


IT Infrastructure Architecture

Infrastructure Building Blocks
and Concepts

Storage

Introduction

- Every day, approximately 15 petabytes of new information is generated worldwide
- The total amount of digital data doubles approximately every two years



History

- Early computers used a very basic persistent storage system, based on punched cards or paper tape
- Drum memory was one of the first magnetic read/write storage systems
 - It was widely used in the 1950s and into the 1960s
 - Consisted of a large rotating metal cylinder that was coated on the outside with magnetic recording material
 - Multiple rows of fixed read-write heads were placed along the drum, each head reading or writing to one track
 - The drum could store 62 kB of data



History – Hard disks

- The first commercial digital disk storage device was part of the IBM RAMAC 350 system, shipped in 1956
 - Approximately 5 MB of data
 - Fifty 61 cm diameter disks
 - Weighed over a ton
- Over the years:
 - Physical size of hard disks shrunk
 - Magnetic density increased
 - Rotation speed increased from 3,600 rpm to 15,000 rpm
 - Seek times lowered as a result of using servo controlled read/write heads instead of stepper motors



History – Tapes

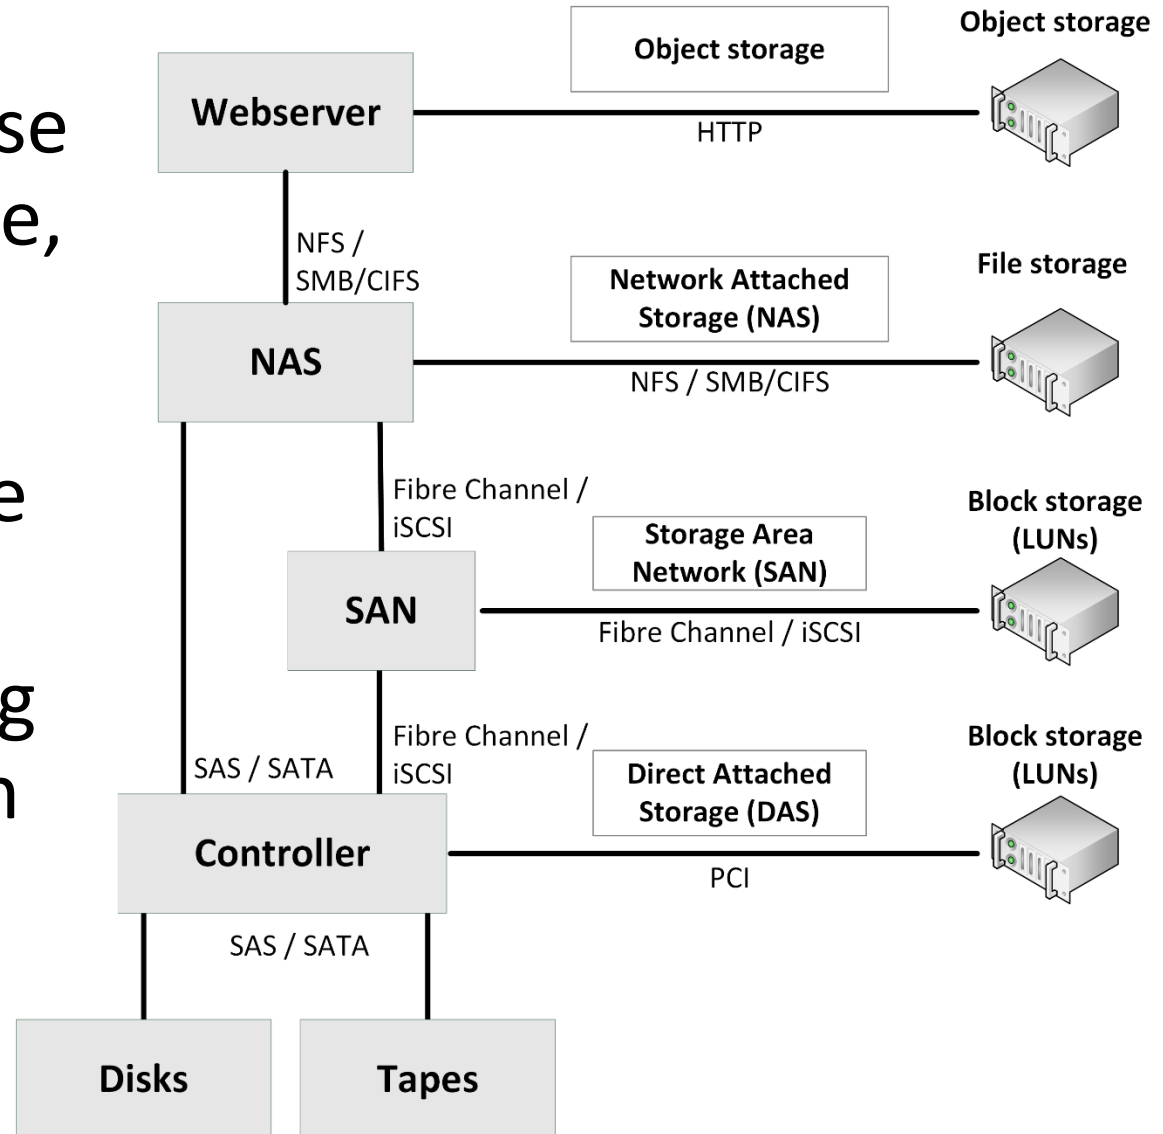
- The IBM 726, introduced in 1952, was one of the first magnetic tape systems
 - 2 MB per 20-centimeter-diameter reel of tape
- Reel tapes were used until the late 1980s, mostly in mainframes
- In 1984, DEC introduced the Digital Linear Tape (DLT)
 - Super DLT (SDLT) tape cartridges can store up to 300 GB of data
- Linear Tape Open (LTO) was originally developed in the late 1990s
 - LTO version 7 was released in 2015 and can hold up to 6 TB of data



Storage building blocks

Storage model

- Most servers use external storage, sometimes combined with internal storage
- A model of storage building blocks is shown on the right



Disks – command sets

- Disks are connected to disk controllers using a command set, based on either ATA or SCSI
 - Advanced Technology Attachment (ATA), also known as IDE, uses a relatively simple hardware and communication protocol to connect disks to computers (mostly PCs)
 - Small Computer System Interface (SCSI) is a set of standards for physically connecting and transferring data between computers (mostly servers) and peripheral devices, like disks and tapes
 - The SCSI command set is complex - there are about 60 different SCSI commands in total
- Serial interfaces replaced the parallel interfaces, but the disk commands are still the same

Mechanical hard disks

- Mechanical disks consist of:

- A vacuum sealed case
- One or more spinning magnetic disks on one spindle
- A number of read/write heads that can move to reach each part of the spinning disks



Mechanical hard disks

- Serial ATA (SATA) disks
 - Low-end high-capacity disks
 - Ideal for bulk storage applications (like archiving or backup)
 - Have a low cost per gigabyte
 - Often used in PCs and laptops
 - Use the SMART command set to control the disk

Mechanical hard disks

- Serial Attached SCSI (SAS) disks
 - Relatively expensive
 - High end disks
 - Spinning disk platters with a rotational speed of 10,000 or 15,000 rpm
 - Typically have 25% of the capacity of SATA or NL-SAS disks
 - Uses the SCSI command set that includes error-recovery and error-reporting and more functionality than the SMART commands used by SATA disks

Mechanical hard disks

- Near-Line SAS (NL-SAS) disks
 - Have a SAS interface, but the mechanics of SATA disks
 - Can be combined with faster SAS disks in one storage array

Solid State Drives (SSDs)

- SSD disks don't have moving parts
- Based on flash technology
 - Flash technology is semiconductor-based memory that preserves its information when powered off
- Connected using a standard SAS disk interface
- Data can be accessed much faster than using mechanical disks
 - Microseconds vs. milliseconds
- Most storage vendors now offer all-flash arrays – storage systems using only SSD disks



Solid State Drives (SSDs)


- SSDs consume less power, and therefore generate less heat, than mechanical disks
- They have no moving parts
- They generate no vibrations that could influence or harm other components, or shorten their lifetime
- The main disadvantage of SSDs is their price per gigabyte
 - Considerably higher than mechanical disks
 - Price per GB is dropping fast

Solid State Drives (SSDs)

- Flash memory can only be rewritten a limited number of times
 - SSD disks “wear out” more rapidly than mechanical disks
 - SSDs keep track of the number of times a sector is rewritten, and map much used sectors to spare sectors if they are about to wear out
 - It is important to monitor the wear level of heavily used SSDs
 - Replace them before they break

Disk capacity - Kryder's law

- The density of information on hard drives doubles every 13 months
- An average single disk drive in 2025 will hold more than 20,000 TB (20 PB) of data



Disk capacity - Kryder's law

- The picture on the right shows 8 bytes core memory and 8 GB SD flash memory
 - An increase of 1,000,000,000 times in 50 years
- To have full benefits of Kryder's law, the storage infrastructure should be designed to handle just in time expansion
 - Buy disks as late as possible!

