IST346

Data Storage

Data Storage

Why Data Storage?

- Information is a the center of all organizations.
- Organizations need to store data. Lots of it.
- What Kinds of Data?
  - Documents and Files
    - (Reports, Proposals, Letters, Presentations, etc.)
  - Databases
    - Customer Information, Orders, Inventory
  - Multimedia
    - Video, Audio
  - Email !!!!!

Before We Begin… Bits and Bytes

- One Byte is approximately a single character:
  `This is 16 bytes` (includes the spaces)
- 1 KB = 2^10 Bytes, Approx. 1000 Bytes
  10-15 sentences of text
- 1 MB = 2^20 Bytes, Approx. 1 Million Bytes
  Typical Song in MP3 format is about 3MB
- 1 GB = 2^30 Bytes, Approx. 1 Billion Bytes
  Mpeg4 Movie in 1080p is around 2GB
- 1 TB = 2^40 Bytes, Approx. 1 Trillion Bytes

Some Storage Stats for the iSchool (2009)

Service	Total Allocation	Average*
Email	405GB	1.4 GB / User*
Personal Files	238GB	1.8 GB / User*
Configuration Data	12GB	94 MB / User*
Shared File Space	202GB	870 MB / User*
Video	145GB	1.1 GB / User*
iLMS (iSchool Blackboard)	430GB	323 MB / User +
TOTAL	**> 1.4 TB**	**~ 1GB / User**

This ONLY represents the information the users generate and does not include operating system files, application files, etc.

*Based on 230 users
+Based on 1550 users

Another way to look at it.

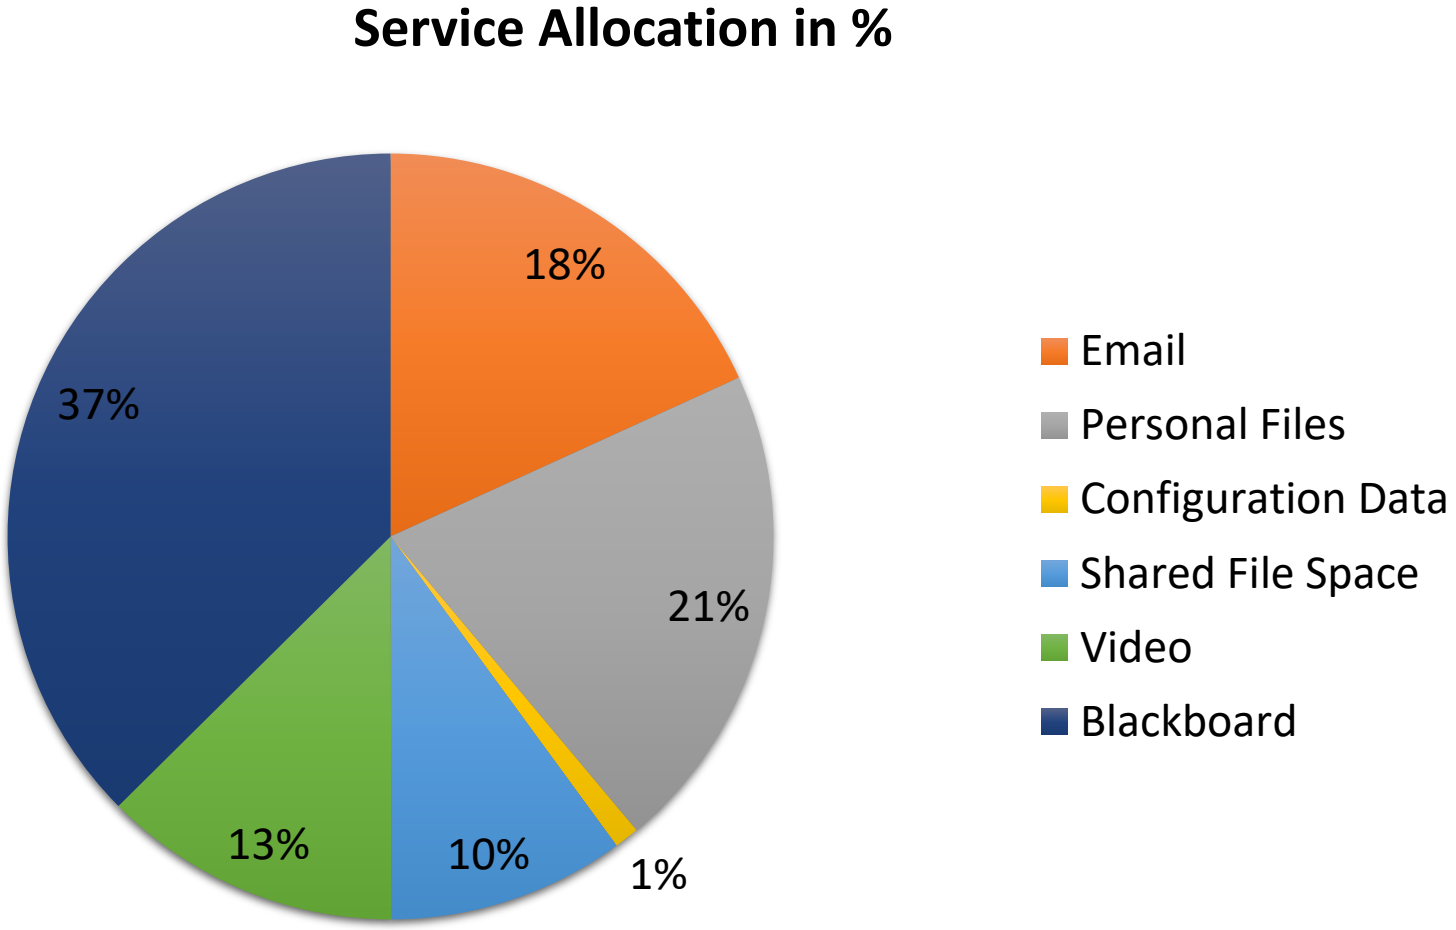

You might be thinking….

- 1GB per user is **NOTHING**, considering:

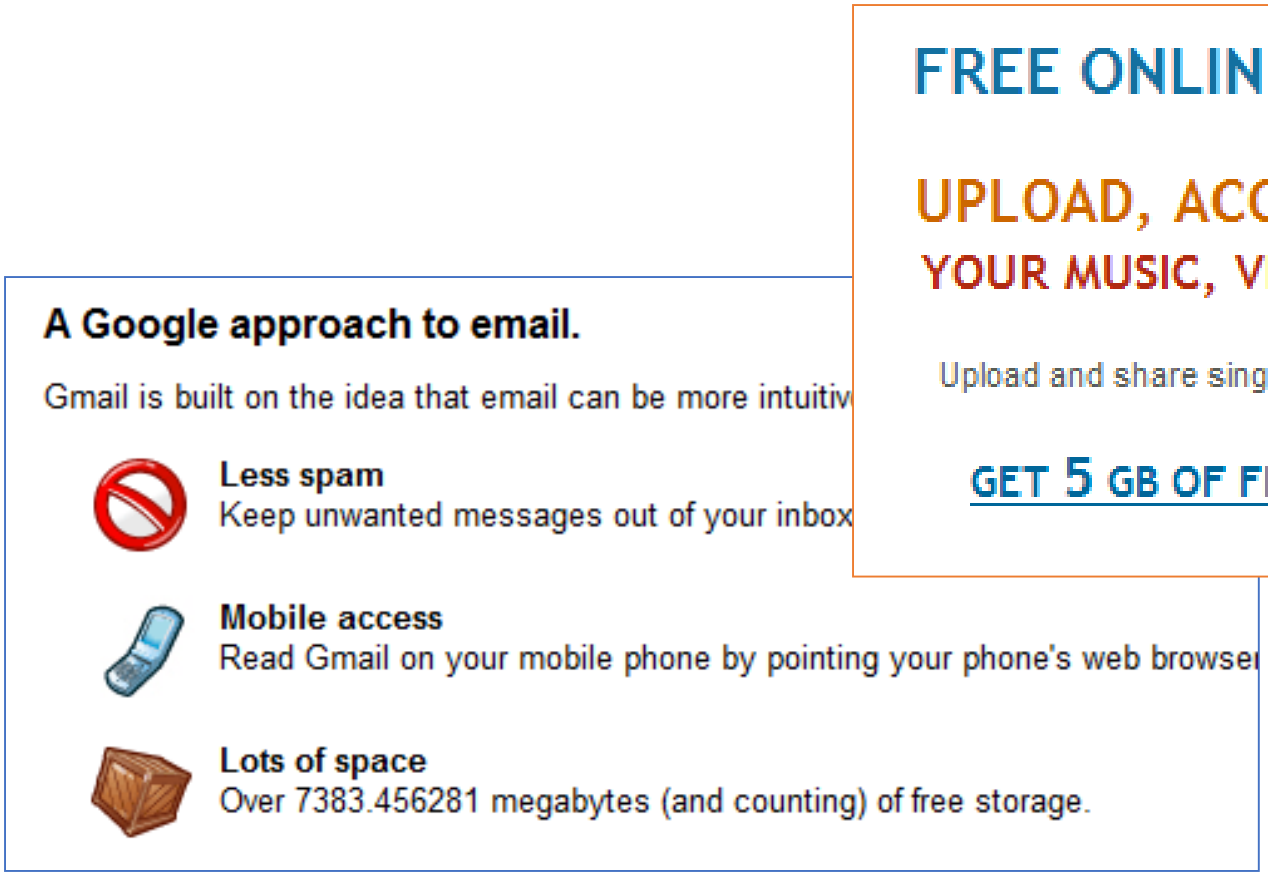

The Fallacy of Enterprise Storage

- For every 1GB of user storage (files, email) you need:

- An additional 1GB for every full backup you do.
- At most another 1GB for fault-tolerance.

- 1GB per user is now 8GB when part of:
  - RAID1 (Mirroring) for fault tolerance (1GB)
  - Weekly full off site backups (4GB)
  - With two months of backups on the shelf (2GB)

The Fallacy of Enterprise Storage

- Disk Drives are Cheap, Storage is not (a simple example)
- 2TB Drive = **$170**
- 8TB Array = **$3,280** (actual 6TB usable in RAID5)
- **19X more expensive, for only 3X the storage!**

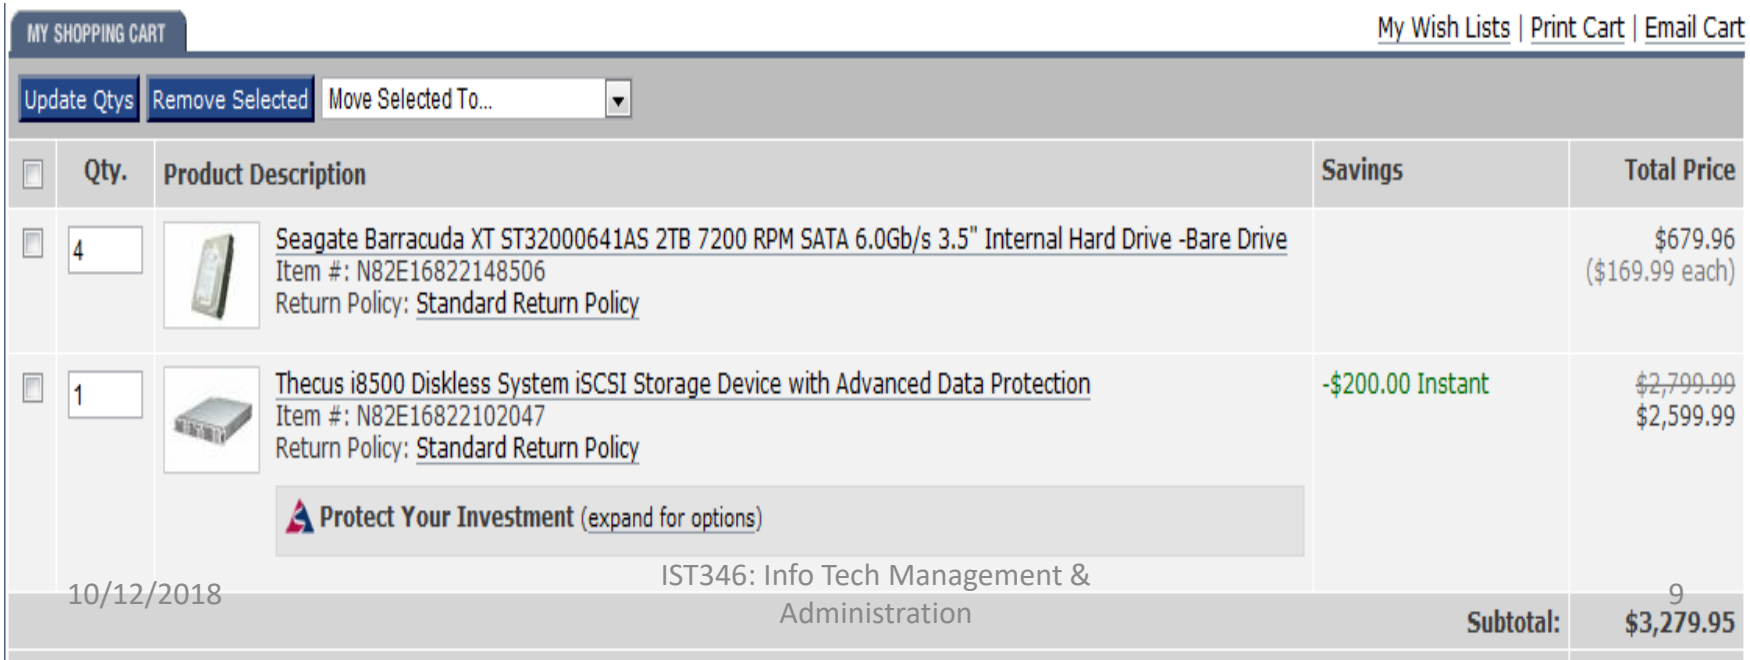

MY SHOPPING CART

My Wish Lists | Print Cart | Email Cart

Update Qtys | Remove Selected | Move Selected To...

	Qty.	Product Description	Savings	Total Price
	4	Seagate Barracuda XT ST32000641AS 2TB 7200 RPM SATA 6.0Gb/s 3.5" Internal Hard Drive -Bare Drive Item #: N82E16822148506 Return Policy: Standard Return Policy		$679.96 ($169.99 each)
	1	Thecus i8500 Diskless System iSCSI Storage Device with Advanced Data Protection Item #: N82E16822102047 Return Policy: Standard Return Policy Protect Your Investment (expand for options)	-$200.00 Instant	~~$2,799.99~~ $2,599.99
			Subtotal:	$3,279.95

Physical Disks

Disk Components

- Traditional disks are made of a few basic components
  - Spindle : common mounting platform for all platters.
  - Platters : disc type material where data is stored
  - Heads : Actual ‘arm’ that moves in and out of the record to read data.
  - Track : location on the platter where data is stored. Each track has the same radius from the spindle on a disk
  - Sector : pie-shaped slice of the platter. Used to sub-divide the platter for quickly locating blocks of data.
  - Cylinder : same tracks on multiple platters.

Disk components

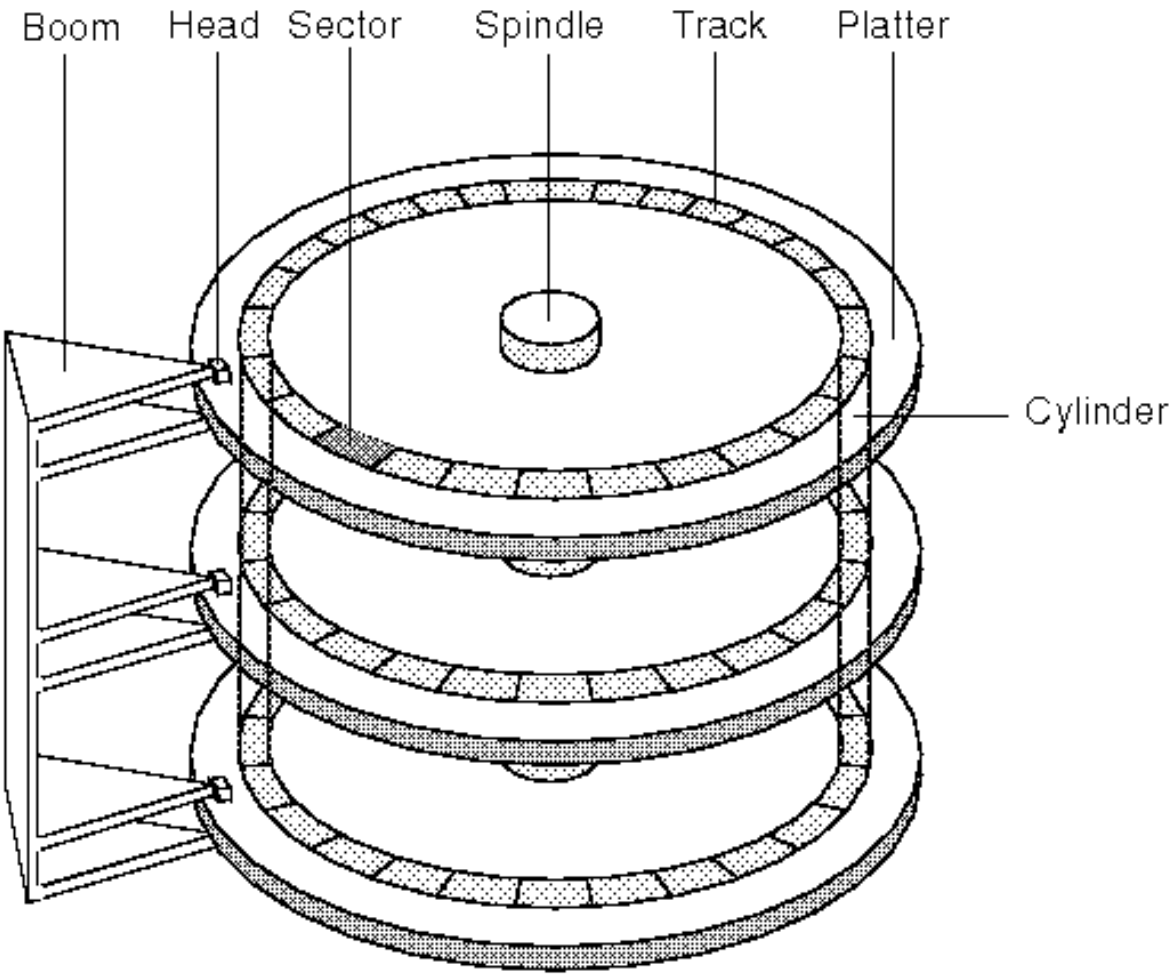

A new storage approach: SSD

- SSD: Solid State Disk/Drive
- Differs from traditional disks as there are no moving parts…no platters, no mechanics, so fewer failures.
- Utilize dram or eeprom memory boards to store bits
- Some use their own CPUs to manage data storage.
- Faster for accessing data than traditional disks since seek time is eliminated
- Much more expensive per GB (for now) than traditional spinning electromagnetic disks

The internals of an SSD

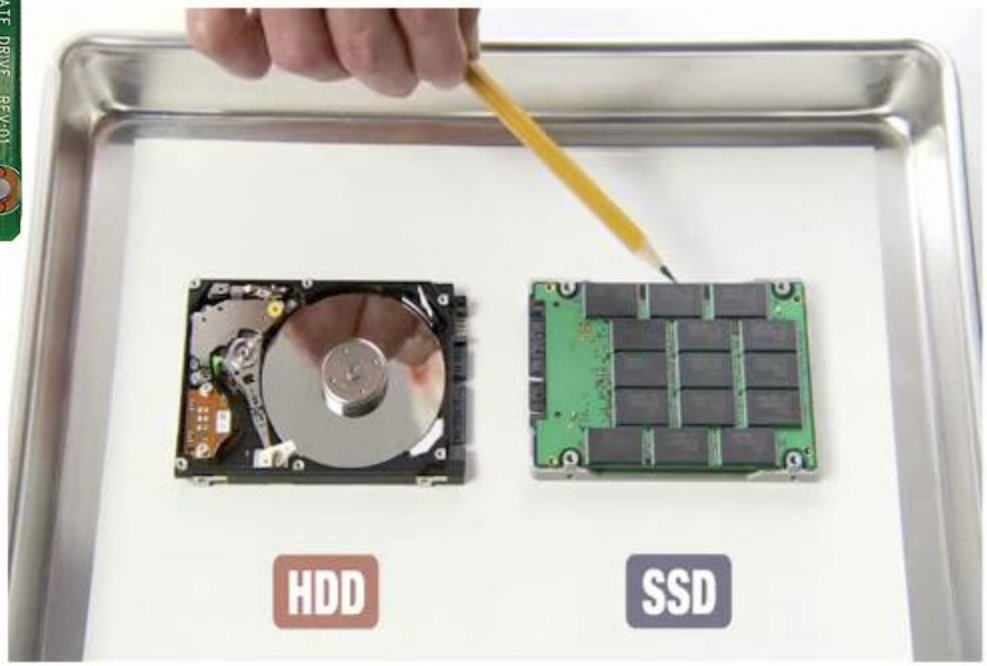

Logical Disks and Storage

Storage Terminology - Volumes

- **Volume** – the chunk of storage as seen by the server.
  - A single hard drive (or logical disk) can be one volume.
  - One hard drive (or LD) can be multiple volumes. We call this **partitioning** the disk.

  - Many hard disks can be combined into one **logical volume**.

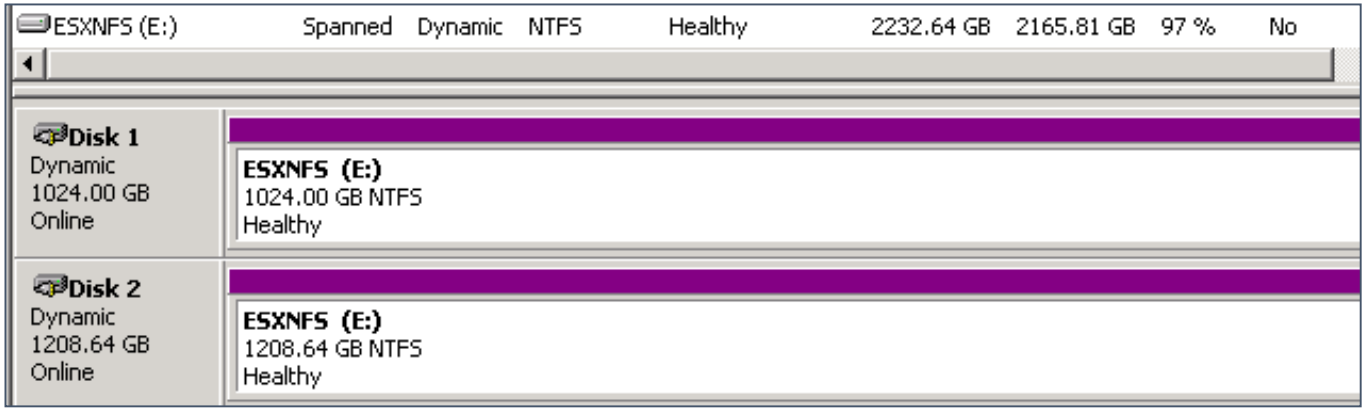

Storage Terminology - File Systems

- Before a volume can be used it must be formatted with a File System.
- A **File System** is the method for storing and organizing files and data on the volume.
- **Formatting** is the act of setting up an empty file system on a volume.
- Some file systems implement **journaling**, which writes to the volume in batches, improving performance and reducing the chances for errors.
- Popular File Systems EXT4, ZFS, NTFS, HFS (Mac), zFAT

Putting it all together **Disk** to **File System**.

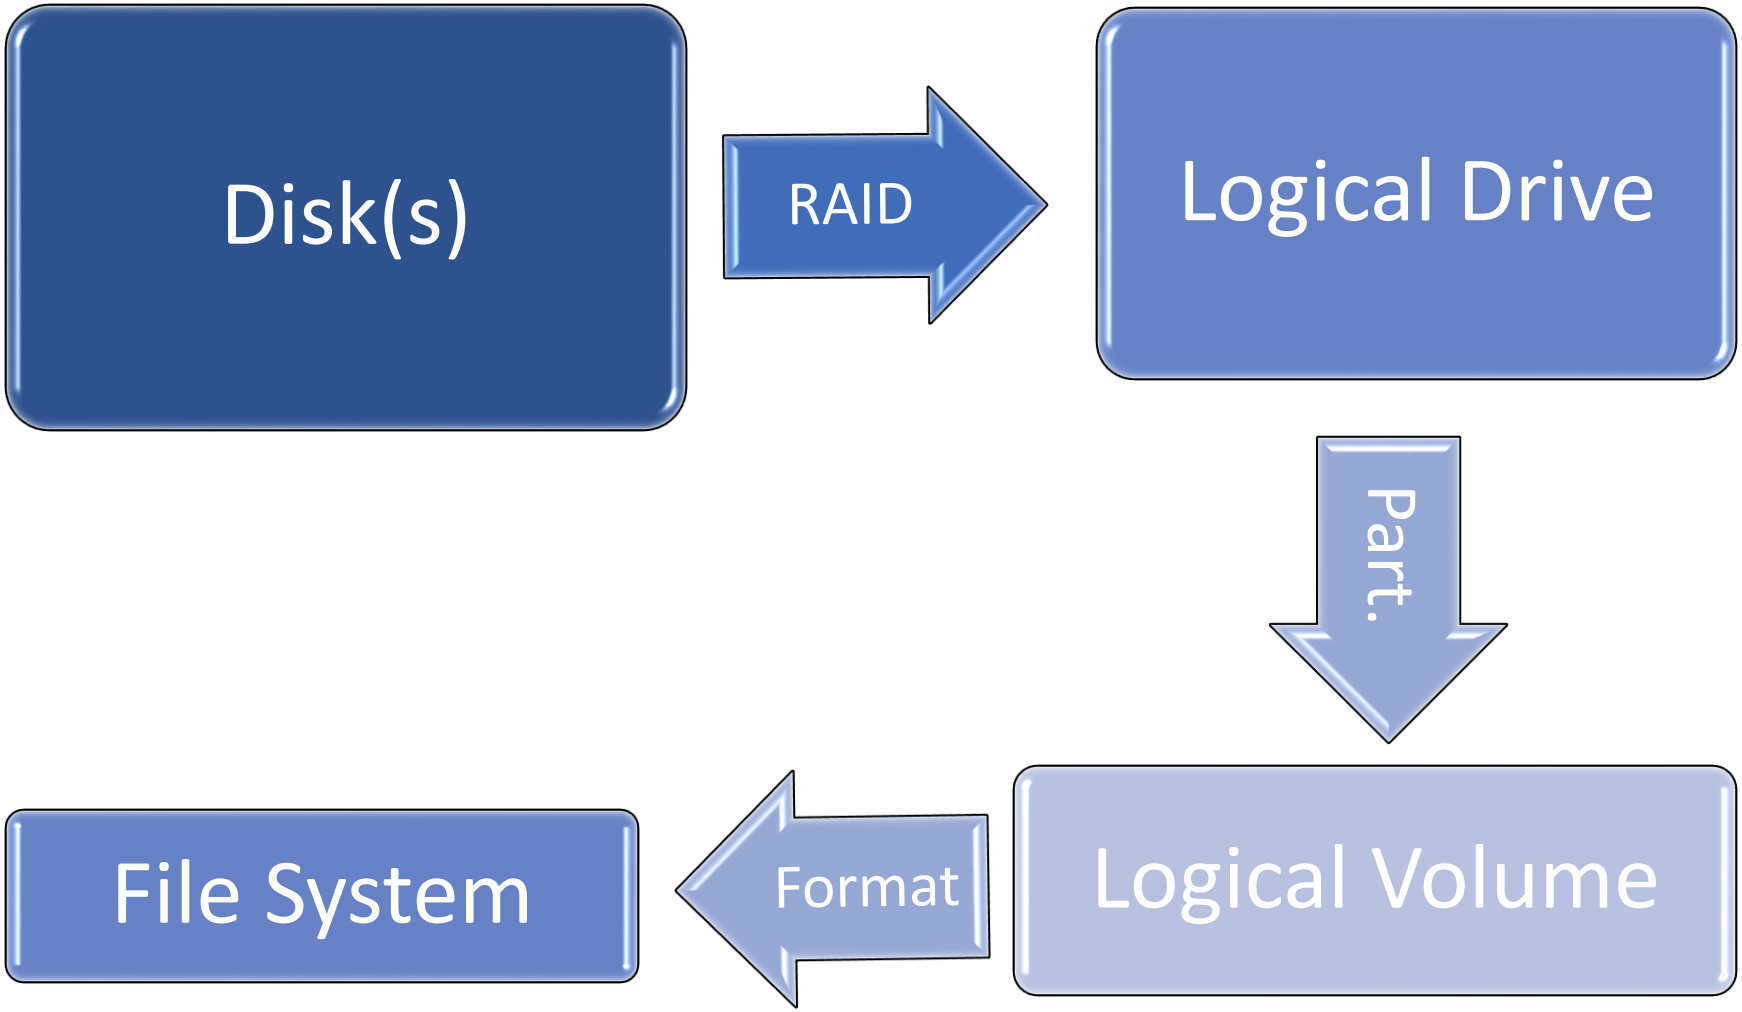

Logical Storage Systems

- DAS – Direct Attached Storage
  - Connects to one host
- NAS – Network Attached Storage
  - Connects to many hosts over common communications network (TCP/IP)
- SAN – Storage Area Network
  - Connects to many hosts over dedicated network for storage.
    - iSCSI or
    - Fibre Channel.

DAS, NAS, and SAN Oh My!

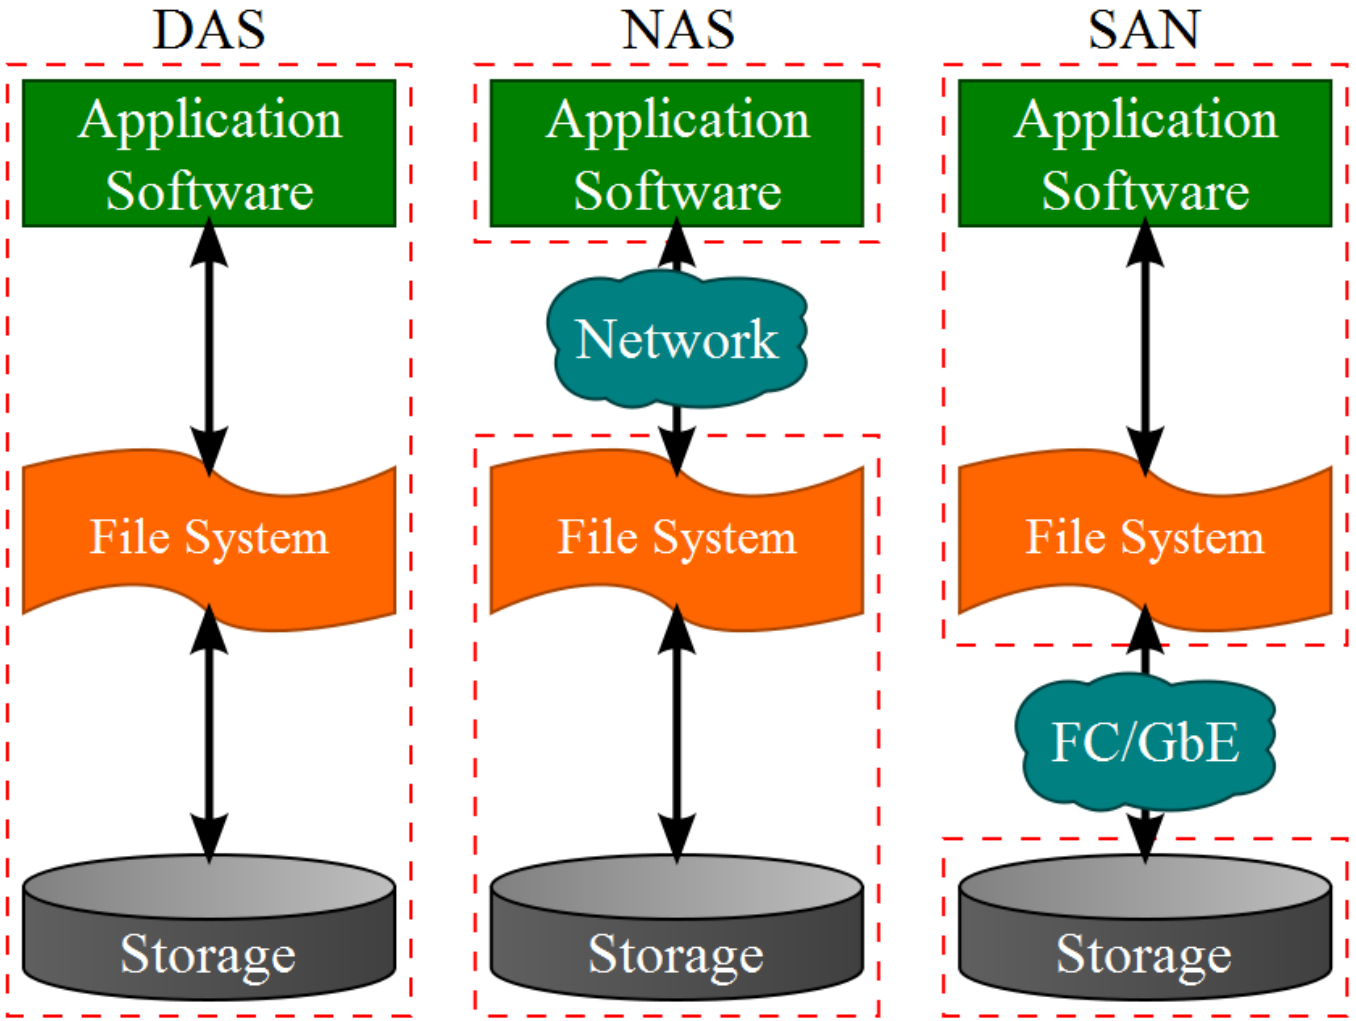

*Taken from Wikipedia Commons

RAID

- Redundant Array of Inexpensive Disks
- A technique to provide storage that is larger, more reliable and faster than what a single disk drive can provide.
- The RAID array of Physical Disks is treated as one logical Disk. (known as a LUN, logical unit number)
- Some RAID supports the use of Hot Spares in the event of a drive failure.

Disk Arrays in Servers

RAID Levels (most common)

Level	Description	Min # Disks	Space Efficiency	Fault Tolerance?
0	Striping	2	N	No
1	Mirroring	2	1	Yes
5	Parity Striping	N	N-1	Yes
6	Dual Parity Striping	N	N-2	Yes
10	Mirrored Stripe	N (Even)	N/2	Yes

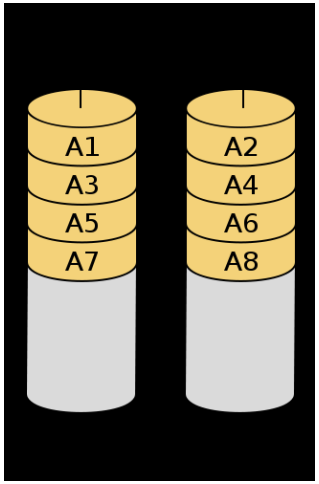

RAID0

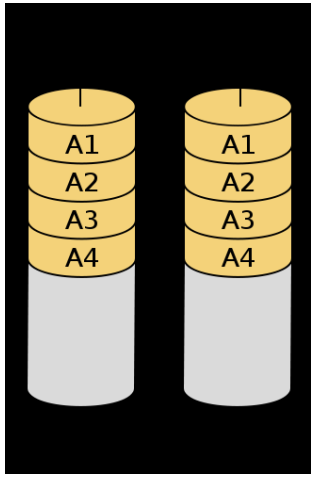

RAID1

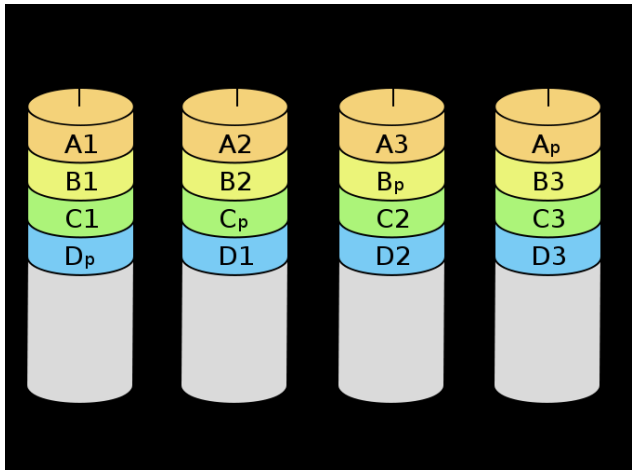

RAID5

Calculating RAID capacity and MTBF

- **Example**:
  One of your servers has a storage array with 8 500GB drives, configured in RAID 5 with one hot spare. According to the manufacturer, the MTBF for the drives are 1.2 Million hours.

- What is the total capacity of the array?

- What is the MTBF? For this server?

Calculating RAID capacity and MTBF

- **Total Capacity**:
  8 total drives – 1 hot spare = 7 drives. RAID 5 has a space efficiency of N-1, so 7-1 = 6. There are 6 drives in used in the array's capacity for a total of 6 * 500GB = 3.0 TB
- Mean Time Between Failure (MTBF):
  1.2 million / 8 drives = 150,000 hours (17 years)

Managing Storage – Considerations?

- Consider storage a Community resource
  - Charge units for their use
- Plan for the future
  - Monitor growth trends, changes
  - Be able to chart storage use over time.
  - Its one thing to know you need more storage, it is an entirely different thing to know why.
- Dedicate resources to cleaning up
  - Help users archive what the no longer need.

NAS File Services

- Network attached storage requires a service to share files with clients
- SMB – Server Message Block
  - Windows File Sharing
  - Samba on Linux
- NFS – Network File System
  - Used by Unix/Linux systems to share files