16th Edition

Understanding Computers

Today and Tomorrow

Comprehensive

Chapter 3
Storage

Deborah Morley
Charles S. Parker

Memory/Storage Hierarchy

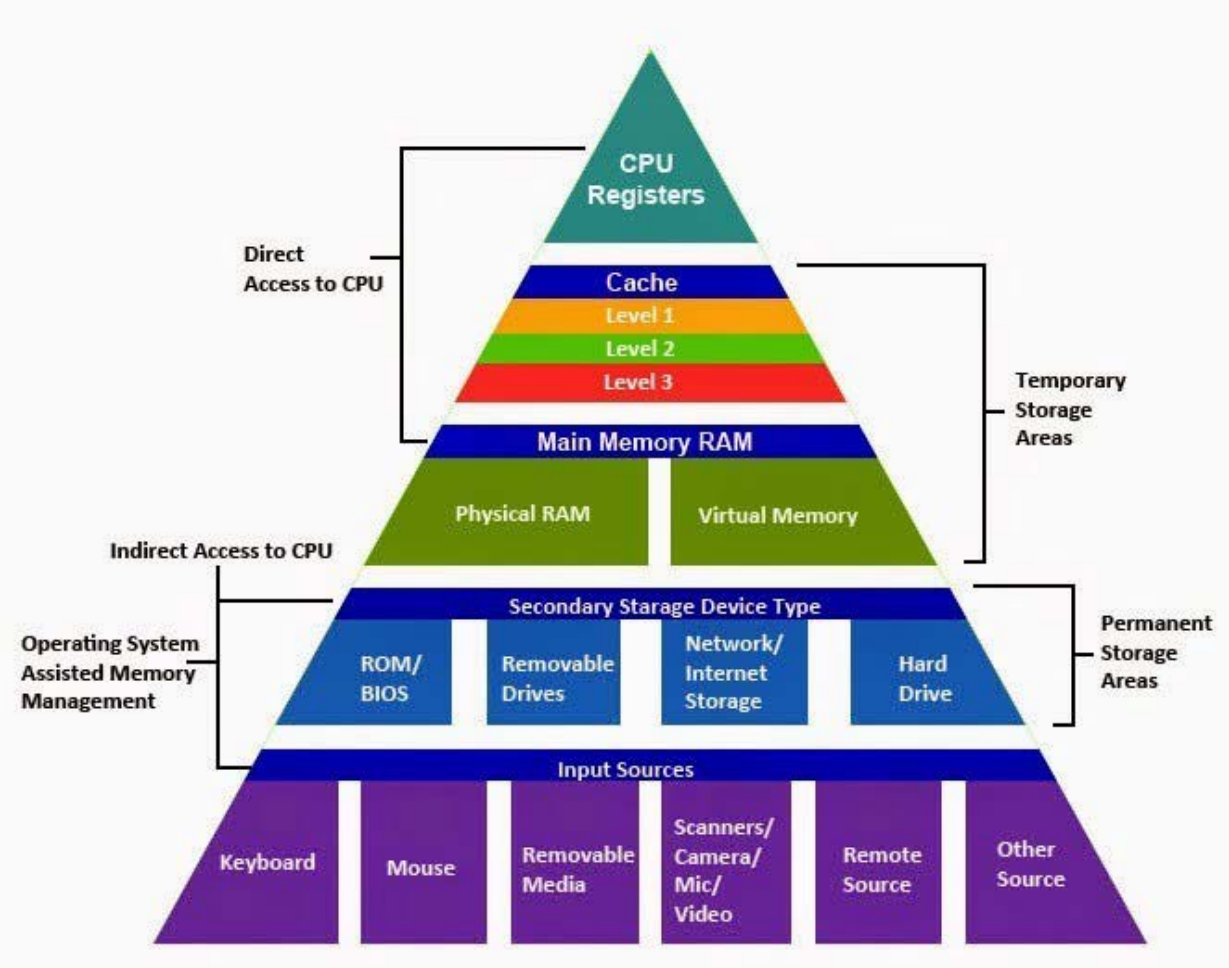

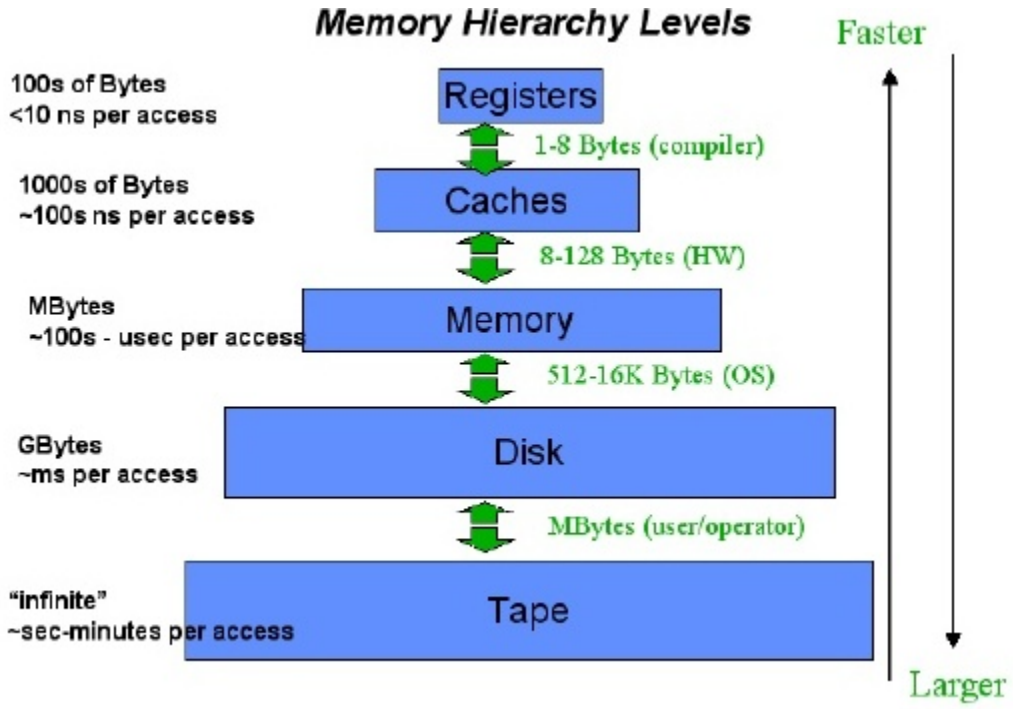
Memory Hierarchy Levels
Faster
Larger
100s of Bytes
<10 ns per access
Registers
1-8 Bytes (compiler)
1000s of Bytes
~100s ns per access
Caches
8-128 Bytes (HW)
MBytes
~100s - usec per access
Memory
512-16K Bytes (OS)
GBytes
~ms per access
Disk
MBytes (user/operator)
"infinite"
~sec-minutes per access
Tape

Storage

Kryder's Law

- Hard drives => exponential increase in density (bits per unit area) of information they are able to store.
- Essentially Moore's Law for storage
  - density of information growing at an even faster rate
  - Increasing by a factor of 1000 in 10.5 years
- Future of Kryder's law
  - If current rates of growth are maintained then in 2 decades, a consumer will be able to store all of the creative works produced by every member of the human species in a $100 storage device, including realtime video capture of one's entire lifetime.

Evaluating Storage Alternatives

- Characteristics to consider:
  - Speed
  - Compatibility
  - Capacity
  - Convenience,
  - Portability
  - (absent from text => cost)
- Each storage alternative involves trade-offs

Logical vs. Physical Representation and Types of Storage Technologies Used

- Logical file representation
  - Individuals view file as one complete unit
  - In a particular folder on a particular drive
- Physical file representation
  - Computers access a particular document stored on a storage medium using its physical location or locations

$\Rightarrow$ OS manages Abstraction

$\Rightarrow$ (and necessarily brings in OS discussion in this chapter)

Volatility and Random vs. Sequential Access

- Volatility
  - Storage media => Nonvolatile
    - Used to save data for later use
- Random vs. sequential access
  - Random access (direct access)
    - Data retrieved from any location on storage medium
    - Virtually all storage devices use random access
    - Note access at cluster level (not byte)
  - Sequential access means
    - Retrieval of data can in the order it is physically stored
    - Note data can only be sorted on one field

Files (Logical)

- File - File/File Systems are OS constructs
  - Anything stored on a storage medium, such as a program, document, digital image, or song
  - Files have a format in accord with their application

  ⇒ *Linux => only regular files and executable
- Filename
  - Name given to a file by the user (or program, eg. a.out)
  - Ext may/may not be associated with application
    - (OS managed)
- Folder (directories)
  - Named place on storage medium
  - Store files

Storage System Characteristics

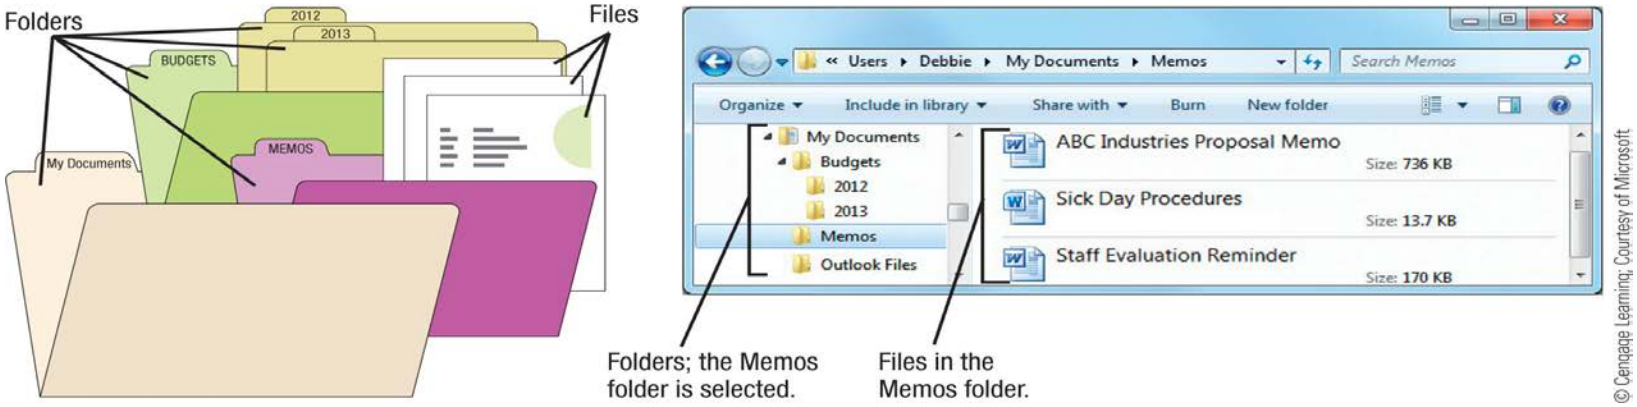

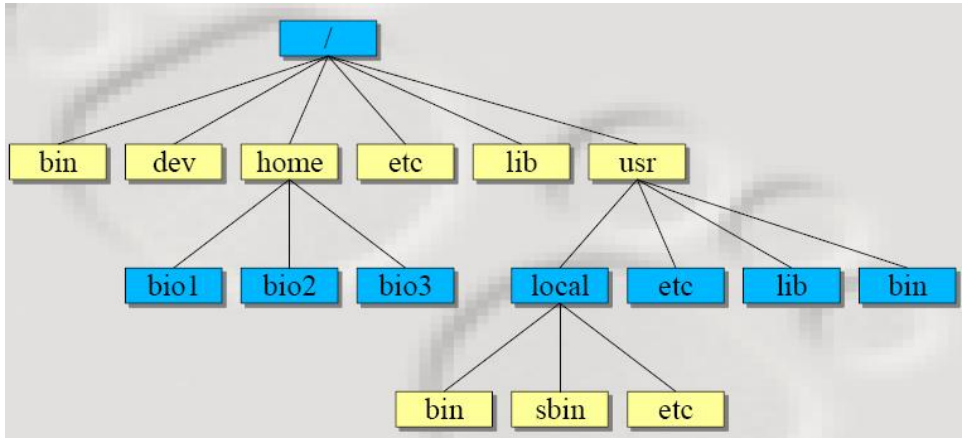

FIGURE 3-2
Organizing data.
Folders are used to organize related items on a storage medium.

Storage System Characteristics: Medium

- A storage system consists of a storage medium and a storage device
  - The **storage medium** is hardware where data stored
    - Magnetic, Electronic, Optical
    - Hard Drive, Flash SSD, DVD
  - The **storage device** is the hardware into which the storage medium is inserted
    - DVD drive, flash memory card reader, etc.
    - Can be internal, external, or remote
    - Storage devices are typically identified by letter
  - Some storage media removable; some not (integrated)

Devices/Mediums

1. Magnetic – Hard Drive (HD)
2. Electronic – Solid State Device (SSD)
3. Optical – Digital Video Disk (DVD)
4. *** DNA - Harvard

Each storage alternative involves trade-offs

Characteristics to consider:

- ⇒Speed
- ⇒Compatibility/convenience
- ⇒Capacity
- ⇒Portability
- ⇒(Again cost)

Hard Drives & Interface Standards

- Hard Drive
  - Stores most programs and data
  - Internal or external
  - Encrypted
- Hardware Interface standards => how drive connects to computer
  - Serial ATA (SATA): most common internal hard drive interface standard
    - eSATA faster than USB
  - Serial attached SCSI (SAS)
  - Fibre Channel
  - Internet SCSI (iSCSI)

Courtesy of Apricorn

FIGURE 3-3
Encrypted hard drives. The data stored on this external hard drive is protected by a fingerprint scanner.

Magnetic Hard Drives (cont'd)

– One/more permanently sealed metal magnetic disks with an access mechanism and read/write heads

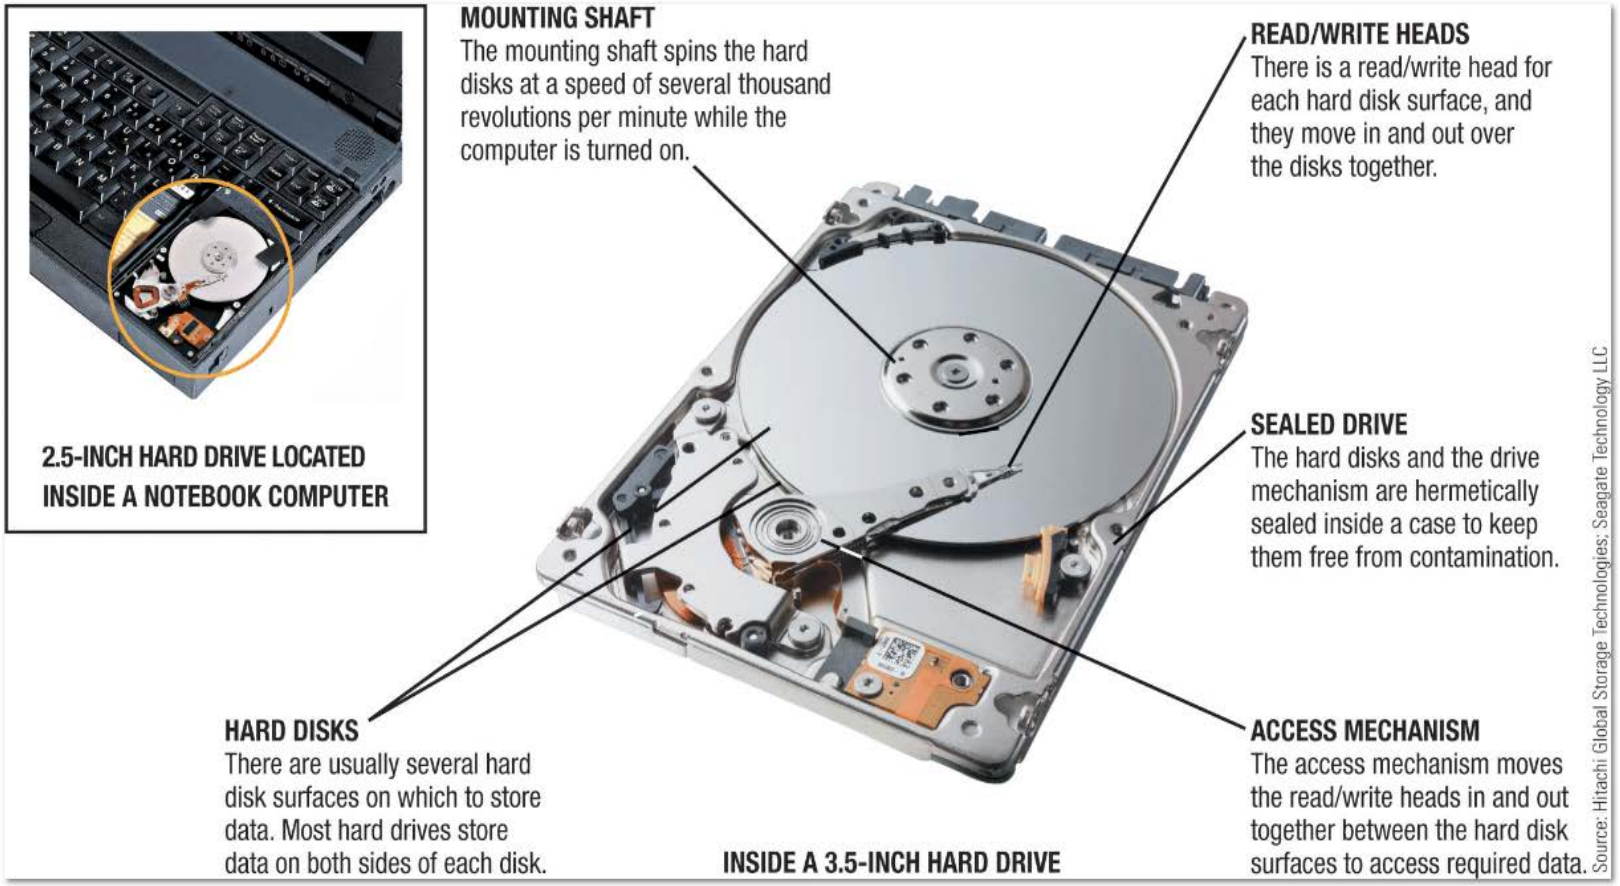

1. Magnetic Hard Drives

- A **magnetic hard drive** or **hard disk drive (HDD)** contains particles on the metal disks inside the drive that are magnetized to represent the data's 0s and 1s

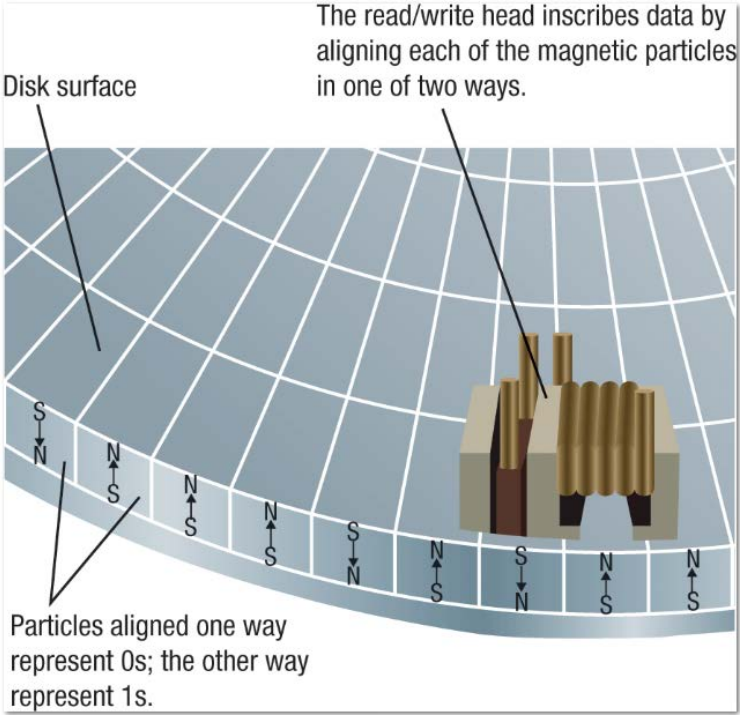

FIGURE 3-4
Storing data on magnetic disks.

Hard Disk Organization

- **Tracks** are concentric paths on the disk where data is recorded
- **Sectors** are small pieces of a track
- **Clusters** consist of one or more sectors
  - Smallest addressable area of a disk
- **Cylinders** are a collection of tracks located in the same location on a set of hard disk surfaces

Examples of Tracks, Sectors, Clusters, and Cylinders

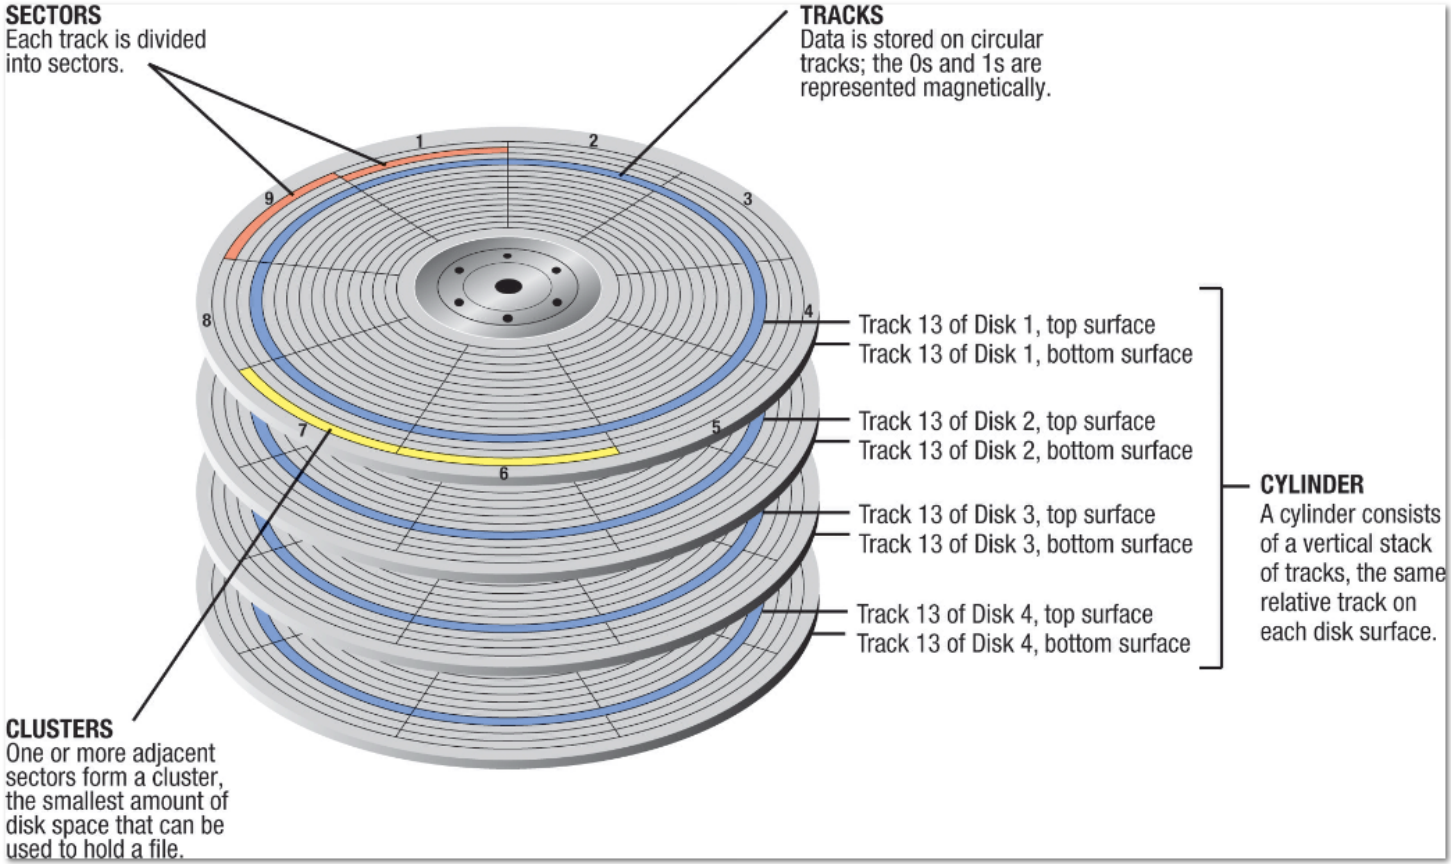

FIGURE 3-6
Magnetic hard disks are organized into tracks, sectors, clusters, and cylinders.

Hard Drive Speed

- **Disk access time** is the total time that it takes for a hard drive to read or write data

  => Recall all processing must be done in memory

  - Access time consists of seek time, rotational delay, and data movement time
  - SSDs don't require seek time or rotational delays

  – Mechanical Latencies
    - **Seek time** is read/write head movement to desired track
      – ~ 3 - 12 ms
    - **Rotational latency** is time incurred waiting for sector to rotate under r/w head
      – 10000 rpm ~ 3 ms
    - Transfer rate (disk to buffer) depends on track location (outside are faster)
      – ~1000 Mb/s
    - Consider mechanical speeds are ms whereas processing is MHz or GHz

CPU/HD Discrepancy (Memory/Storage Hierarchy)

Computer speeds are esoteric to humans(nano seconds). To put the Memory Hierarchy and CPU/Disc discrepancy in context – Lets treat CPU machine cycle as second

- CPU = nanosecond (.000000001)
- HD = milliseconds (.001)
- Difference in scale is 1,000,000
- So what is this difference if clock cycle was second?
- 1000000/60/60/24 => 11.5 days

Improve => Disk Cache

Disk Cache

⇒**Disk cache**

⇒Memory used in conjunction with HD => improve performance

- Typically RAM-based disk cache inside hard drive case
- Speed up performance
- Save battery life

**Takes advantage of locality of reference

Hard Drives

- File system (OS) determines(logical constructs)
  - partition size
  - cluster size
  - maximum drive size

Hard Drive Partitioning

- Partitioning divides the physical capacity of a single drive logically into separate areas, called partitions
  - Each partition functions as an independent hard drive
  - Referred as logical drives
  - Increases efficiency (smaller drives use smaller clusters)
- Partitions are used to create:
  - A recovery partition
  - A new logical drive for data
  - A dual boot system

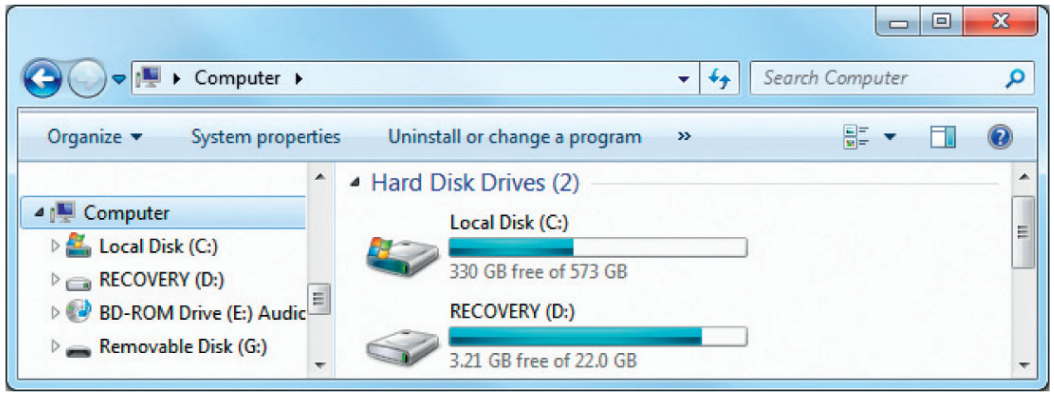

Courtesy of Microsoft

Hard Drive File Systems and Interface Standards

- File system determines the partition size, cluster size, maximum drive size, and maximum file size
  - Windows => FAT, FAT32, and NTFS
  - Mac OSX => Hierarchical File System Plus (HFS+)
    - extended and "Journaled" hard drive
    - Often compresses parts of OS X
  - Linux OS => EXT4
- Hardware Interface standards => how drive connects to computer
  - Serial ATA (SATA): most common internal hard drive interface standard
  - Serial attached SCSI (SAS)
  - Fibre Channel
  - Internet SCSI (iSCSI)

Another Categorization
Internal and External Hard Drives

- Internal hard drives are permanent storage devices located inside the system unit
  - Removed only if a problem develops
- External hard drives transport large amounts of data from one computer to another, for backup, and for additional storage
  - Full-sized external hard drives are often used for backup
  - Portable external hard drives: smaller and easier to transport
  - Most connect with a USB connection, although some may connect through wired or wireless networking connections

2. Flash Memory Storage Systems

- **Flash memory** is a chip-based storage medium that represents data using electrons
  - Nand vs. Nor
    - Nand block accessible (fast write-inexpensive)
    - Nor byte accessible (executable, fast read-expensive)
- **Embedded flash memory** refers to flash memory chips embedded into products
  - Smartphones, tablets, smart watches, and even sunglasses and wristwatches
  - Usually the primary storage for mobile devices such as tablets and smartphones

Flash Memory Cards and Readers

- A **flash memory card** is a small card containing one or more flash memory chips, a controller chip, and metal contacts to connect the card to the device or reader being used
  - Available in variety of formats (not interchangeable)
  - Secure Digital (SD) is one of the most widely used types of flash memory media
  - Most common type of storage media for digital cameras, smartphones, and other portable devices
- Many devices today have a built-in flash memory card reader; an external reader via USB port is also used
- Adapters allow the use of smaller flash memory cards in a larger slot of the same type (microSD to SD, etc.)

USB Flash Drives

- **USB flash drives** (USB drives or flash drives) consist of flash memory media integrated into a self-contained unit that plugs into and is powered by a USB port
  - Designed to be very small and very portable
  - Available in a host of formats
    - Low-profile drives, custom shapes, micro drives, etc.
  - Can be built into a consumer product
  - Additional related hardware becoming available
    - USB duplicator systems

Smart Cards

- A **smart card** is a credit card-sized piece of plastic that contains some computer circuitry (processor, memory, and storage)
  - Stores a small amount of data (about 64 KB or less)
  - Commonly used to store prepaid amounts of digital cash or personal information
  - Smart card readers are built into or attached to a computer, door lock, vending machine, or other device
  - Some smart cards store biometric data
  - Use of mobile smart cards is an emerging trend

Solid-State Drives (SSDs)

- A **solid-state drive (SSD)** uses flash memory technology to store data
  - Uses less power and has no moving parts
  - Much faster than magnetic hard drives, but more expensive
  - The norm for netbooks, mobile devices, and other portable devices

FIGURE 3-7
Solid-state drives (SSDs). Contain only flash memory.

Solid-State Hybrid Drives (SSHDs)

- A **solid-state hybrid drive (SSHD)** or **hybrid drive** uses a combination of magnetic disks and flash memory chips
  - The data that is most directly associated with performance is stored in the flash memory
  - Nearly as fast as solid-state drives (SSDs)
  - Slightly more expensive than magnetic hard disk drives (HDDs)

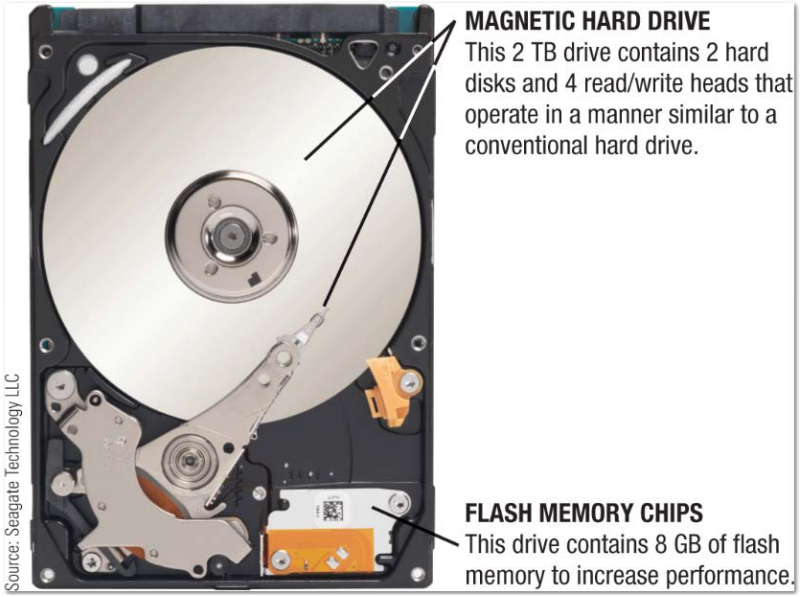

FIGURE 3-8
Solid-state hybrid drives (SSHDs). Contain both magnetic hard disks and flash memory.

Technology and You

Tiny PCs

- The size of a USB flash drive
- Typically connect to a TV via an HDMI port
- May also have built-in storage and a microSD
- Capabilities vary; smart TVs can display and stream Internet content
- The newest tiny PCs are fully functioning "computers-on-a-stick"

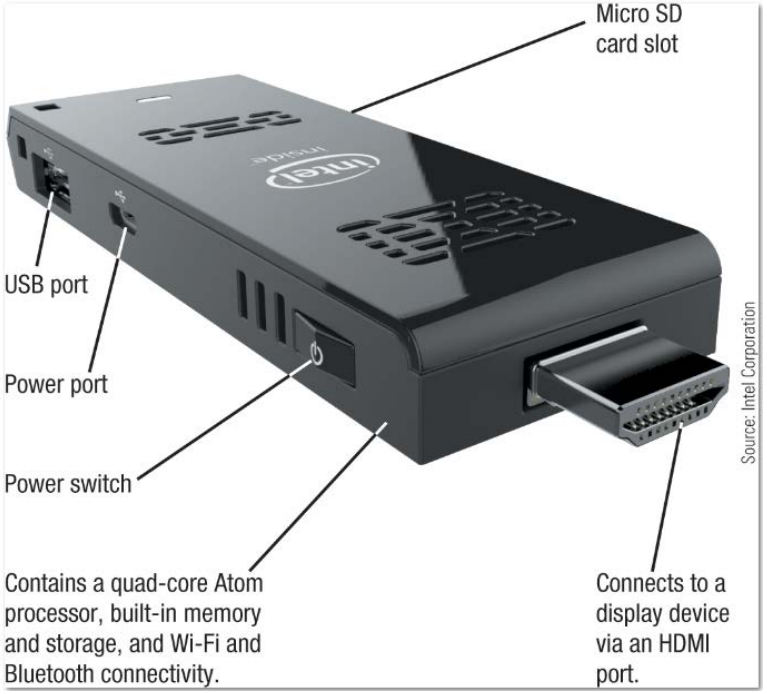

Intel Compute Stick

3. Optical Drives

- Optical discs are read by **optical drives**
  - Optical drive must support the type of optical disc being used
  - Almost always backward-compatible
  - Recording data onto a optical disc is called burning; requires burning software
  - Optical drives can be internal or external
    - External drives typically connect via USB port
    - External drives can be used with netbooks and other devices without an optical drive

FIGURE 3-12
External optical drives. Can be connected as needed, typically via a USB port, such as to the netbook shown here.

Optical Discs

- **Optical discs** are thin circular plastic discs
  - Are read from and written to using laser beams
  - Divided into sectors like magnetic discs but use a single spiral track (groove)
  - Have a relatively large capacity and are durable
  - Used for backup purposes and for storing and transporting music, photos, video, etc.
  - CDs, DVDs, RW vs R, etc

Representing Data on an Optical Disc

- Pits and lands are used to represent 1s and 0s
- The transition between a pit and a land represents a 1; no transition represents a 0
- Read-only optical disc
  - Surface of disc is molded or stamped to represent data
- Recordable or rewritable disc
  - The reflectivity of the disc is changed using a laser beam to represent the data
  - Different types of optical discs use different types of laser beams

How Recorded Optical Discs Work

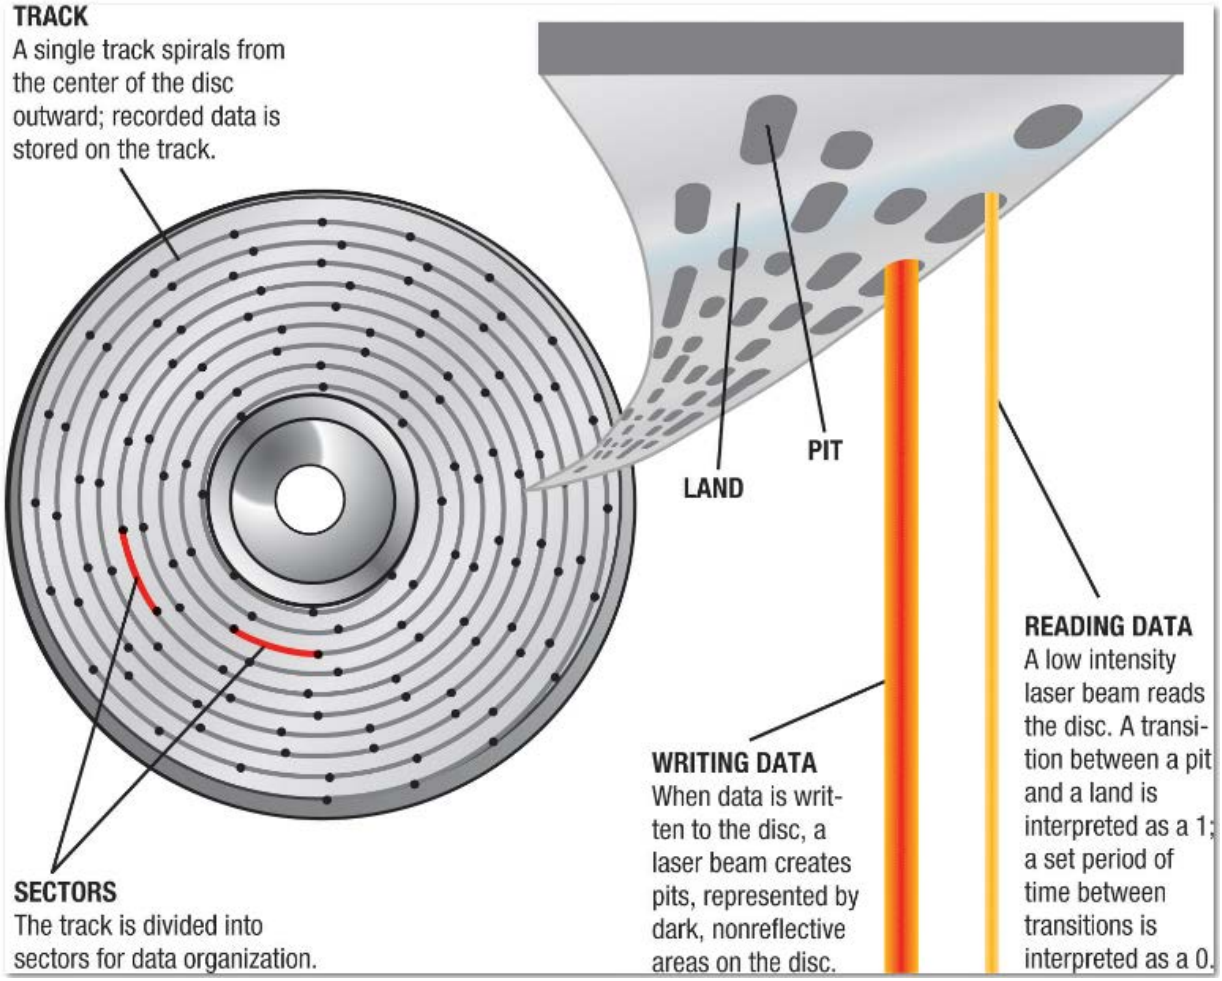

FIGURE 3-11
How recorded optical discs work.

Other Types of Storage Systems

- Remote storage refers to using a storage device that is not connected directly to the user's computer
- **Network storage:** Using a storage device via a local network
  - Works in much the same way as using local storage
  - **Network attached storage (NAS)** devices are high performance storage systems connected individually to a network
  - A **storage area network (SAN)** consist of separate network of hard drives or other storage devices that are attached to the main network

Cloud Storage

- **Cloud storage** (**online storage**) is accessed via the Internet
  - Cloud applications (Flickr, Facebook, Google Docs, etc.)
  - Online storage sites (Box, Dropbox, OneDrive, etc.)
  - Growing in importance because more and more applications are Web-based
  - Increasingly used for backup purposes
  - Files can be synched between PC and cloud storage
  - Many online storage sites offer some free storage
  - Business cloud storage is available; businesses can also create private clouds

RAID

- **RAID** (**redundant arrays of independent discs**) is a method of storing data on two or more hard drives that work together to record redundant copies
  - Used to protect critical data on large storage systems
  - Helps to increase fault tolerance
  - Different levels of RAID:
    - RAID 0 = disk striping (spread files over two or more hard drives)
    - RAID 1 = disk mirroring (duplicate copy)
    - Other levels use a combination or striping and mirror

Two Primary RAID Techniques

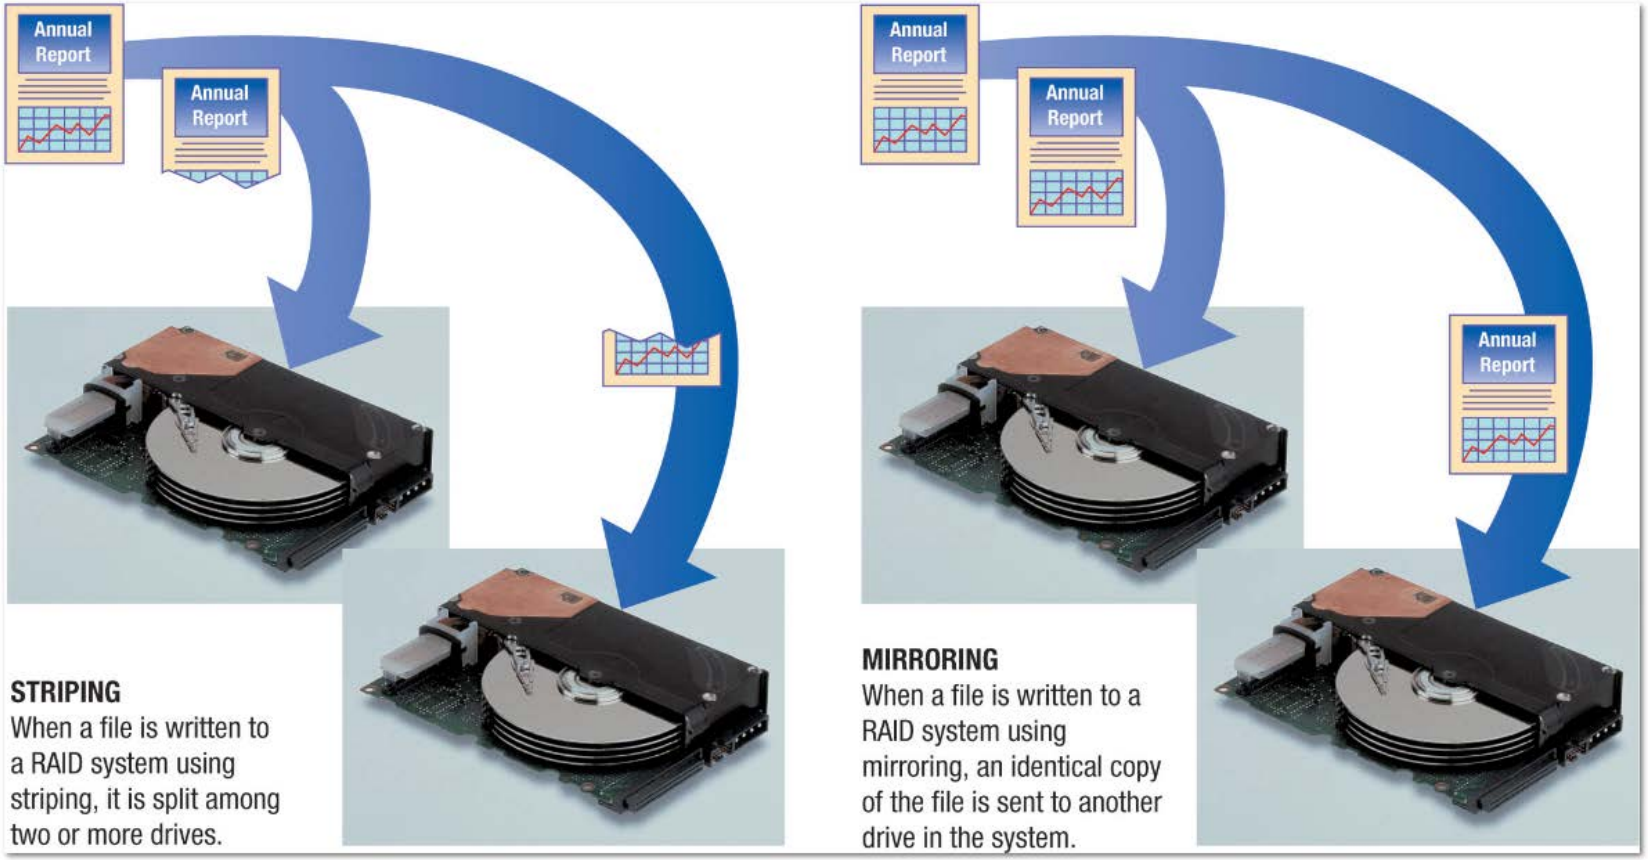

FIGURE 3-23
RAID. Two primary RAID techniques are striping and mirroring.

Archival Storage Systems

- Data archiving is the process of identifying and moving data that is no longer actively being used from regular storage systems to a separate long-term archival storage system
- Options for data archival systems:
  - Large hard drives, such as a helium hard drive (10 TB)
  - **Magnetic tape**
    - Typically cartridge tapes; can be tape libraries
  - Higher capacity, archival Blu-ray Discs that are becoming available; so are optical jukeboxes
  - Cloud storage

Magnetic Tape

- Magnetic Tape Systems
  - Plastic tape with a magnetizable surface that stores data as a series of magnetic spots
  - Primarily used for backup and archival purposes
  - Sequential access only
  - Low cost per megabyte
  - Most tapes today are in the form of cartridge tapes

Courtesy of Imation

Backup (OS functionality)

Each type begins with Full backup – differ in successive backups

- Full (or Reference) Backup
  - Backs up entire/selected files
- Differential Backup
  - Provides a backup of all files that have changed since a full backup was performed
    - Backups continually increase in size
- Incremental Backup
  - Provides a backup of files that have changed or are new since the last incremental backup.
- Now - What are the performance issues of each backup and restore?

Evaluating Your Storage Alternatives

- Product Characteristics to Consider:
  - Speed
  - Compatibility
  - Storage capacity
  - Convenience
  - Portability
  - (Cost)
- Most Users Require:
  - Hard drive
  - Recordable or rewritable optical drive
  - Flash memory card reader
  - USB port(s)