JOHN ABBOTT
CEGEP/COLLEGE

Technical Support (420-1N6-AB)

Hardware – Servers and Fault Tolerance

Fall 2023

Adapted for Fall 2025

Outline

- What is a Server?
- Server Hardware
- Fault Tolerance
  - Redundancy
  - Data loss
    - RAID
    - Advanced file systems

What Is A Server?

- A server is a computer or device on a network that manages resources or provides functionality.
  - To private users inside an organization
  - Or to public users via the Internet.
- Servers can provide various functionalities often referred as “Services”. Such as:
  - Sharing data.
  - Sharing/Hosting applications.
  - Sharing hardware.
  - Performing computation.

Client/Server Model

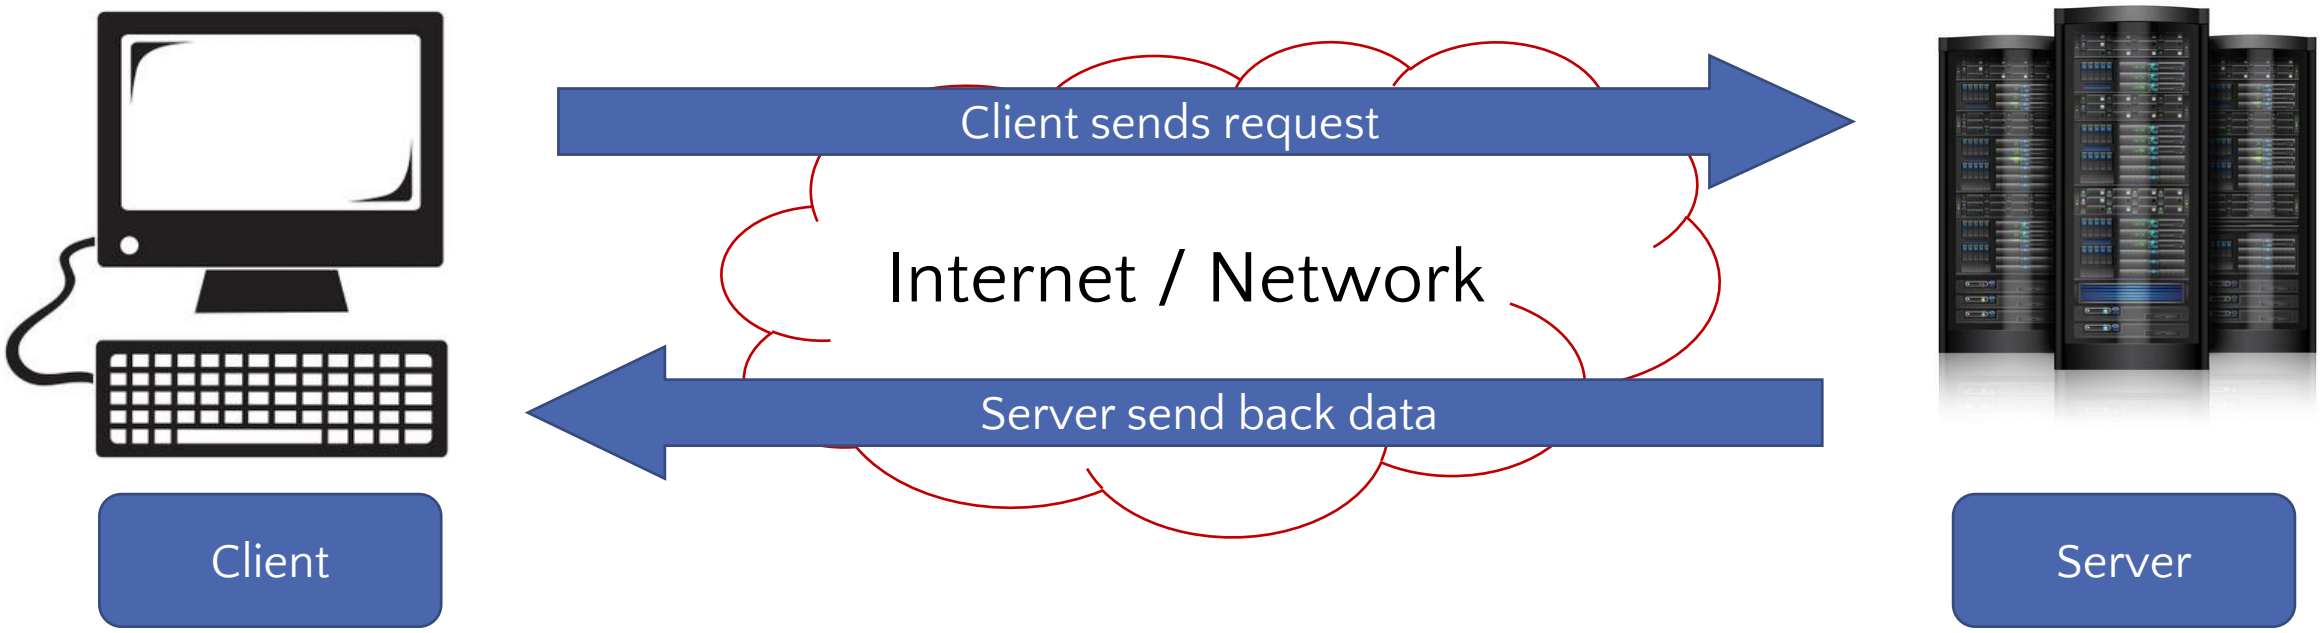

Examples of Servers

- File server
  - Stores and distributes files.
- Print server
  - Allows the management and distribution of printing functionality.
- Hypervisor server
  - Hosts virtual machine guest operating systems for remote/shared access
- Application server
  - Runs applications that are shared by a large number of users.
  - Clients do not need to have enough resources nor have to maintain the applications.
- Mail server
  - Transfers and stores mail over corporate networks and across the Internet.
- Web server
  - Hosts webpages and provides access to them across a network to the user's web browser.

Server Hardware

Differences Between Servers and Workstations

- Servers perform a service (or services) for clients that are connected.
  - They may provide access to shared files, printers, other network resources, or may provide email and web hosting services.
- Workstations are personal computers.
  - They are used for applications such as word processing, graphical design, development, browsing the web, etc.

Server Hardware Components

- In many ways server hardware really isn't all that different from a workstation PC hardware.

- Servers use the same basic components such as:
  - CPUs.
  - Main memory.
  - Secondary storage.
  - Power supplies.

Server Hardware

- Servers are
  - Noisy
  - Power hungry
  - Need good Internet access
  - Require increased security.

- Servers are usually stored in dedicated *server centers* (or *data centers*).

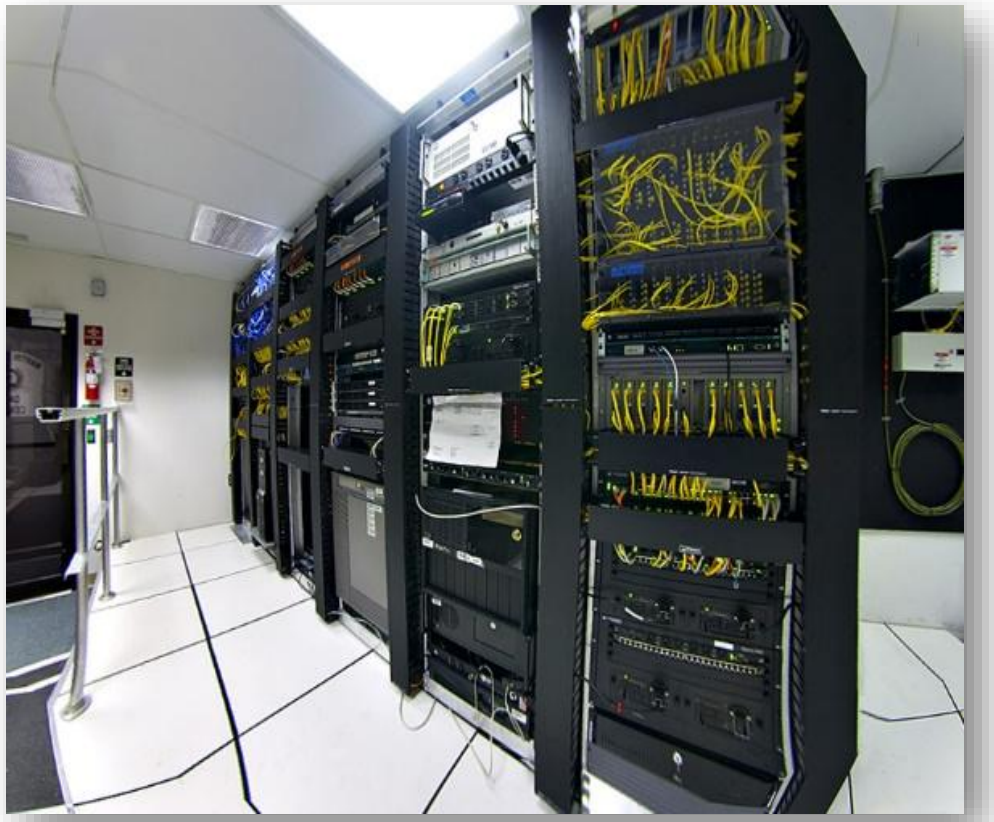

Types of Servers

- (Form Factors)

Tower Servers	Rack Servers	Blade Servers
		

Tower Servers

- Look like PCs.
- Each tower server is a standalone machine.
- Occupy space and tent to be noisy.
- Usually used in small data centers.

Rack Servers

- Servers mounted within a rack.
- Uniform width, various heights.
  - Server heights is measured using rack units.
  - A standard rack mount server is referred to as a 1U server meaning that it is 1 rack unit in size.
- Each rack can accommodate multiple servers which are typically stacked one on top of the other.
- Mounted to rack using screws.

Blade Servers

- Mount inside a special rack (chassis).
- Blade servers tend to be vendor proprietary.
  - You can't for example insert a Dell blade server into an HP chassis.
- Lack certain components that rack servers have, e.g. power supplies.
- The chassis can include a power supply, cooling unit, etc.

Server Models

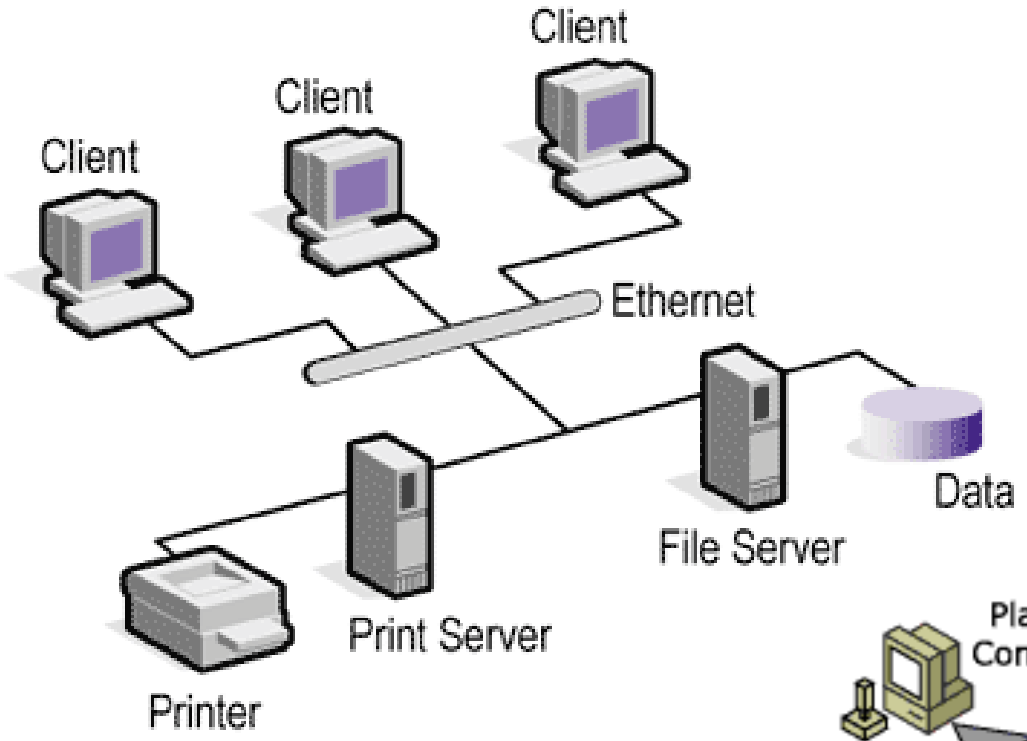

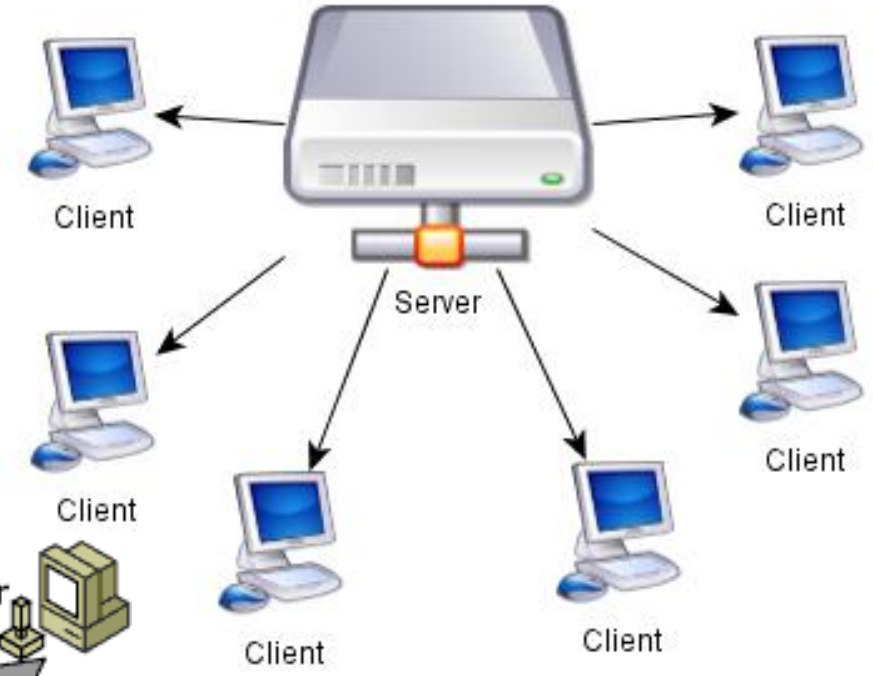

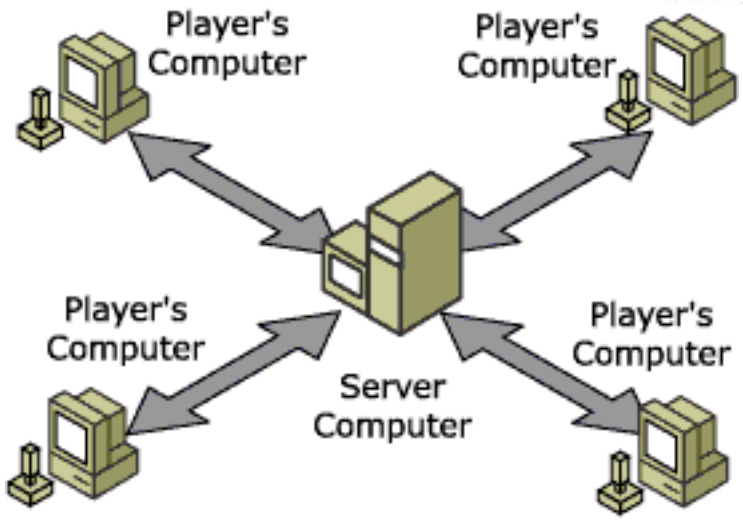

Server Hardware vs Computer Hardware

- In the hardware sense, the word *server* means hardware intended for hosting software applications under *intense demand*.
- Most PCs can act as a server, but a dedicated server will contain features making it more suitable for production environments where *performance* and *resilience* are important.
- Server hardware components are much more expensive.
  - There is a tradeoff between performance and expense.

Server Hardware Performance Features

- Advanced CPUs
  - More cores
  - Extra cache
  - Hyperthreading
- Server hard drives
  - Resistant to vibrations.
  - Large is size.
  - Fast access time.

Fault Tolerance

- **Fault tolerance,** or **redundancy** is the property that enables a system to continue operating properly in the event of the failure of one (or more) of its components.
- Fault tolerance is crucial in servers' environments.
- Variety of fault tolerance mechanisms:
  - Redundant files, volumes or drives.
  - Redundant hardware components or servers.
  - Hardware components with recovery options.

Examples of Fault Tolerance

Power: Redundant power supply (RPS) / Uninterruptible power supply (UPS)

- RPS: using more than one power supply to provide redundancy in case one fails
- UPS: A kind of power bar that contains a battery (see image to the right). Allows low power operations to continue running and high power operations to shut down gracefully in the event of a power outage/brown-out

Examples of Fault Tolerance

Network: Redundant network interface cards.

- Attaching extra NIC (network interface cards) to servers allows duplication/parallization of network connections
- Improvement in performance (Combined speed) and provides fault tolerance

Examples of Fault Tolerance

- **Memory:** Memory Modules with error checking (ECC).
  - RAM with ECC support automatically uses parity (like Assignment 1) to detect and correct the most-common kinds of internal data corruption.
  - NOTE: not all motherboards are compatible with ECC RAM (consumer-grade motherboards often cheap-out on "non essential" features like this -- it is controversial!)

- **Storage:** RAID (Redundant Array of Independent Disks)
  - Configure a group of storage disks to **stripe, mirror**, and/or **apply parity** to data in storage.
  - Each of these actions either improves performance or provides fault tolerance
  - The next set of slides explains all of the possible RAID configurations

ECC RAM contains an extra chip to store parity information for error correction

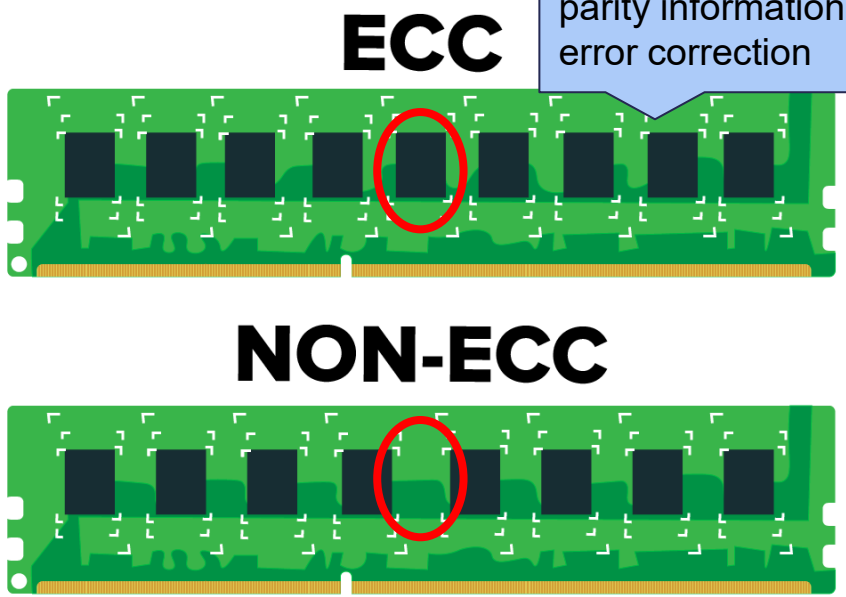

Examples of Fault Tolerance

Storage: RAID (Redundant Array of Independent Disks)

- Configure a group of storage disks to **stripe, mirror**, and/or **apply parity** to data in storage.
- Each of these actions either improves performance or provides fault tolerance
- The next set of slides explains all of the possible RAID configurations

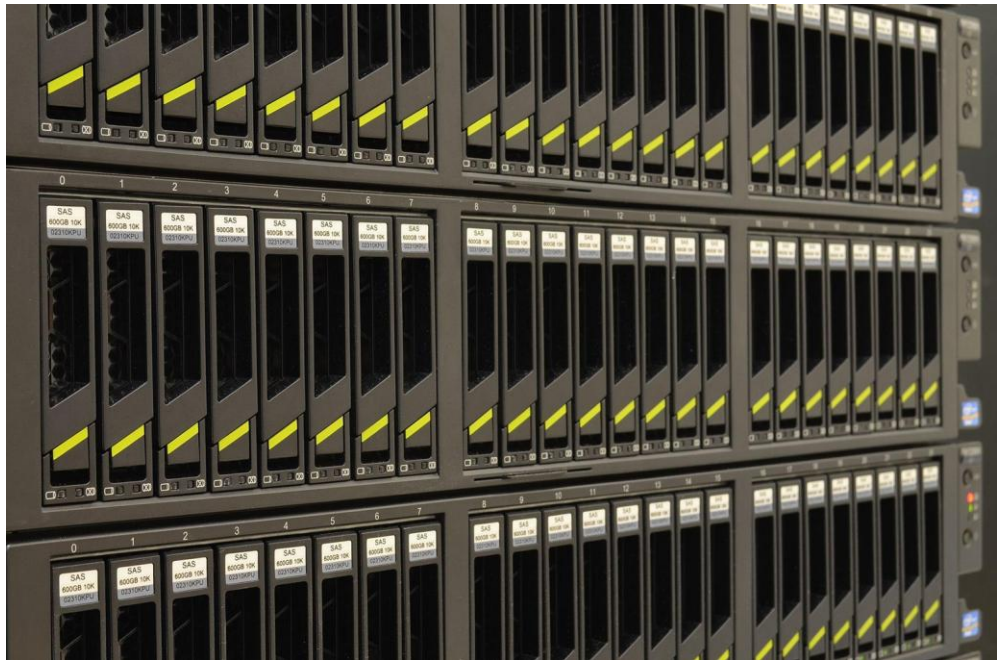

File storage servers: each contain 24 hard disk drives. Configuring them with RAID gives **storage redundancy**. Image from Wikipedia.

NOT RAID: JBOD
"Just a bunch of disks" (No Fault Tolerance)

What's the simplest way to combine a storage drives? That is **JBOD**, Just a Bunch of Disks.

With JBOD, data is **spanned** across all available disks

- Nothing special is done to read or write data: wherever there is space on a disk, data can be written.
- No fault tolerance (no parity bits or data mirroring)
- Combined storage:
  - This is the simplest configuration for **combining the storage of many disks** to appear as one volume.
  - Each disk can be any size; the **total storage is just the total size of the disks** (e.g.: 1TB + 2TB + 3TB = 6TB total storage)

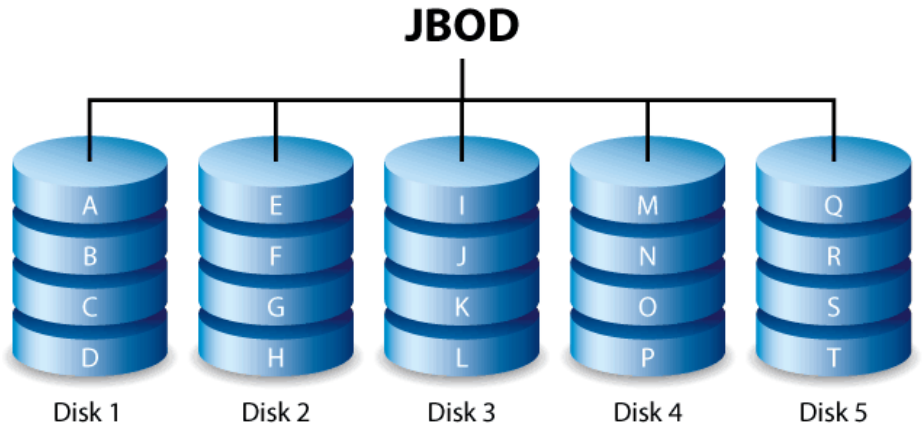

note: for disk graphs that look like this, you can consider that each letter is a different file block (A is one file block, B is another, etc.)

RAID 0: Striping for improved performance

RAID 0 is pretty similar to JBOD, with one key difference:

Data is **striped** across all available disks. When new data is written to a disk, it is **cut up into blocks and spread across all disks in the RAID configuration**. All files are written to all disks (see right-hand side: A is one file, A1/A2/A3/etc. are blocks containing file A data)

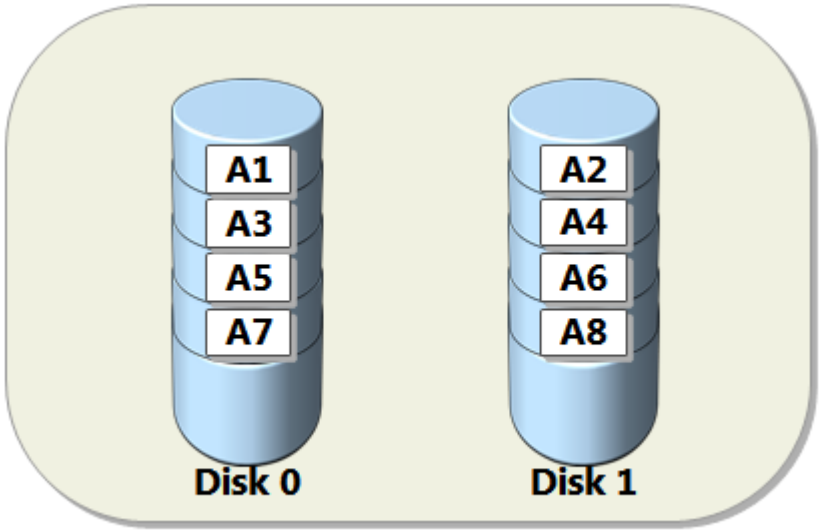

Properties:

- No fault tolerance (no parity bits or data mirroring)
- On HDD: better read/write speed compared to one-huge disk
  - more access points to data improves read/write speed

Combined storage:

- **total capacity = number of disks * size of smallest disk**
- **Why?** Data is striped across all disks -- when the smallest disk is full, no more data can be added. **In RAID, it is very common to make sure all disks are the same size to prevent wasting space.**
- https://www.seagate.com/ca/en/internal-hard-drives/raid-calculator/ nice tool for visualizing RAID storage capacity

RAID 1: Mirroring for fault tolerance

- Disk mirroring where data is copied onto more than one drive simultaneously
  - (NOT striped: **copied**: note that the file A is divided into four parts, and all four parts are present on both disks)
- Simplest and cheapest way to implement **fault tolerance.**
- Minimum of 2 disks required: each additional disk becomes “another mirror”
  - additional disks provide more protection, but do not provide more storage
- Combined storage:
  - **Total capacity = size of smallest disk**
  - RAID 1 is almost exclusively used with just two disks and no more

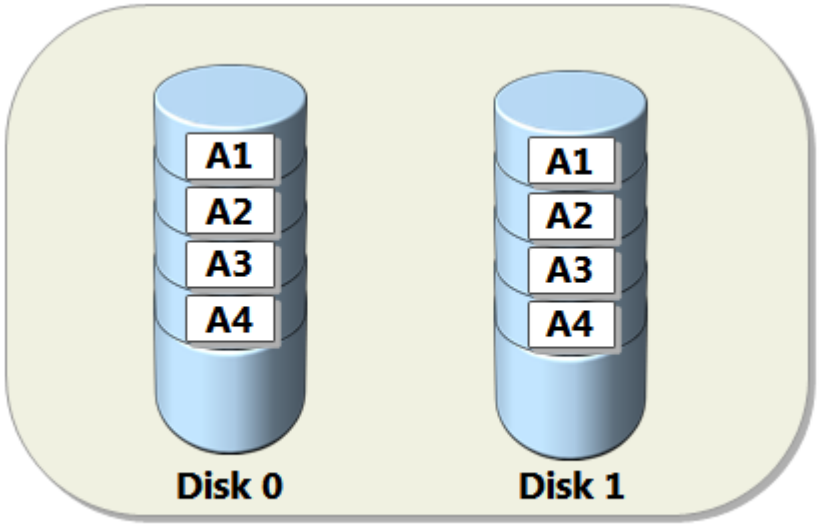

RAID 5: Striping + 1 disk of parity

- A common RAID configuration for business servers.
- Data and parity (additional data used for recovery) is striped across 3 or more disks, with the equivalent of 1 disk of storage being dedicated to storing parity data.
- **If any one disk fails**, data can be re-created automatically using the parity bits across the remaining disks. Once any more than one disks fails, data is lost permanently.
- Great solution for fault-tolerance, but slows down server performance
- Combined storage:
  - Total capacity = (number of disks - 1) * size of smallest disk
  - Why? One disk is dedicated to parity, the rest are dedicated to storage

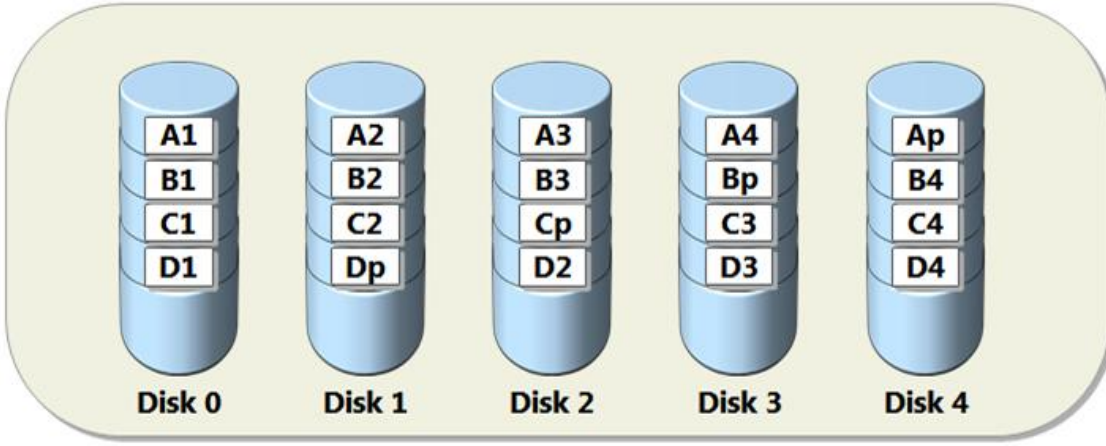

RAID 6: Striping + 2 disks of parity

- Identical to RAID 5 but uses one more parity block
  - more robust: two disks can die
  - less storage:
    - Total capacity = (number of disks - 2) * size of smallest disk

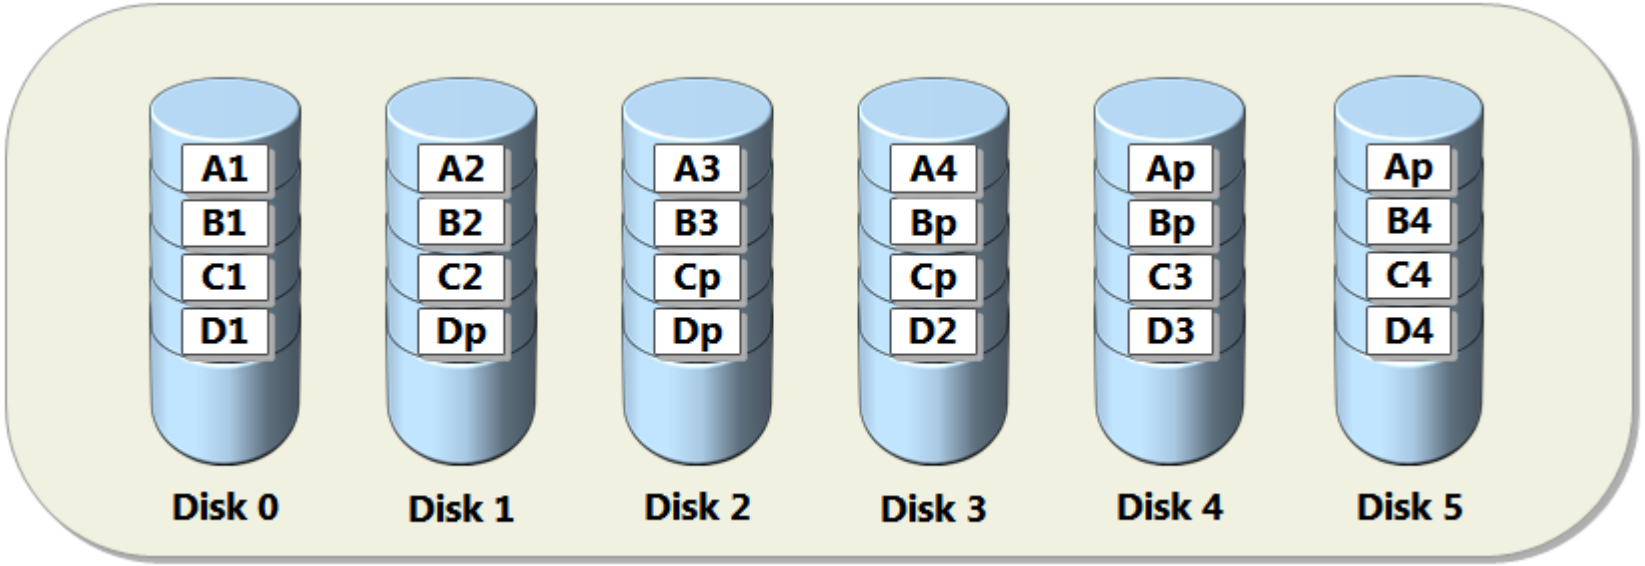

RAID 10 (1+0): Striping + Mirroring

- Combination of RAID 0 and 1, combines mirroring of 1 with striping of 0.
- Gives best performance but requires twice as many disks (minimum of 4 disks).

- Combined storage:
  - **Total capacity = (number of disks * size of smallest disk) / 2**
  - Why? Half of the disks are dedicated to mirroring data

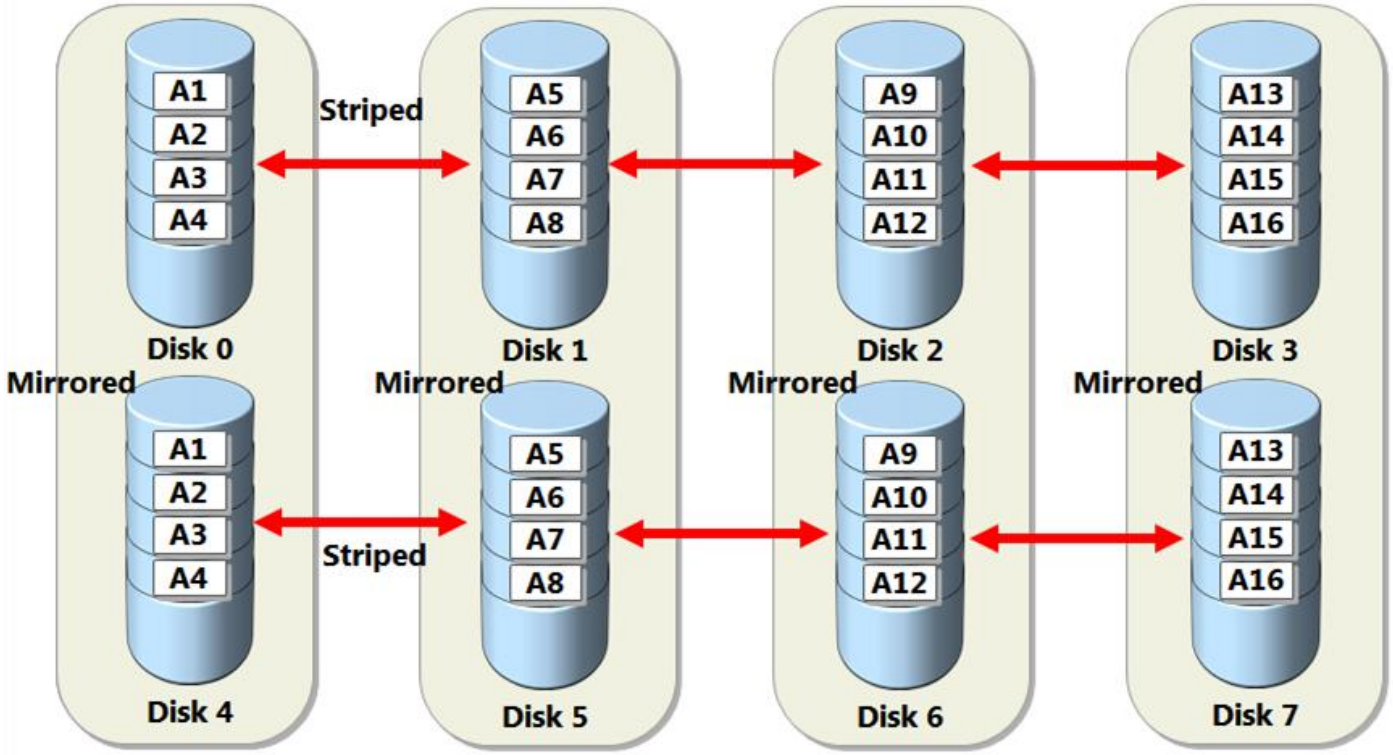

Technical Terms: Spanning

- **Disk spanning** is the process of combining multiple storage disks into one volume.
- JBOD is an example of disk spanning. When one disk is full, the next disk is used
- There is no extra efficiency or fault tolerance for spanning: spanning simply allows you to combine disks to create larger storage volumes.

Technical Terms: Striping

- **Disk striping** is the process of dividing data into blocks and spreading the data blocks across multiple storage devices.
  - Each file is therefore stored on ALL disks that are striped
  - This is very different from **spanning**!

- Different storage systems can stripe in different ways (e.g. at byte or block level).

- Striping provides high performance (faster reading and writing of data), but low resiliency (failure in any physical drive would result in data loss).

Technical Terms: Mirroring

- **Mirroring** involves writing the same data to more than one disk.
- It is a good choice for applications that require high performance and high availability such as transactional applications.
- Data can be read from both disks simultaneously, making read operations fast. However, write operations are slower because they have to be carried out twice.

Q & A