleaved
  - ⑥ Distributed block-interleaved

- Traditional
  - ① None. Striping
  - ② Mirror
  - ③ ECC
  - ④ Bit-interleaved
  - ⑤ Block-interleaved
  - ⑥ Distributed block-interleaved
  - ⑦ P + Q parity

- Traditional
  - ① None. Striping
  - ② Mirror
  - ③ ECC
  - ④ Bit-interleaved
  - ⑤ Block-interleaved
  - ⑥ Distributed block-interleaved
  - ⑦ P + Q parity
- 0+1 or 1+0 combination of levels 0 and 1

- Media
  - Removable disks
  - Tapes
- Robotic tape libraries
- Futuristic?
  - Holographic
  - MEMS

# How does the OS deal with removable media?

- API
  - Disks unchanged
  - Tapes only get raw I/O
- File names left to user
- Can manage robotic jukebox directly like any other device
  - Staging old files to tape
  - NOT utilizing taps for virtual memory

# What to consider ...

## 1 Speed



# What to consider ...

- 1 Speed
  - Bandwidth



# What to consider ...

- 1 Speed
  - Bandwidth
  - Latency



# What to consider ...

- 1 Speed
  - Bandwidth
  - Latency
- 2 Reliability



# What to consider ...

- 1 Speed
  - Bandwidth
  - Latency
- 2 Reliability
  - Hard drive > optical storage > tape

# What to consider . . .

## 1 Speed

- Bandwidth
- Latency

## 2 Reliability

- Hard drive > optical storage > tape
- Head crashes are near fatal

# What to consider . . .

- 1 Speed
  - Bandwidth
  - Latency
- 2 Reliability
  - Hard drive > optical storage > tape
  - Head crashes are near fatal
- 3 Cost

# What to consider . . .

## 1 Speed

- Bandwidth
- Latency

## 2 Reliability

- Hard drive > optical storage > tape
- Head crashes are near fatal

## 3 Cost

- Traditionally: DRAM > disk > tape

# What to consider ...

## 1 Speed

- Bandwidth
- Latency

## 2 Reliability

- Hard drive > optical storage > tape
- Head crashes are near fatal

## 3 Cost

- Traditionally: DRAM > disk > tape
- Now: DRAM > tape  $\geq$  disk!!!