UNIVERSITY of VIRGINIA

CS4740
CLOUD COMPUTING

Intro to Cloud

Prof. Chang Lou, UVA CS, Fall 2025

AGENDA

- What is "Cloud"?
- Where does the idea of Cloud Computing come from?
- What does Cloud look like?

QUIZ

- Which company is the leading player of today's Cloud market?
  - A. Amazon
  - B. Alibaba Cloud
  - C. Google
  - D. Microsoft

QUIZ

— Which company is the leading player of today's Cloud market?

The Big Three Stay Ahead in Ever-Growing Cloud Market

Worldwide market share of leading cloud infrastructur service providers in Q2 2025*

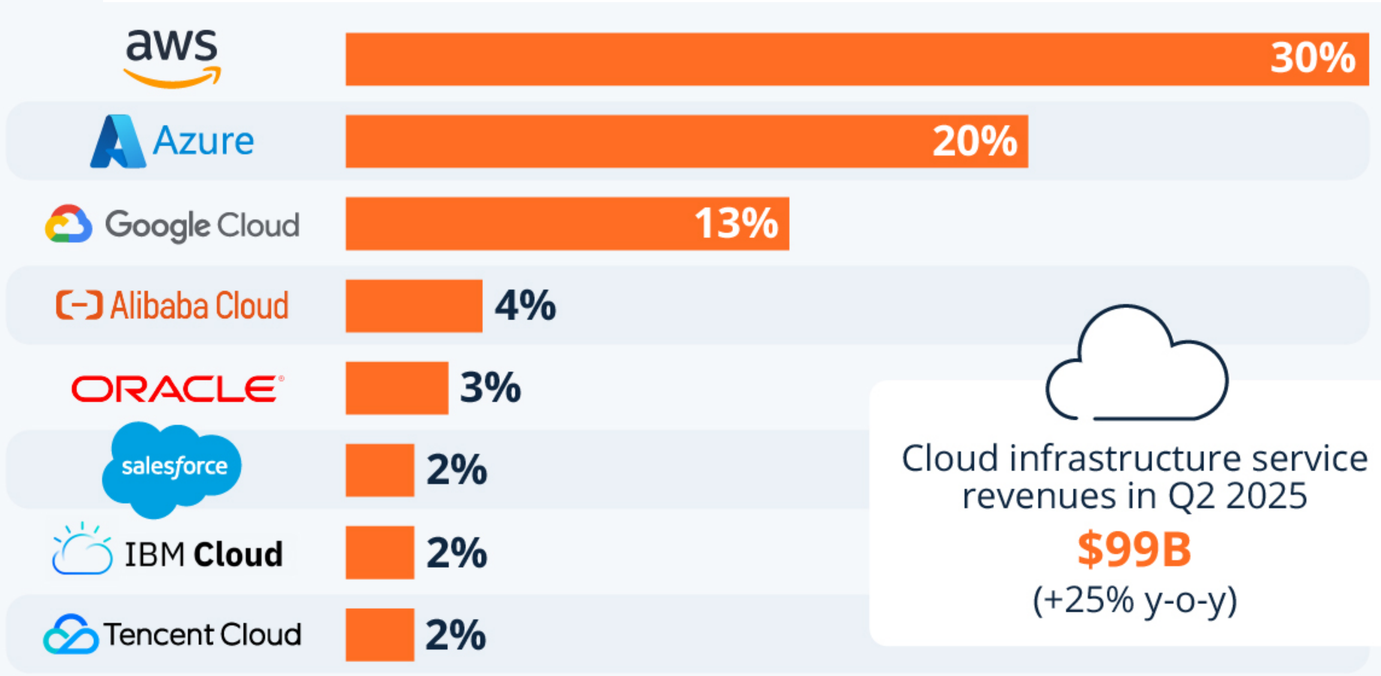

* Includes platform as a service (PaaS) and infrastructure as a service (IaaS) as well as hosted private cloud services

Source: Synergy Research Group

(THERE WAS) THE HYPE!

- Forrester in 2010 – Cloud computing will go from $40.7 billion in 2010 to $241 billion in 2020.
- Today: cloud market is $676B (expected to reach $2.2T by 2032)

- Companies and even Federal/state governments using cloud computing now: fbo.gov

MANY CLOUD PROVIDERS

- AWS: Amazon Web Services
  - EC2: Elastic Compute Cloud
  - S3: Simple Storage Service
  - EBS: Elastic Block Storage
- Microsoft Azure
- Google Cloud/Compute Engine/AppEngine
- Rightscale, Salesforce, EMC, Gigaspaces, 10gen, Datastax, Oracle, VMWare, Yahoo, Cloudera, and many many more!
- They become the backbone of modern computing.

TWO CATEGORIES OF CLOUDS

- Can be either a (i) **public** cloud, or (ii) **private** cloud
- Private clouds are accessible only to company employees
- Public clouds provide service to any paying customer:
  - Amazon S3 (Simple Storage Service): store arbitrary datasets, pay per GB-month stored.
    - Recently: 0.09c to 2 c per GB month
  - Amazon EC2 (Elastic Compute Cloud): upload and run arbitrary OS images, pay per CPU hour used
    - Recently: 2 c per CPU hr to 40c per CPU hr (depending on strength), only CPUs not GPUs
  - Google cloud: similar pricing ranges as above
  - Google AppEngine/Compute Engine: develop applications within their appengine framework, upload data that will be imported into their format, and run

WHY WE NEED CLOUD?

WHY WE NEED CLOUD?

- Customers Save Time and $$$!
  - Dave Power, Associate Information Consultant at Eli Lilly and Company: “With AWS, Powers said, a new server can be up and running in three minutes (it used to take Eli Lilly seven and a half weeks to deploy a server internally) and a 64-node Linux cluster can be online in five minutes (compared with three months internally). … It's just shy of instantaneous.”
  - Jim Swartz, CIO, Sybase: “At Sybase, a private cloud of virtual servers inside its datacenter has saved nearly $US2 million annually since 2006, Swartz says, because the company can share computing power and storage resources across servers.”

WHY WE NEED CLOUD?

- Scalability
- Flexibility
- Enhanced security
- Fault tolerance

- Q: any downside?

What is cloud?

WHAT IS CLOUD?

— It’s a cluster!

— It’s a supercomputer!

— It’s a datastore!

WHAT IS CLOUD?

— It’s a cluster!

— It’s a supercomputer!

— It’s a datastore!

— Cloud = **Lots of storage + compute cycles nearby**

WHAT IS CLOUD?

- A **single-site** cloud (aka “Datacenter”) consists of
  - Compute nodes (grouped into racks) (2)
  - Switches, connecting the racks
  - A network topology, e.g., hierarchical
  - Storage (backend) nodes connected to the network (3)
  - Front-end for submitting jobs and receiving client requests (1)
  - Software Services
- A geographically **distributed** cloud consists of
  - Multiple such sites
  - Each site perhaps with a different structure and services

*(1,2,3: Often called “three-tier architecture”)

A SAMPLE CLOUD TOPOLOGY

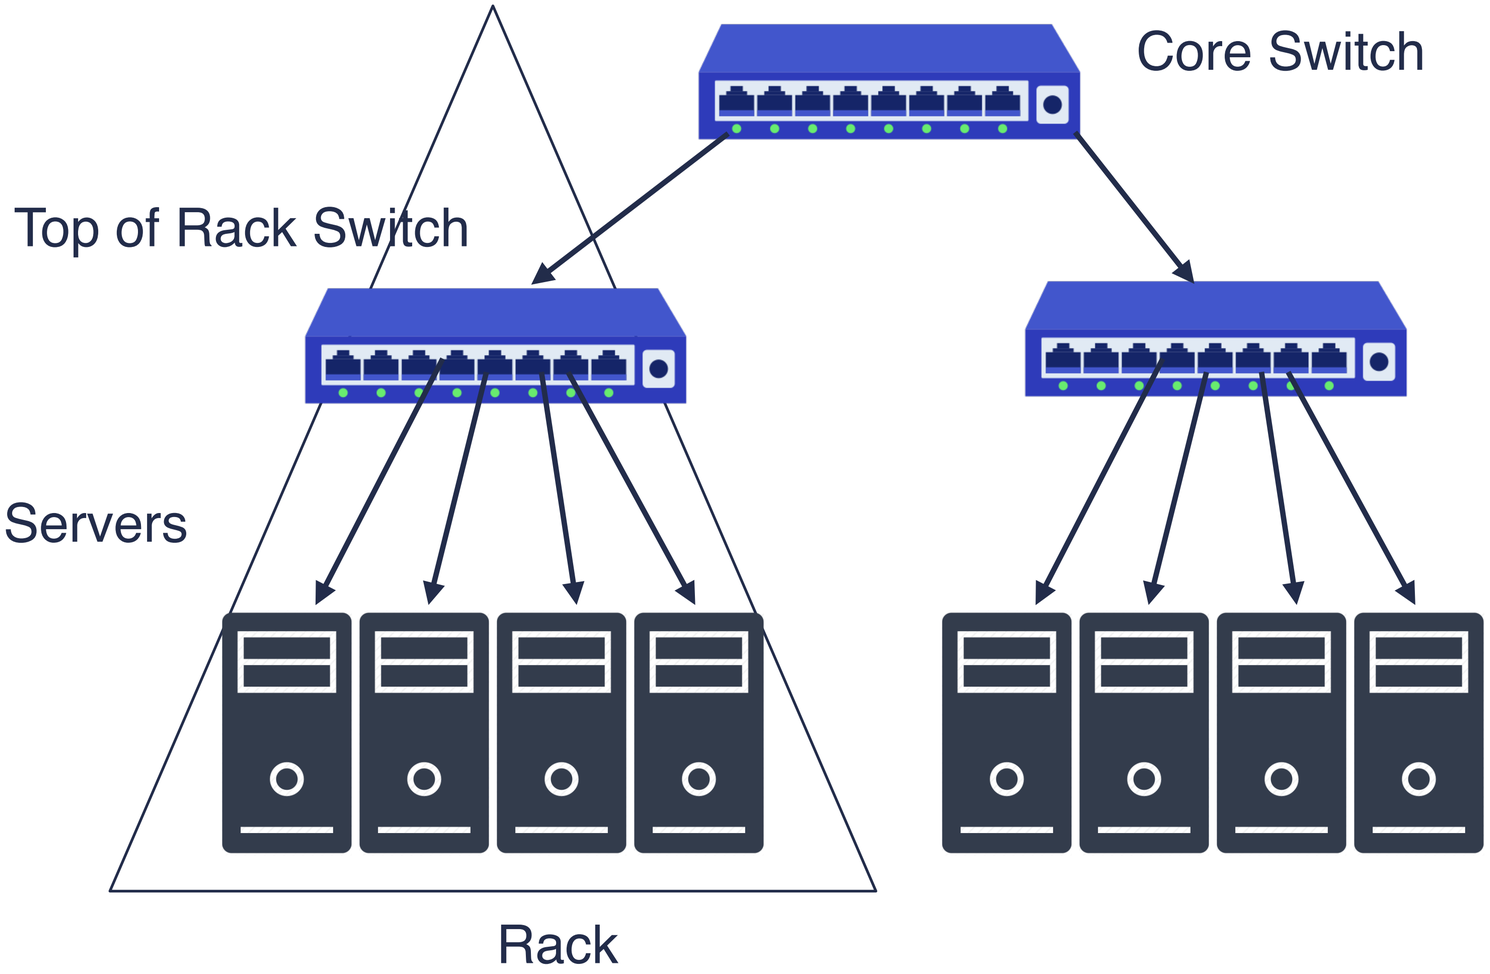

A CLOUDY HISTORY OF TIME

A CLOUDY HISTORY OF TIME

- First large datacenters: ENIAC, ORDVAC, ILLIAC
- Many used vacuum tubes and mechanical relays

A CLOUDY HISTORY OF TIME

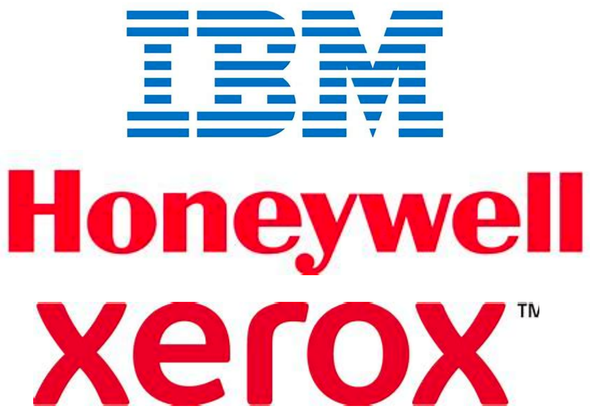

- Data Processing Industry
  - 1968: $70 M. 1978: $3.15 Billion
- Timesharing Industry (1975):
  - Market Share: Honeywell 34%, IBM 15%
  - Xerox 10%, CDC 10%, DEC 10%, UNIVAC 10%
  - Honeywell 6000 & 635, IBM 370/168, Xerox 940..

A CLOUDY HISTORY OF TIME

- Why time-sharing started to decline?
- Answer: Rise of PC

1940 1950 1960 1970 1980 1990 2000 2010

First datacenters!

Timesharing Companies & Data Processing Industry

A CLOUDY HISTORY OF TIME

- Grids (1980s-2000s):
  - GriPhyN (1970s-80s)
  - Open Science Grid and Lambda Rail (2000s)
  - Globus & other standards (1990s-2000s)

A CLOUDY HISTORY OF TIME

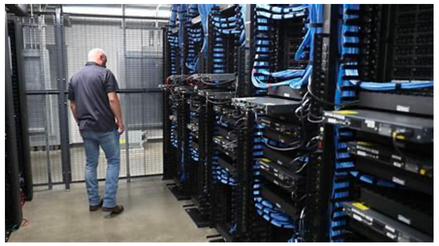

CRAY
THE SUPERCOMPUTER COMPANY

- Berkeley NOW Project
- Supercomputers
- Server Farms (e.g., Oceano)

- Difference between Grids and Clusters?

1940 1950 1960 1970 1980 1990 2000 2010

First datacenters!

Timesharing Companies & Data Processing Industry

Grids, Clusters

A CLOUDY HISTORY OF TIME

- P2P Systems (90s-00s)
  - Many Millions of users
  - Many GB per day

Peer-to-Peer systems

PCs (not distributed)

1940 1950 1960 1970 1980 1990 2000 2010

First datacenters!

Timesharing Companies & Data Processing Industry

Grids, Clusters

A CLOUDY HISTORY OF TIME

- P2P Systems (90s-00s)
  - Many Millions of users
  - Many GB per day

Peer-to-Peer systems

PCs (not distributed)

1940 1950 1960 1970 1980 1990 2000 2010

First datacenters!

Timesharing Companies & Data Processing Industry

Grids, Clusters

Cloud

TRENDS: TECHNOLOGY

- Doubling Periods – storage: 12 mos, bandwidth: 9 mos, and (what law is this?) cpu compute capacity: 18 mos
- Then and Now
  - Bandwidth
    - 1985: mostly 56Kbps links nationwide
    - Today: Tbps links widespread
  - Disk capacity
    - Today's PCs have TBs, far more than a 1990 supercomputer

TRENDS: USERS

- Then and Now:
  - 1990: biologists were running small single-molecule simulations
  - Today: CERN's Large Hadron Collider producing many PB/year

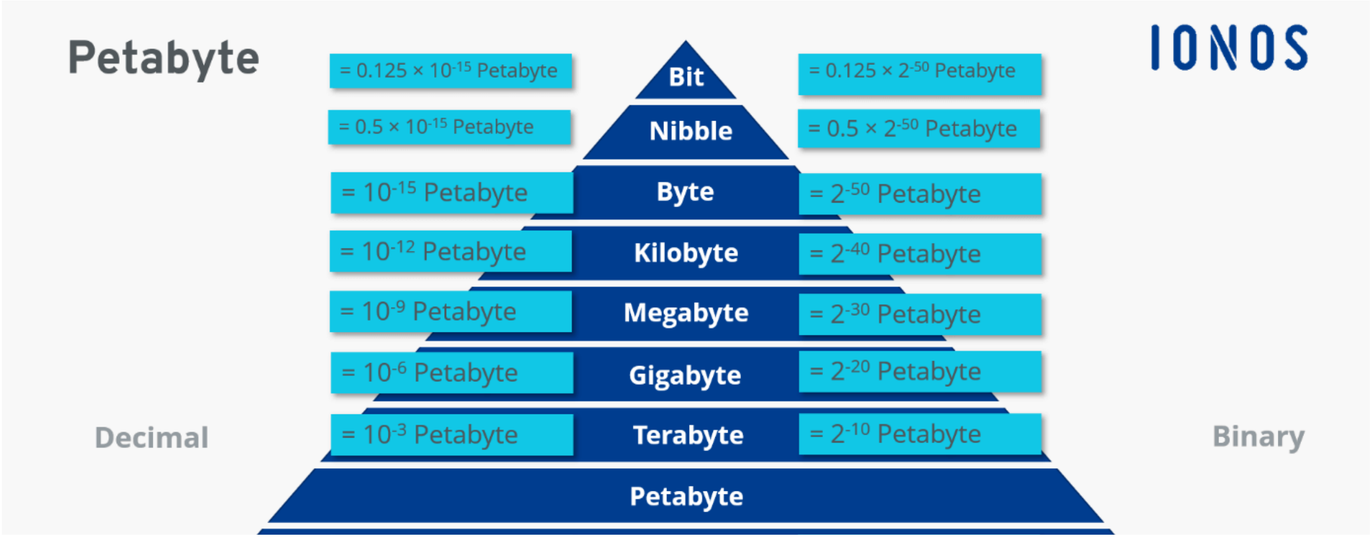

PROPHECIES

- In 1965, MIT's Fernando Corbató and the other designers of the Multics operating system envisioned a computer facility operating “like a power company or water company”.
- Plug your **thin client** into the computing Utility and Play your favorite Intensive Compute & Communicate Application
  - Discussion: Have today’s clouds brought us closer to this reality?