جامعة الامير سلطان
PRINCE SULTAN UNIVERSITY

WHAT ARE DISTRIBUTED SYSTEMS?

CS435 Distributed Systems

Basit Qureshi PhD, FHEA, SMIEEE, MACM
https://www.drbasit.org/

TOPICS

- What is a Distributed System
- Data Center World!
- Dist. Systems Goals
- Types of Dist. Systems
- Applications of Dist. Systems

WHAT IS A DISTRIBUTED SYSTEM

1) Multiple computers
2) Connected by a network
3) Working together

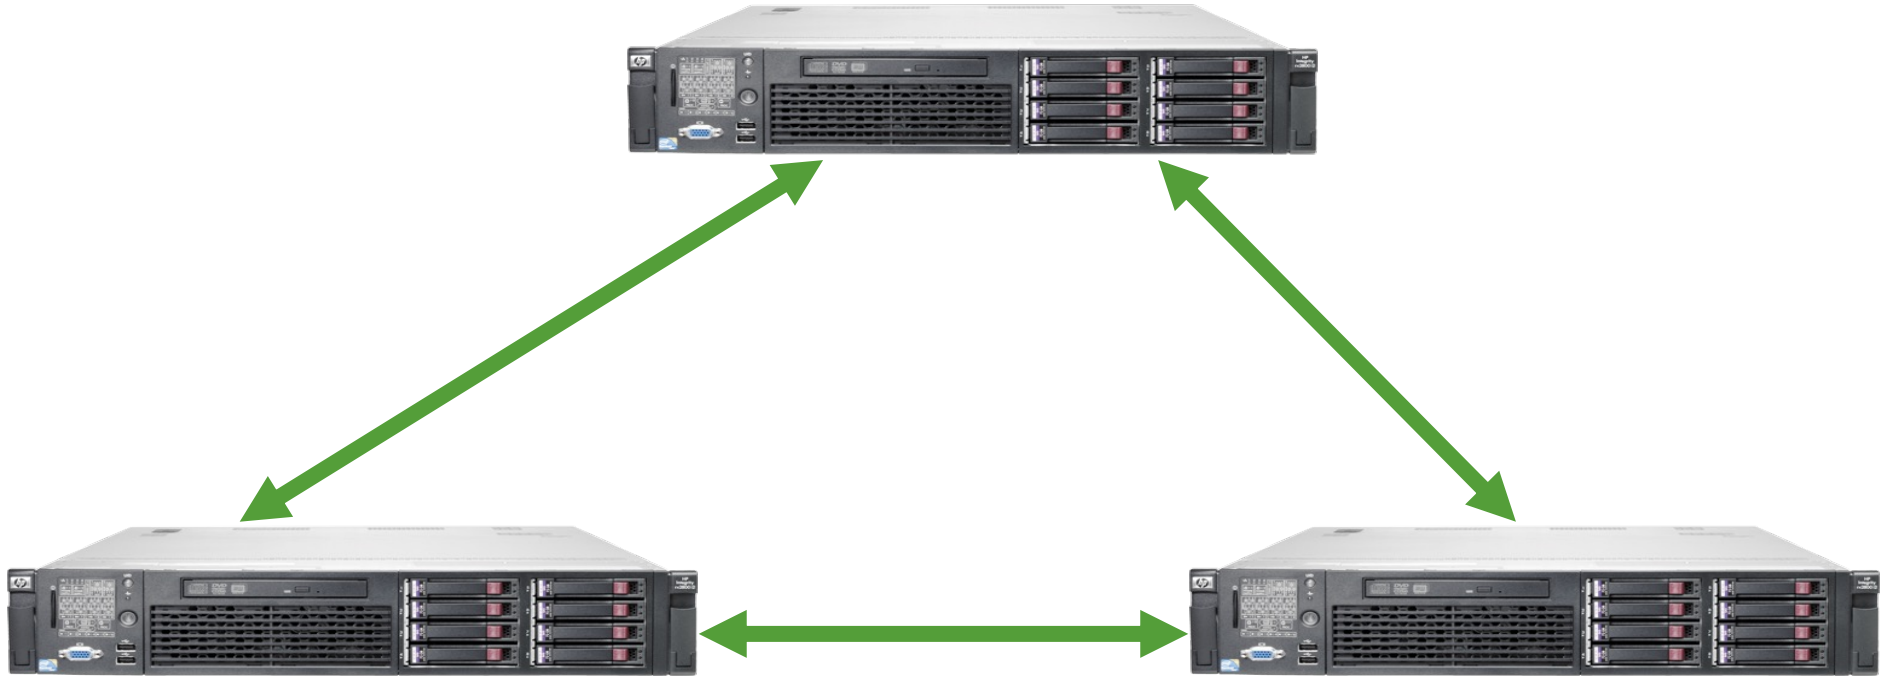

WHAT IS A DISTRIBUTED SYSTEM

1) **Multiple computers**
2) **Connected by a network**
3) **Working together**

WHY?

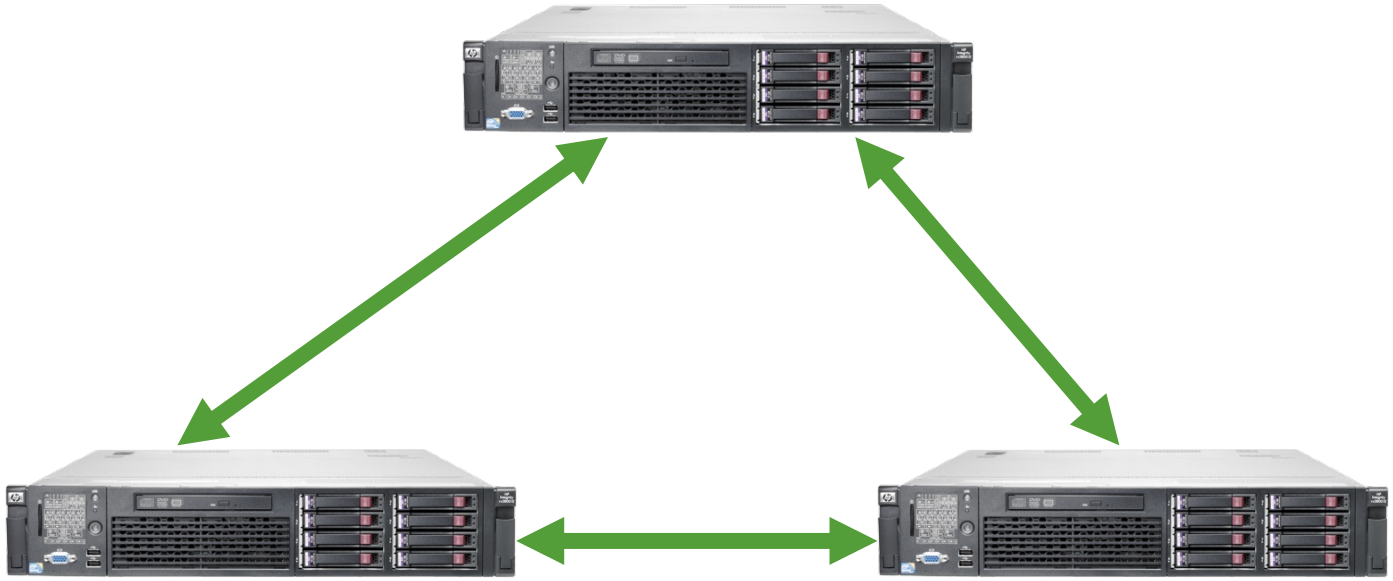

Limited computation/storage/…
Limited Computing Power
Physical location (edge)
Resolution of Failure
.
.
.

WHAT IS A DISTRIBUTED SYSTEM

- Application?
  - Web Search
  - Shopping
  - File Sync
  - Social Networks
  - Music
  - Ride Sharing
  - Video streaming
  - Online gaming
  - Online payments
  - and on and on and on

DATA CENTER WORLD

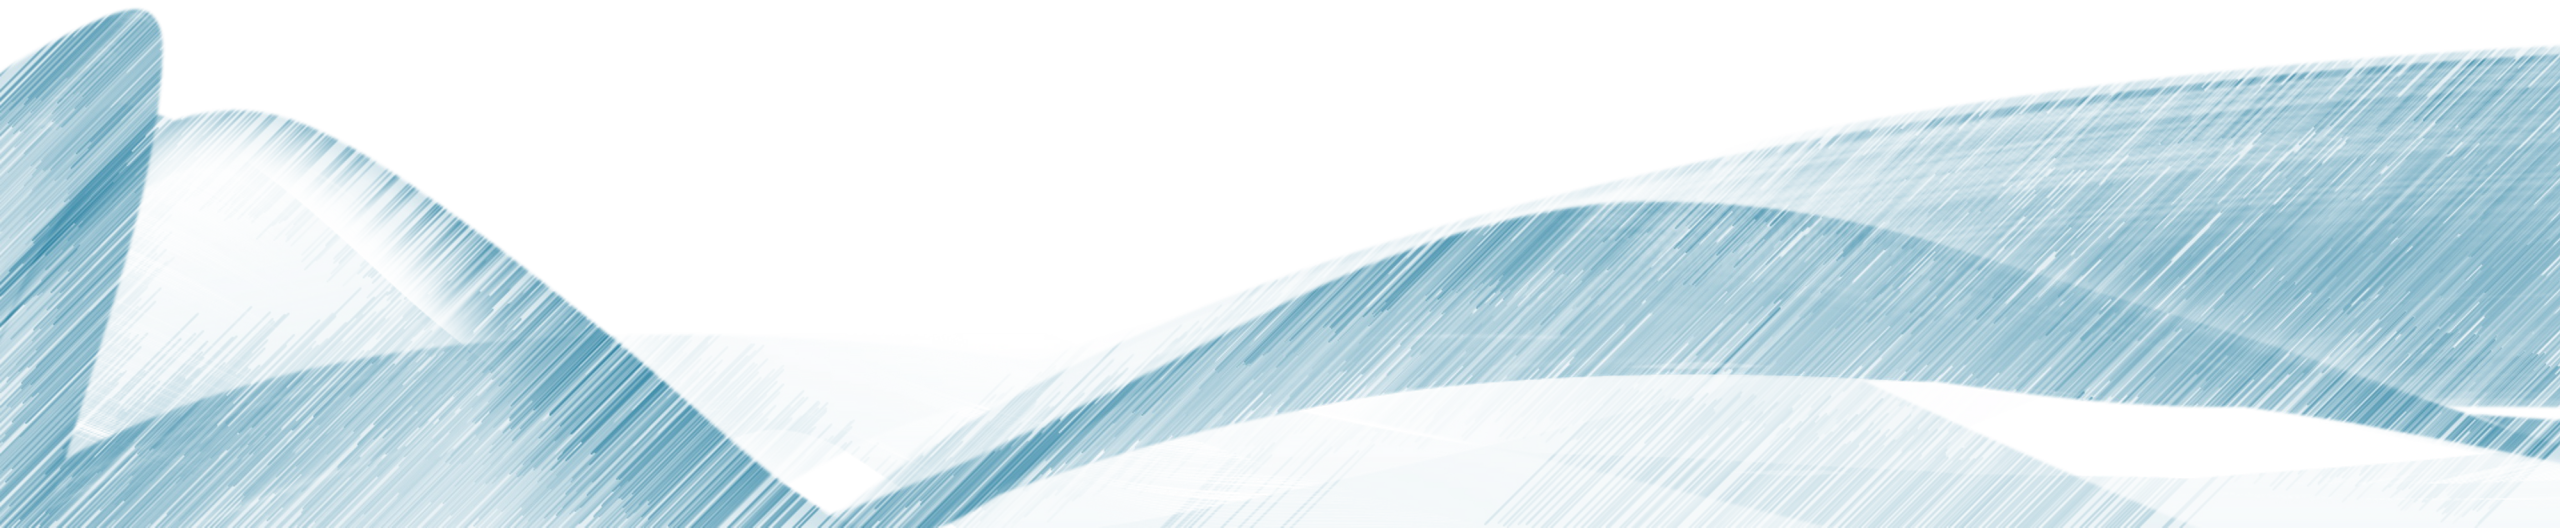

DATA CENTERS

Data centers

- Hundreds/thousands of servers
- Network gear (cables, switches, routers)
- Racks, Floors
- Cooling Units

DATA CENTERS

DATA CENTERS

Google

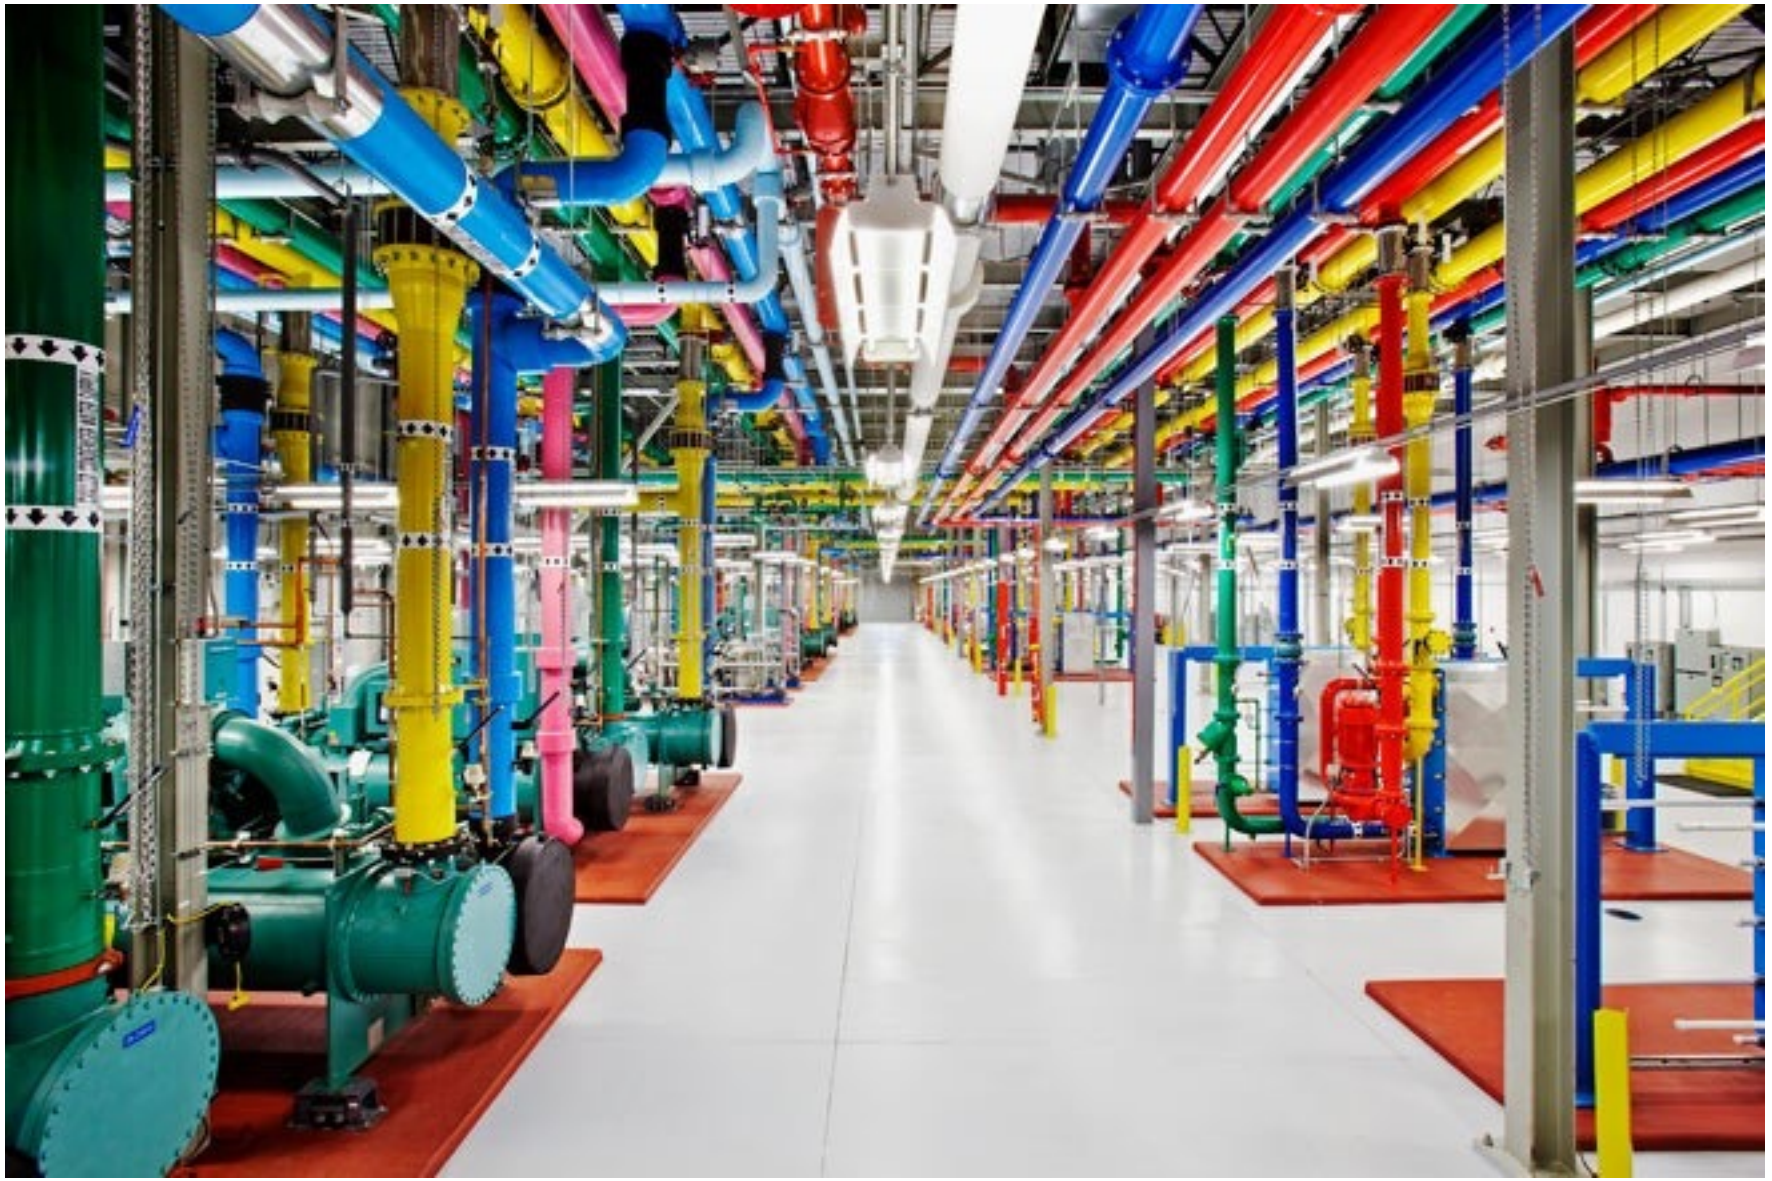

DATA CENTERS

DATA CENTERS

Microsoft

https://blogs.microsoft.com/green/2017/11/08/building-operating-greener-datacenters-commitment-leed-gold/

DATA CENTERS

- 100,000s of physical servers
- 10s MW energy consumption
- Facebook Prineville: $250M physical infra, $1B IT infra
- STC Datacenters: $1B IT infra
- 18 Data centers in Saudi Arabia

https://www.datacentermap.com/saudi-arabia/

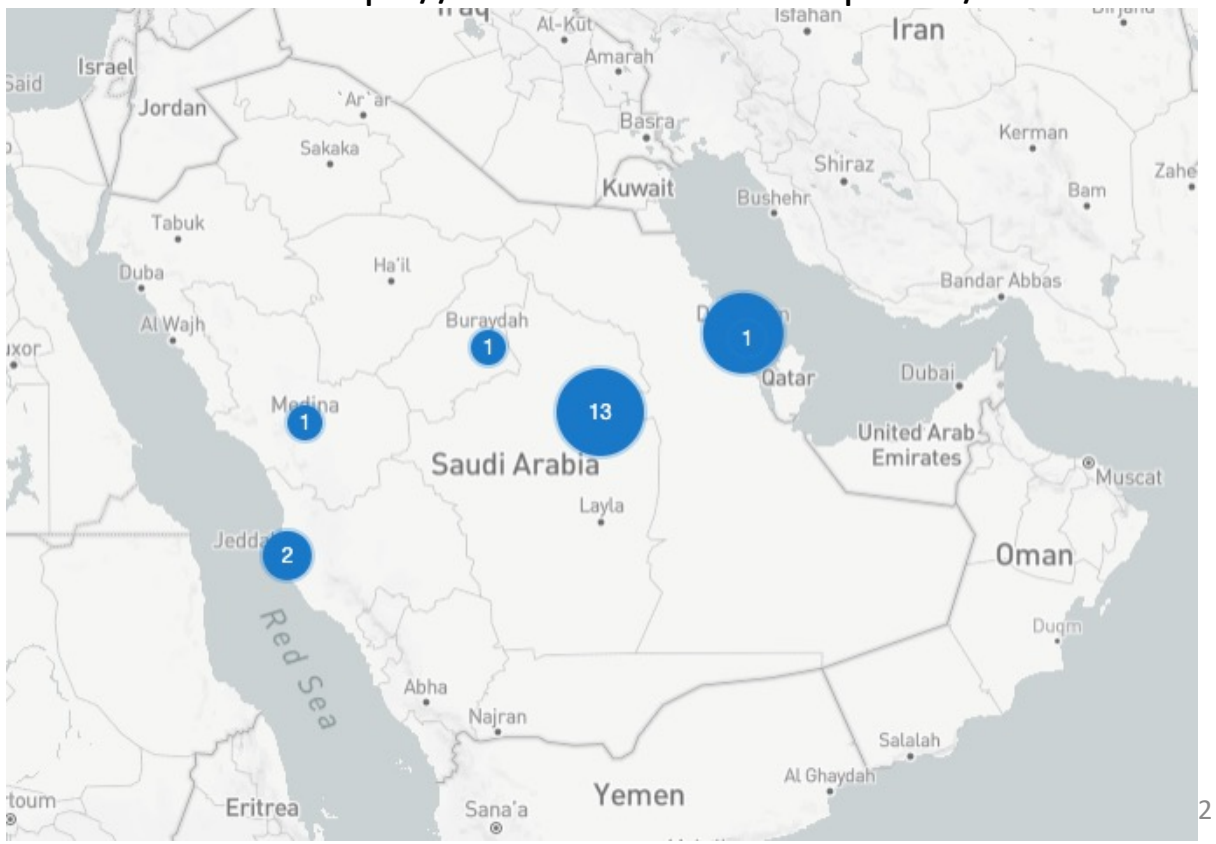

DATA CENTERS

What is a Data Center?

- A data center is a facility used to house computer systems and associated components, such as networking and storage systems, cooling, uninterruptable power supply, air filters…
- A data center typically houses a large number of heterogeneous networked computer systems
- A data center can occupy one room of a building, one or more floors, or an entire building

DATA CENTERS

Data Center Components

- Air conditioning
- Keep all components in the manufacturer's recommended temperature range
- Redundant Power
  - UPS/Generators
  - Multiple power feeds
- Fire protection
- Physical security
  - CCTV/Access Control
- Monitoring Systems
  - Connectivity
  - Multiple ISPs/Leased Lines

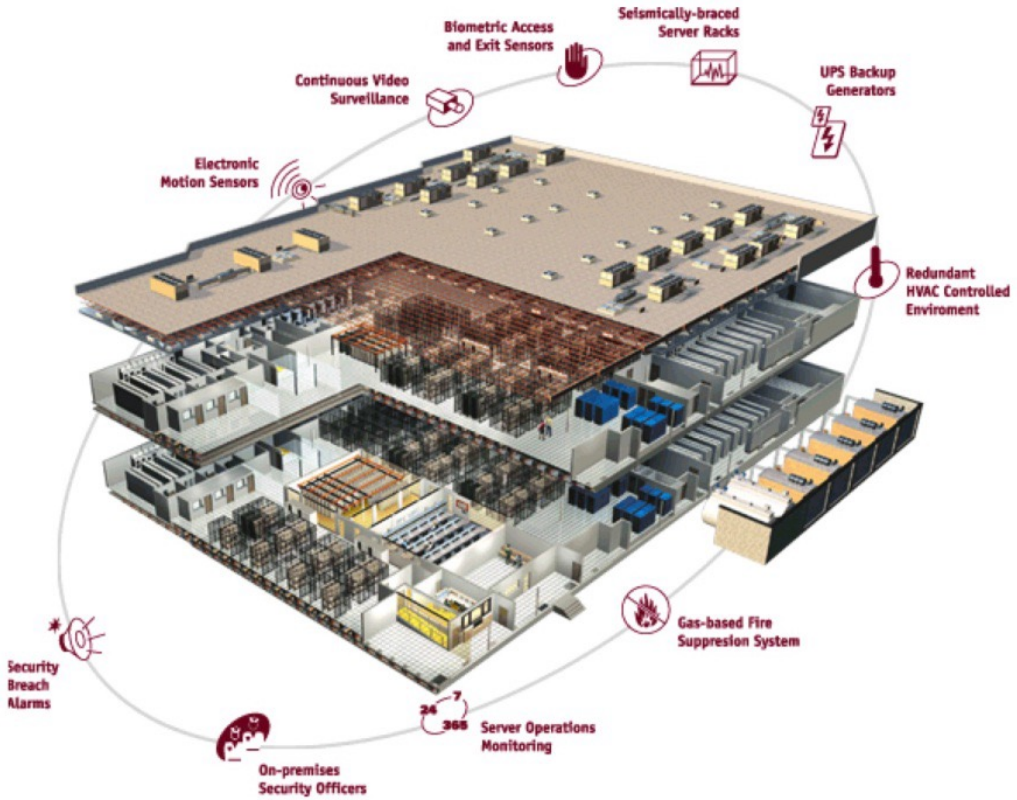

DATA CENTERS

- **Rack-mount servers**
  - Wide, flat standalone servers
  - designed to be stacked on top of each other in a rack
  - Each rack-mount server has its own power supply, cooling fans, network switches, and ports, along with the usual processor, memory, and storage.

- **Blade servers**
  - Fits a chassis to hold blades
  - Contains processors, network controllers, memory and sometime storage;
  - Contains the power supply, network management and other resources for all the blades in the chassis.

DATA CENTERS

Racks

- Equipment (e.g., servers) are typically placed in racks
- Equipment are designed in a modular fashion to fit into rack units (1U, 2U etc.)
- A single rack can hold up to 42 1U servers

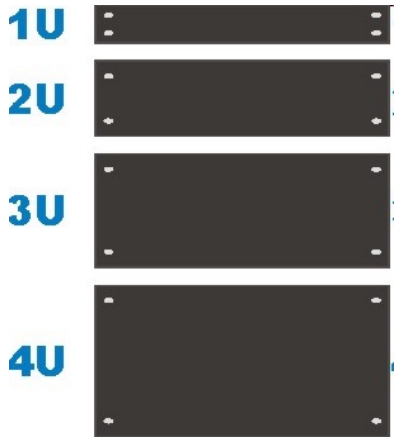

1U server

7U blade server

DATA CENTERS

Blades and Blade Enclosures

- A blade server is a stripped down computer with a modular design
- A blade enclosure holds multiple blade servers and provides power, interfaces and cooling for the individual blade servers

DATA CENTERS

RACK

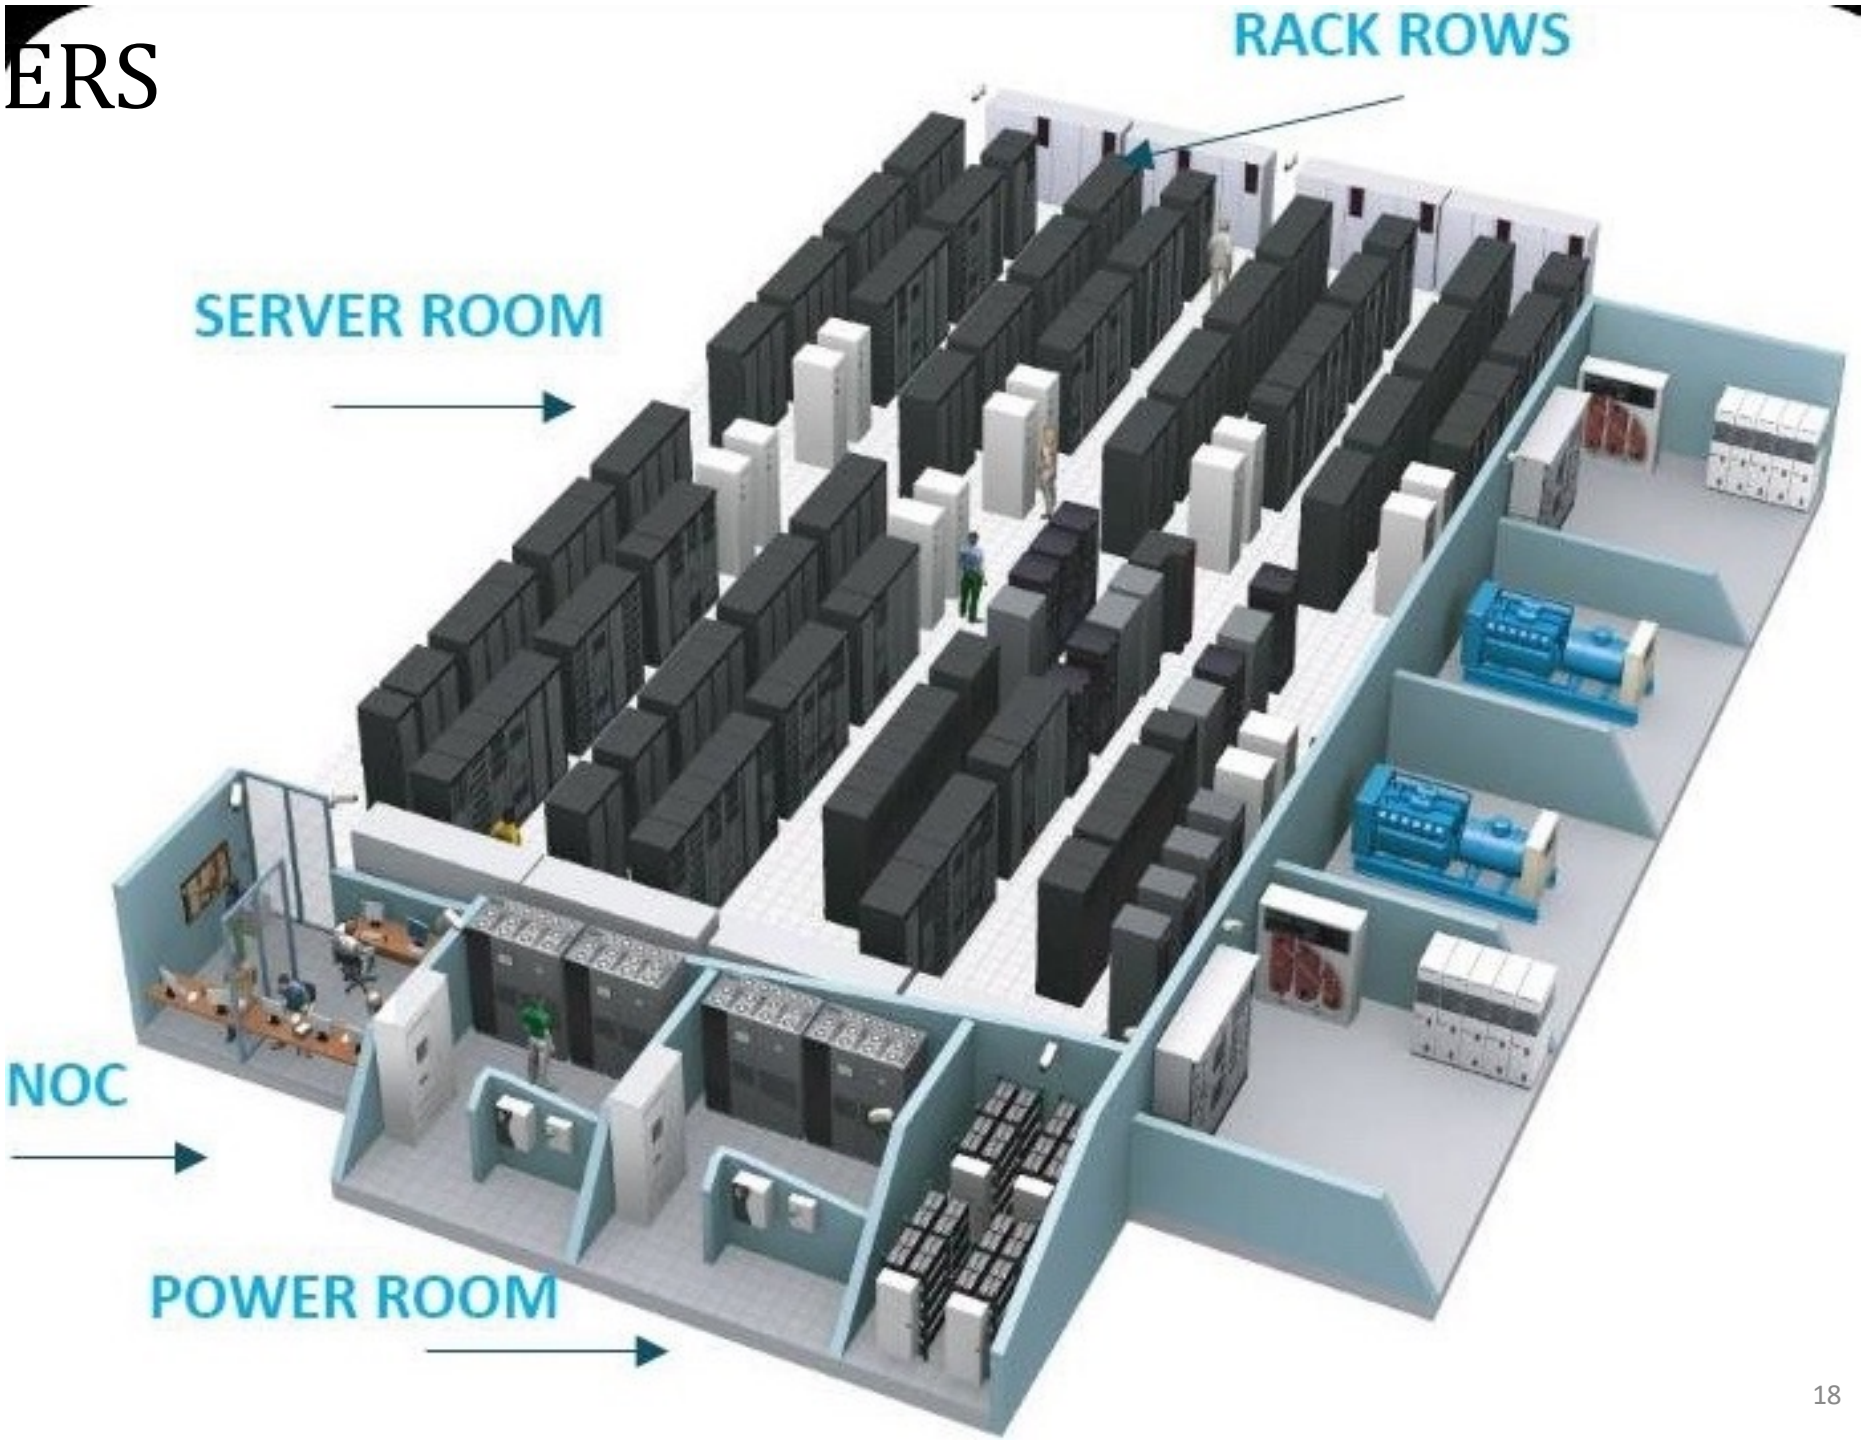

UNDERSTANDING DIST SYS GOALS

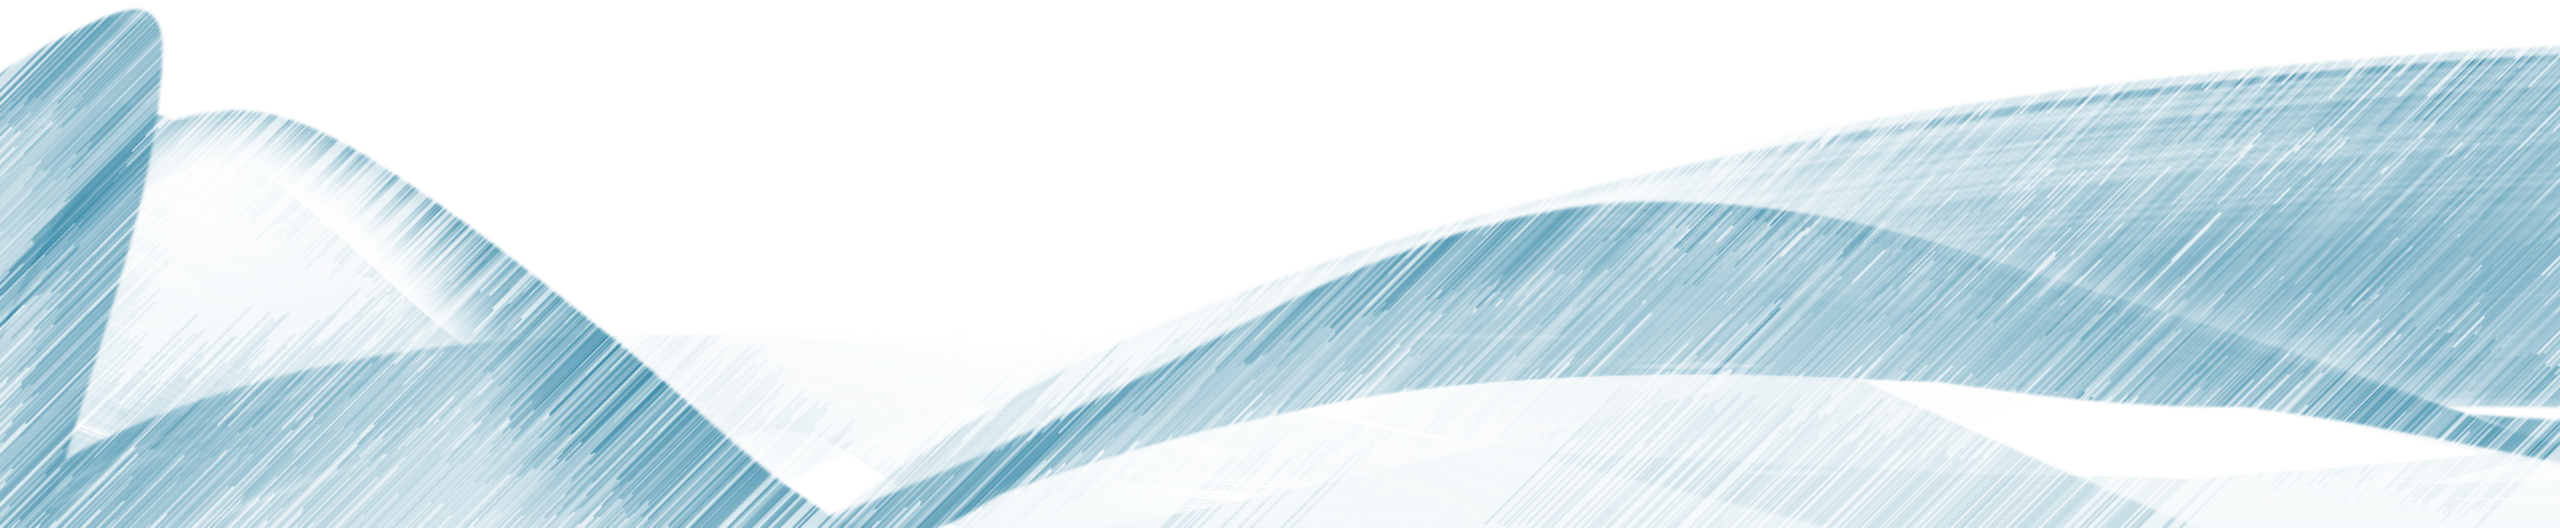

DISTRIBUTED SYSTEMS

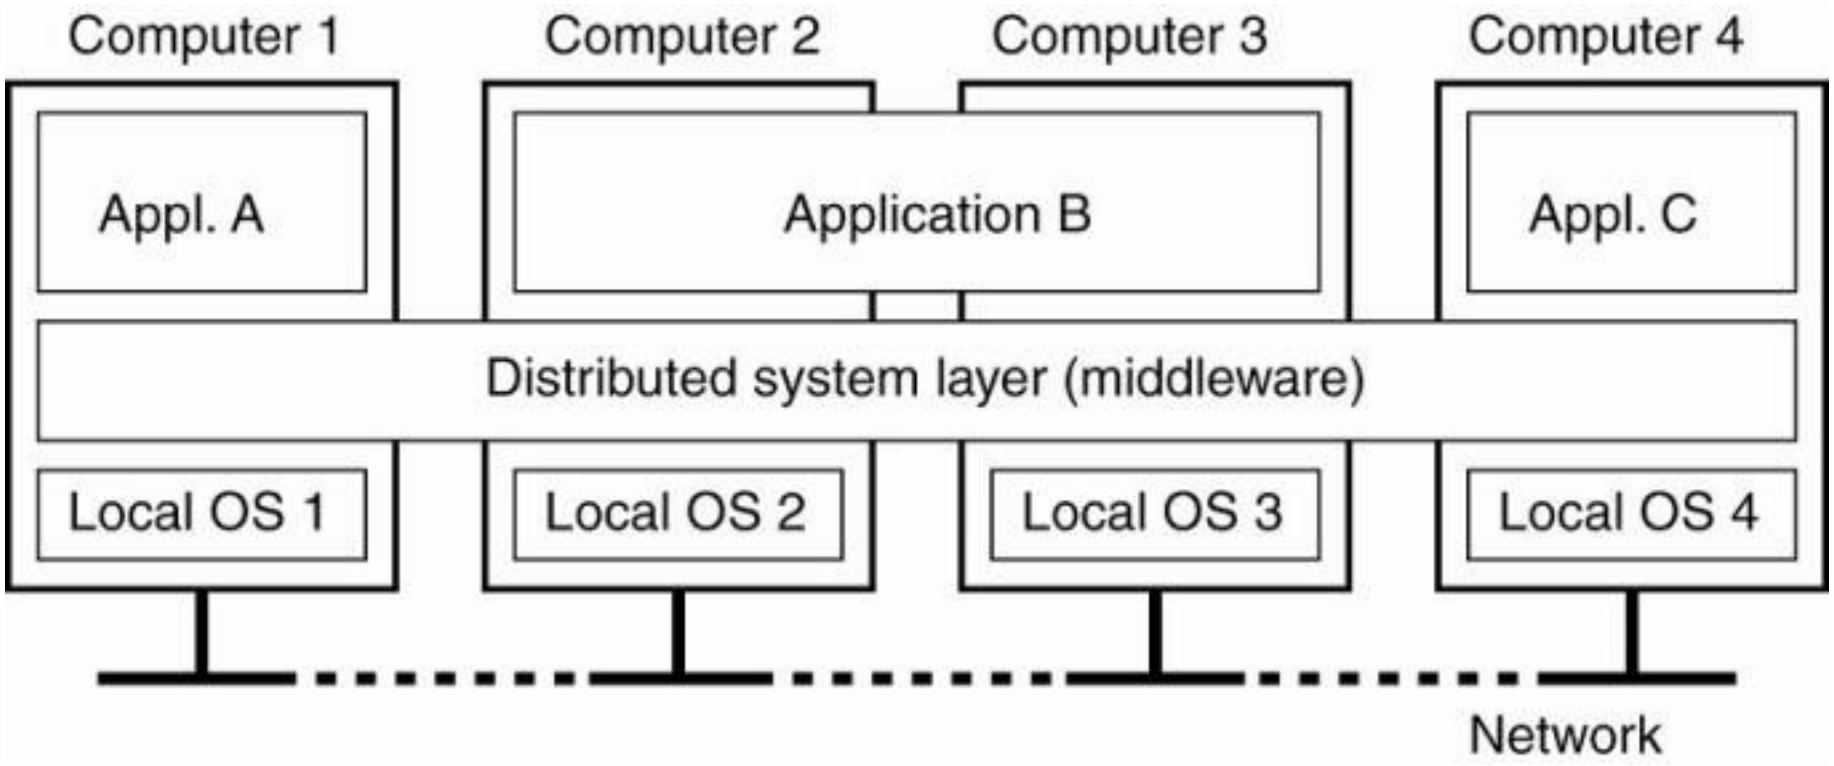

Multiple applications, Multiple servers, Networked together

Pretty much everywhere and everything computing now

Service with higher-level abstractions/interface (Dist. Databases, File-Systems, etc)

DISTRIBUTED SYSTEMS GOALS

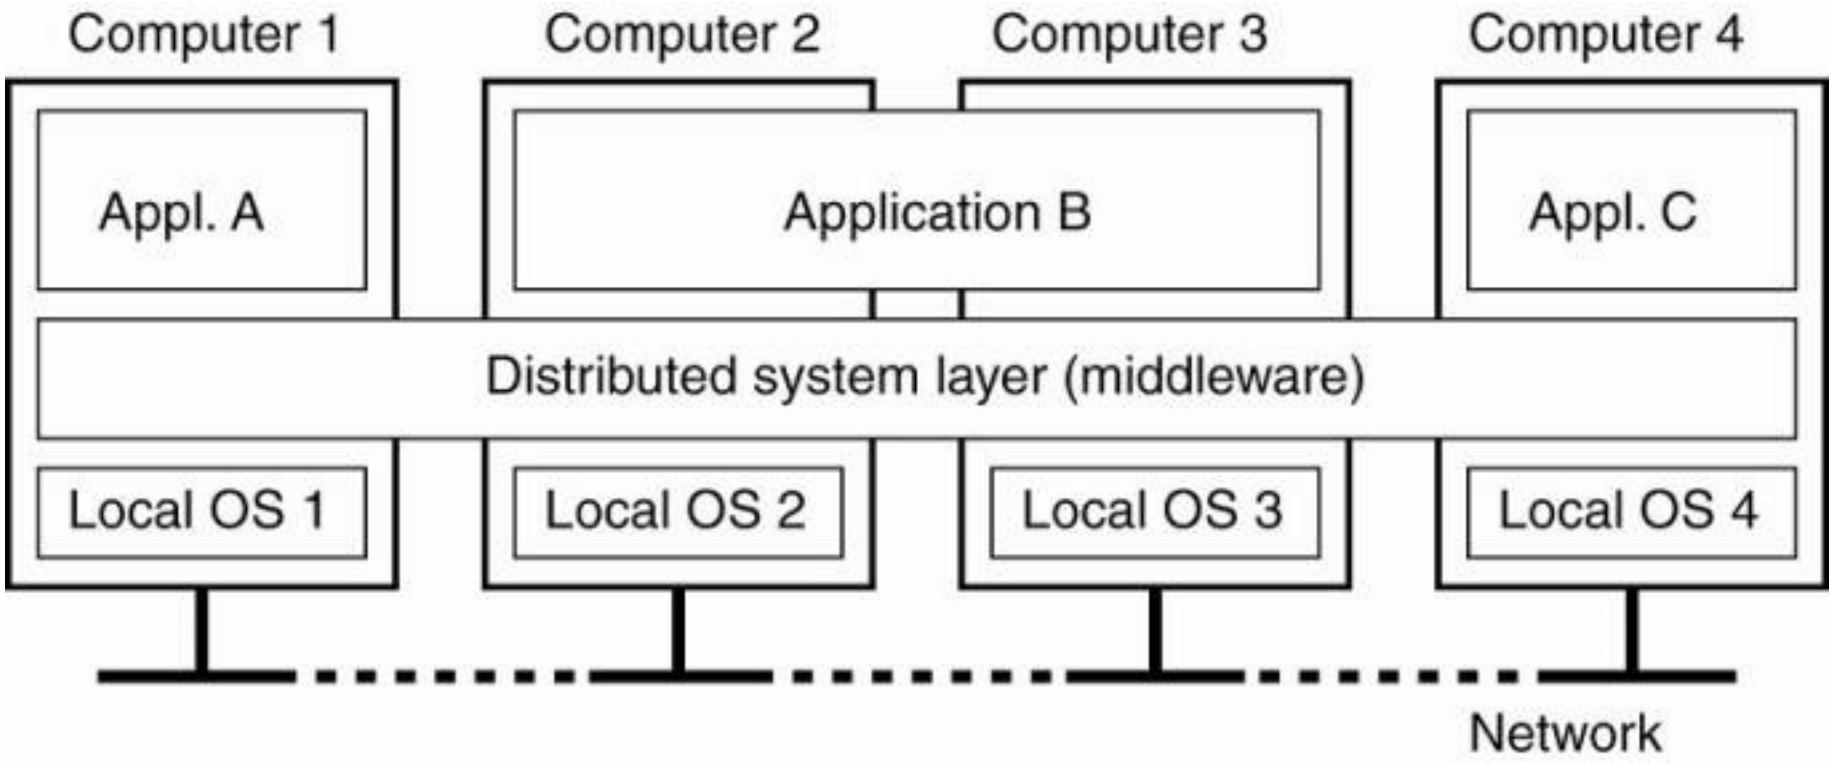

- **Scalability (Scale up/down size/volume)**
- **Consistency (Performance)**
- **Reliability (Fault Tolerance/Failure)**
- **Availability (No DNS)**
- **Complexity (Transparency)**

UNDERSTANDING DISTRIBUTED SYSTEMS GOALS

A simple webservice application

- Server Machine
- Client Borwsers
- Internet/Network

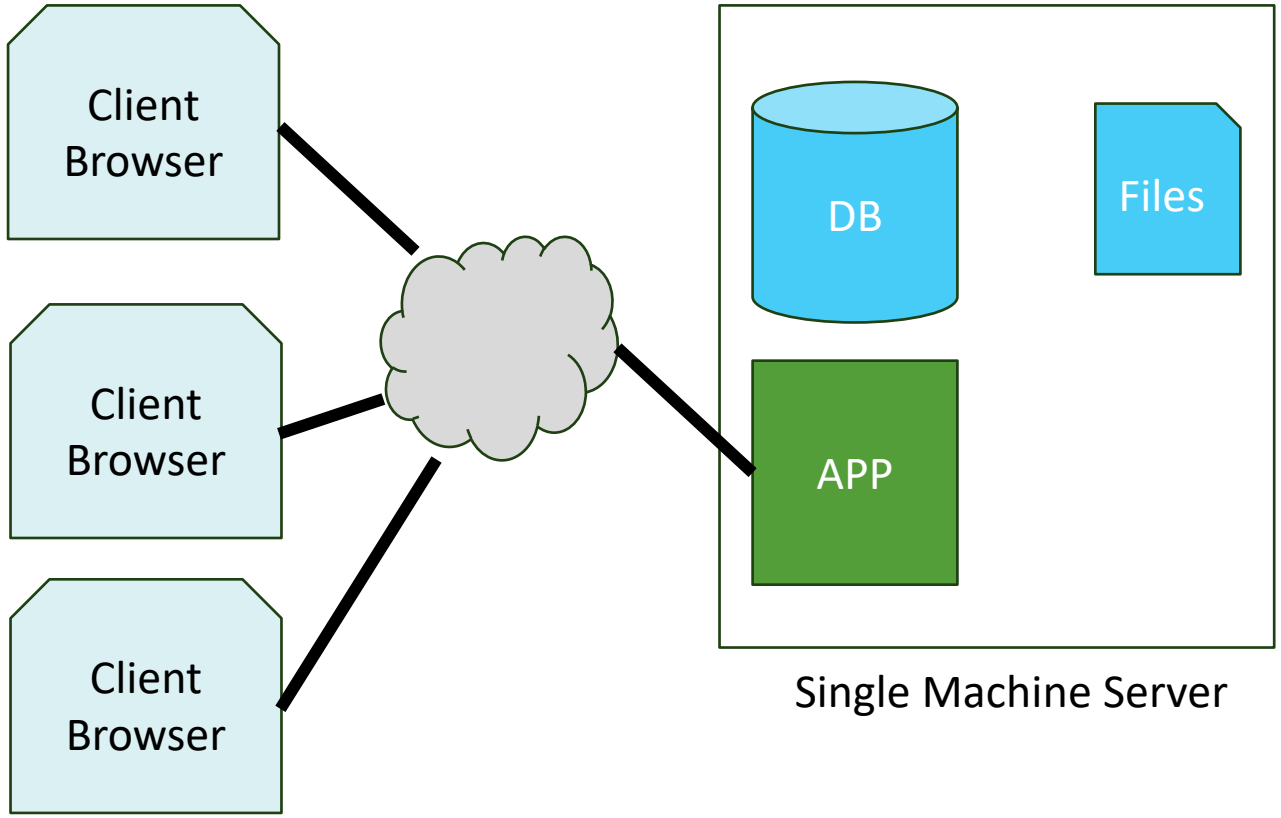

UNDERSTANDING DISTRIBUTED SYSTEMS GOALS

- Challenges
  - Increase File space
  - DB size?
  - APP size?
  - APP load?
  - # of Net access?

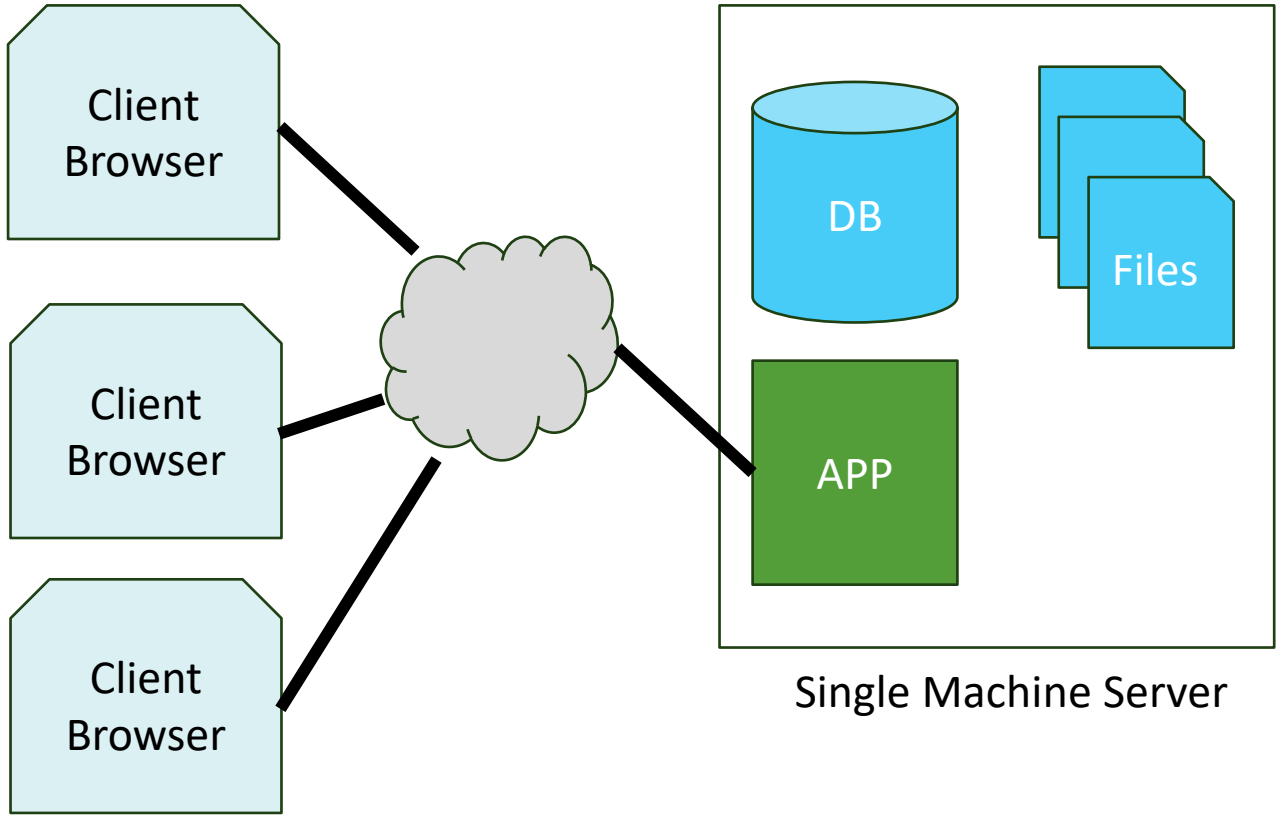

UNDERSTANDING DISTRIBUTED SYSTEMS GOALS

- Solution
  - 3 servers, each for APP, DB and Files
- Challenges
  - APP srvr is down? (Maintenance, Power-out etc) -> Availability
  - DB srvr is down? -> Data unavailability/ Data Durability
  - File srvr down?

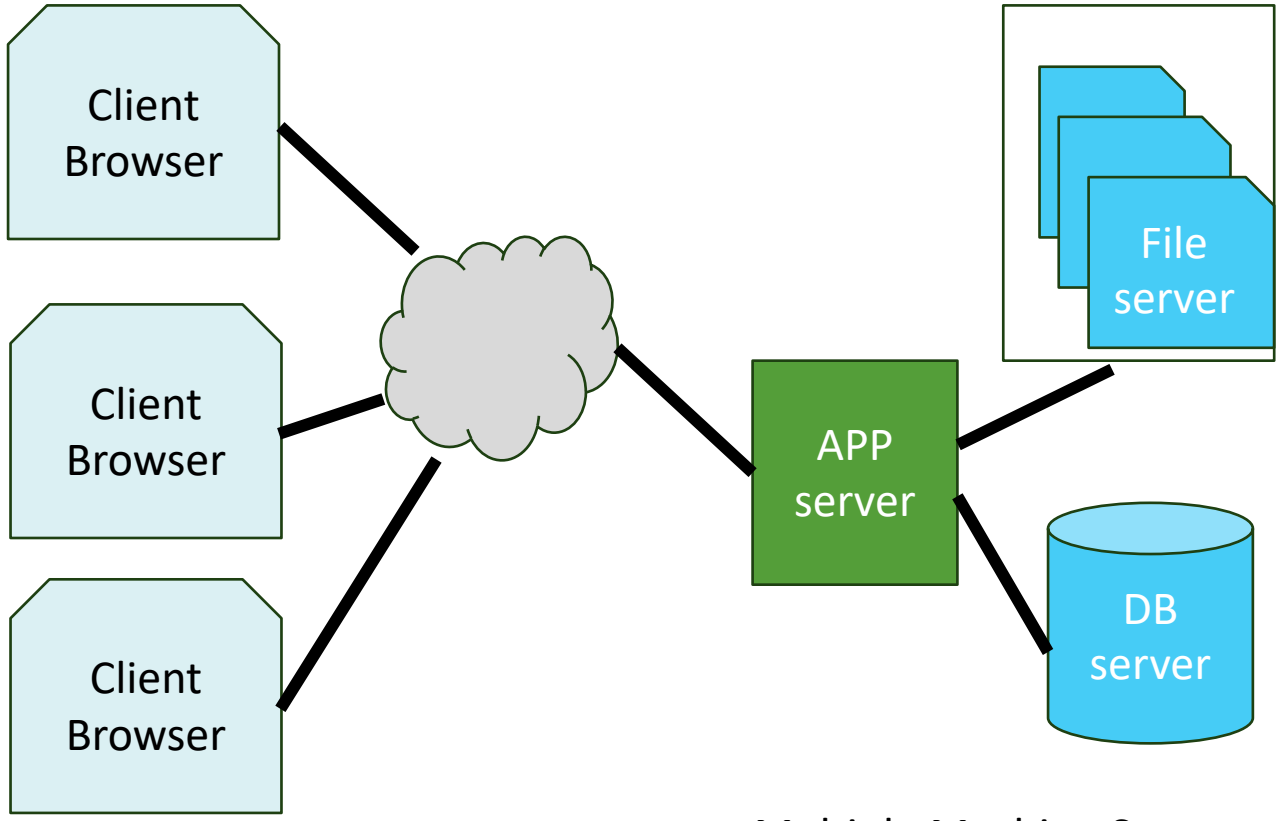

Multiple Machine Servers

UNDERSTANDING DISTRIBUTED SYSTEMS GOALS

- Solution
  - Add APP server(s), DB and Files
- Challenges
  - DNS problem?
  - Which APP srvr is primary?
  - How to balance load?

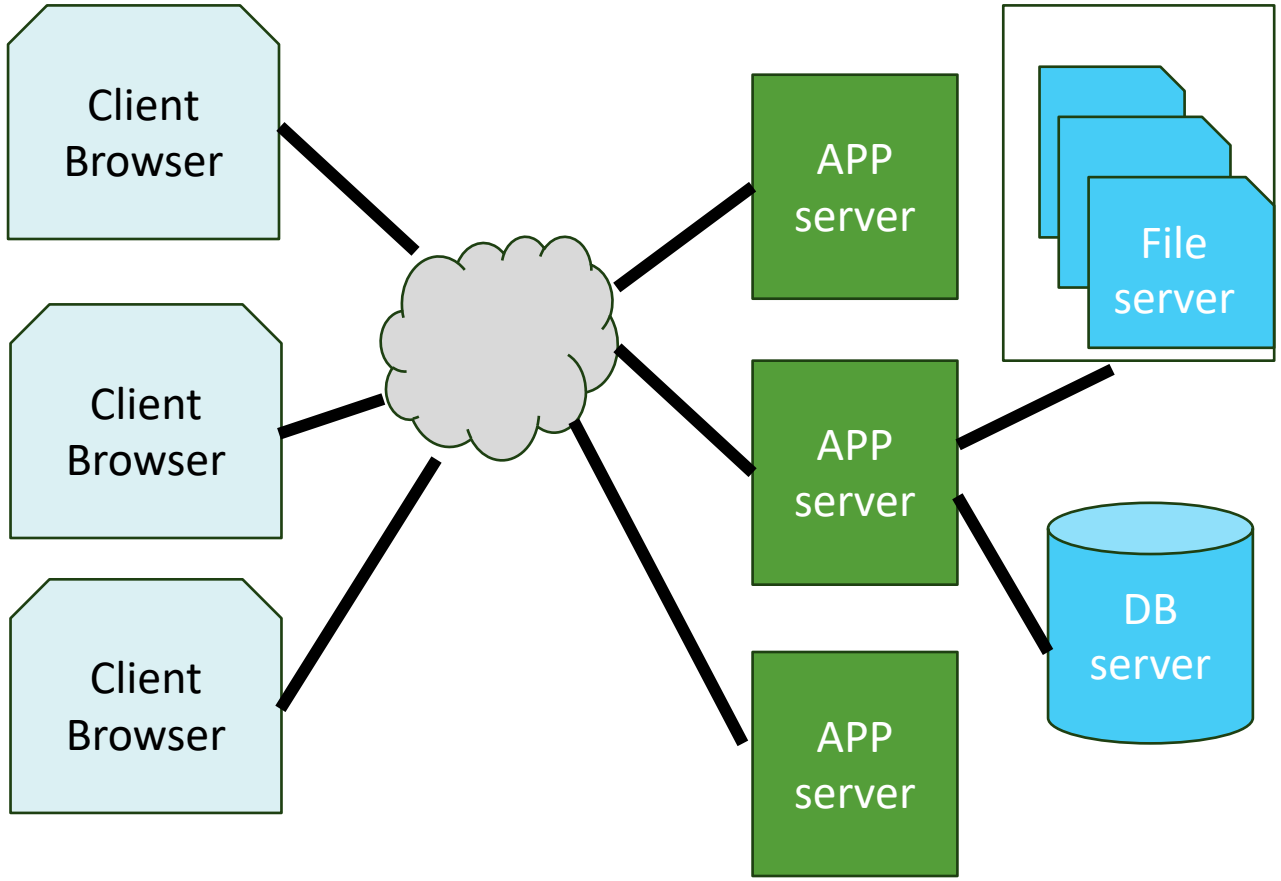

UNDERSTANDING DISTRIBUTED SYSTEMS GOALS

- Solution
  - Add Load Balancer
- Challenges
  - Each App srvr connects to DB
  - Multiple Access
  - Data integrity
  - Locks/Race-conditions
  - What if DB server crashes?

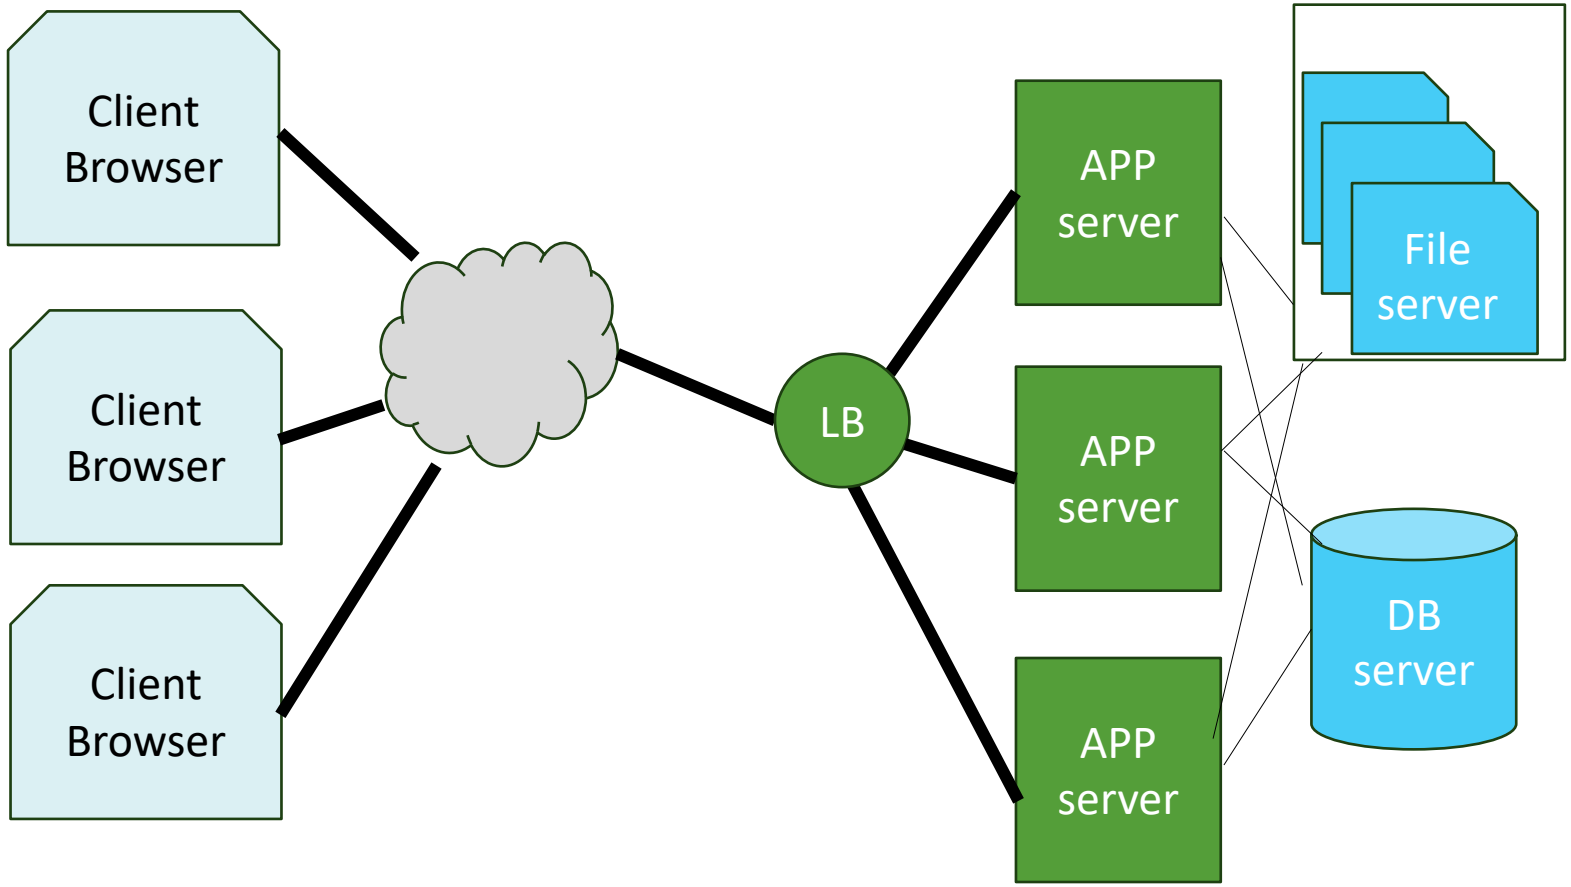

UNDERSTANDING DISTRIBUTED SYSTEMS GOALS

- Solution
  - Add DB servers
- Challenges
  - Which DB server if Primary? (All or one)
  - Master/Slave Arch
  - Read/Write issues
  - Load balancing
  - Caching?
  - Data replication?
  - Failure?

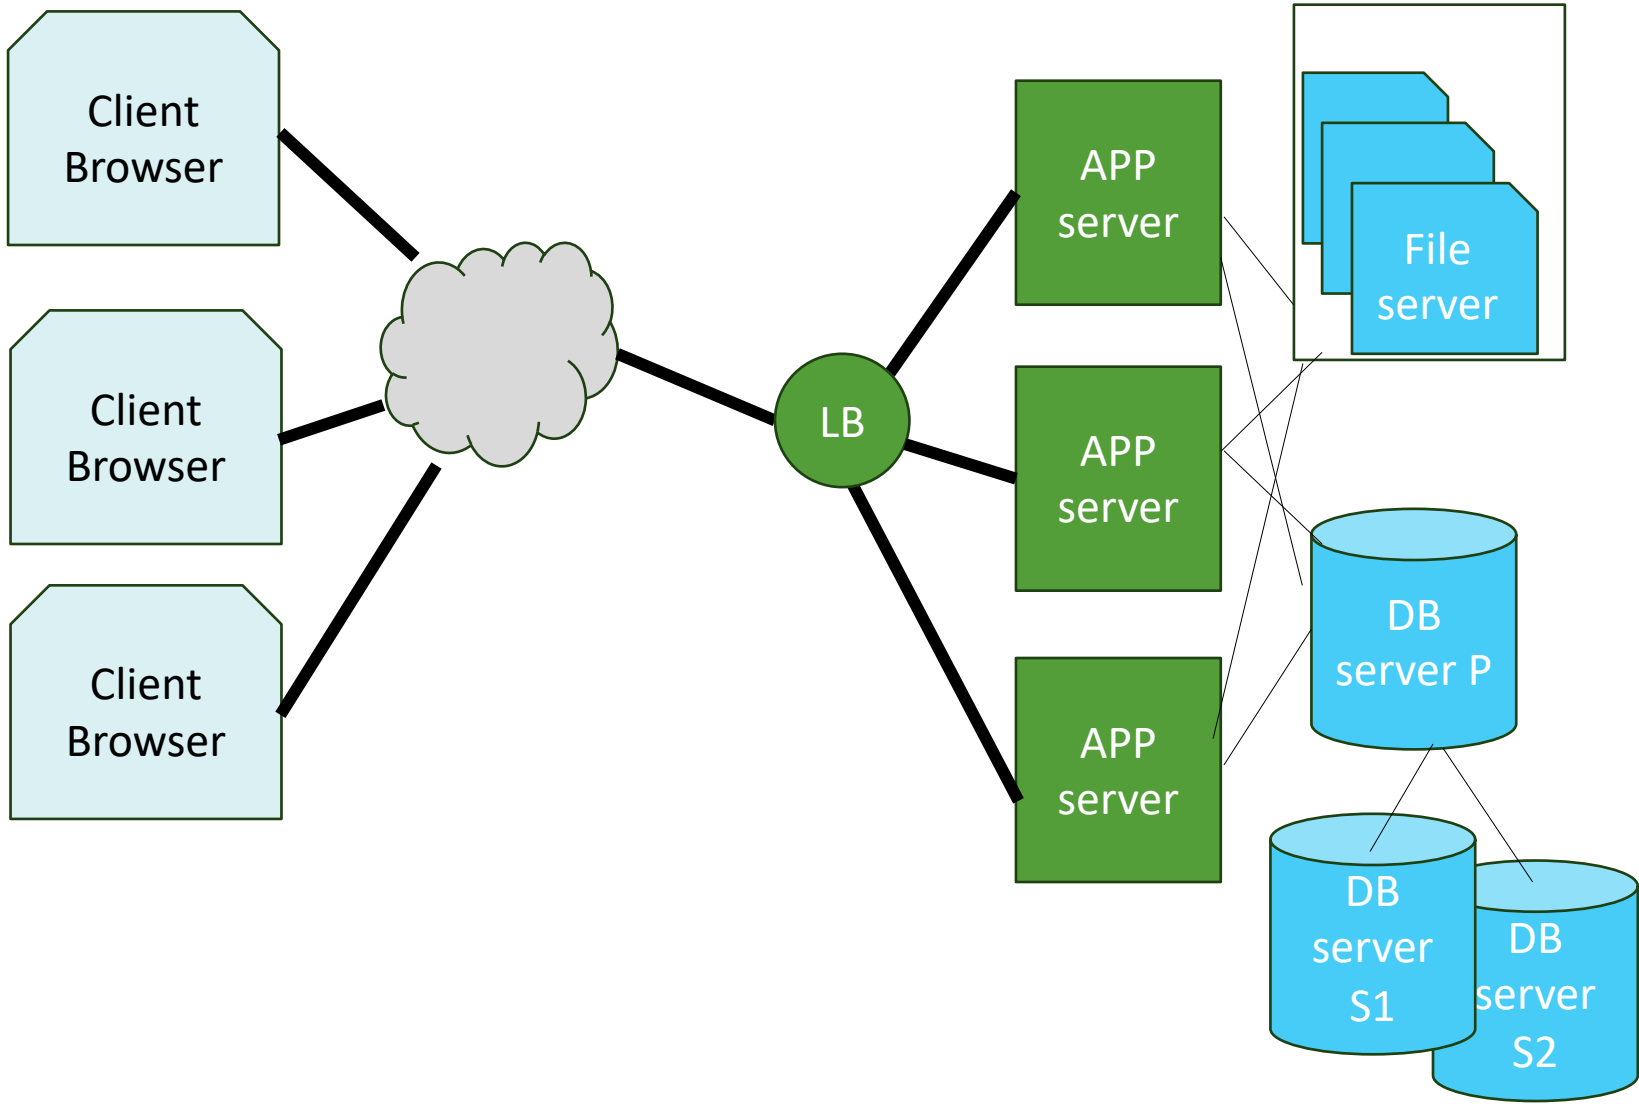

UNDERSTANDING DISTRIBUTED SYSTEMS GOALS

- Solution
  - Shard/Slice the DB servers
  - Logical Representation
- Challenges
  - What about File Server(s)?

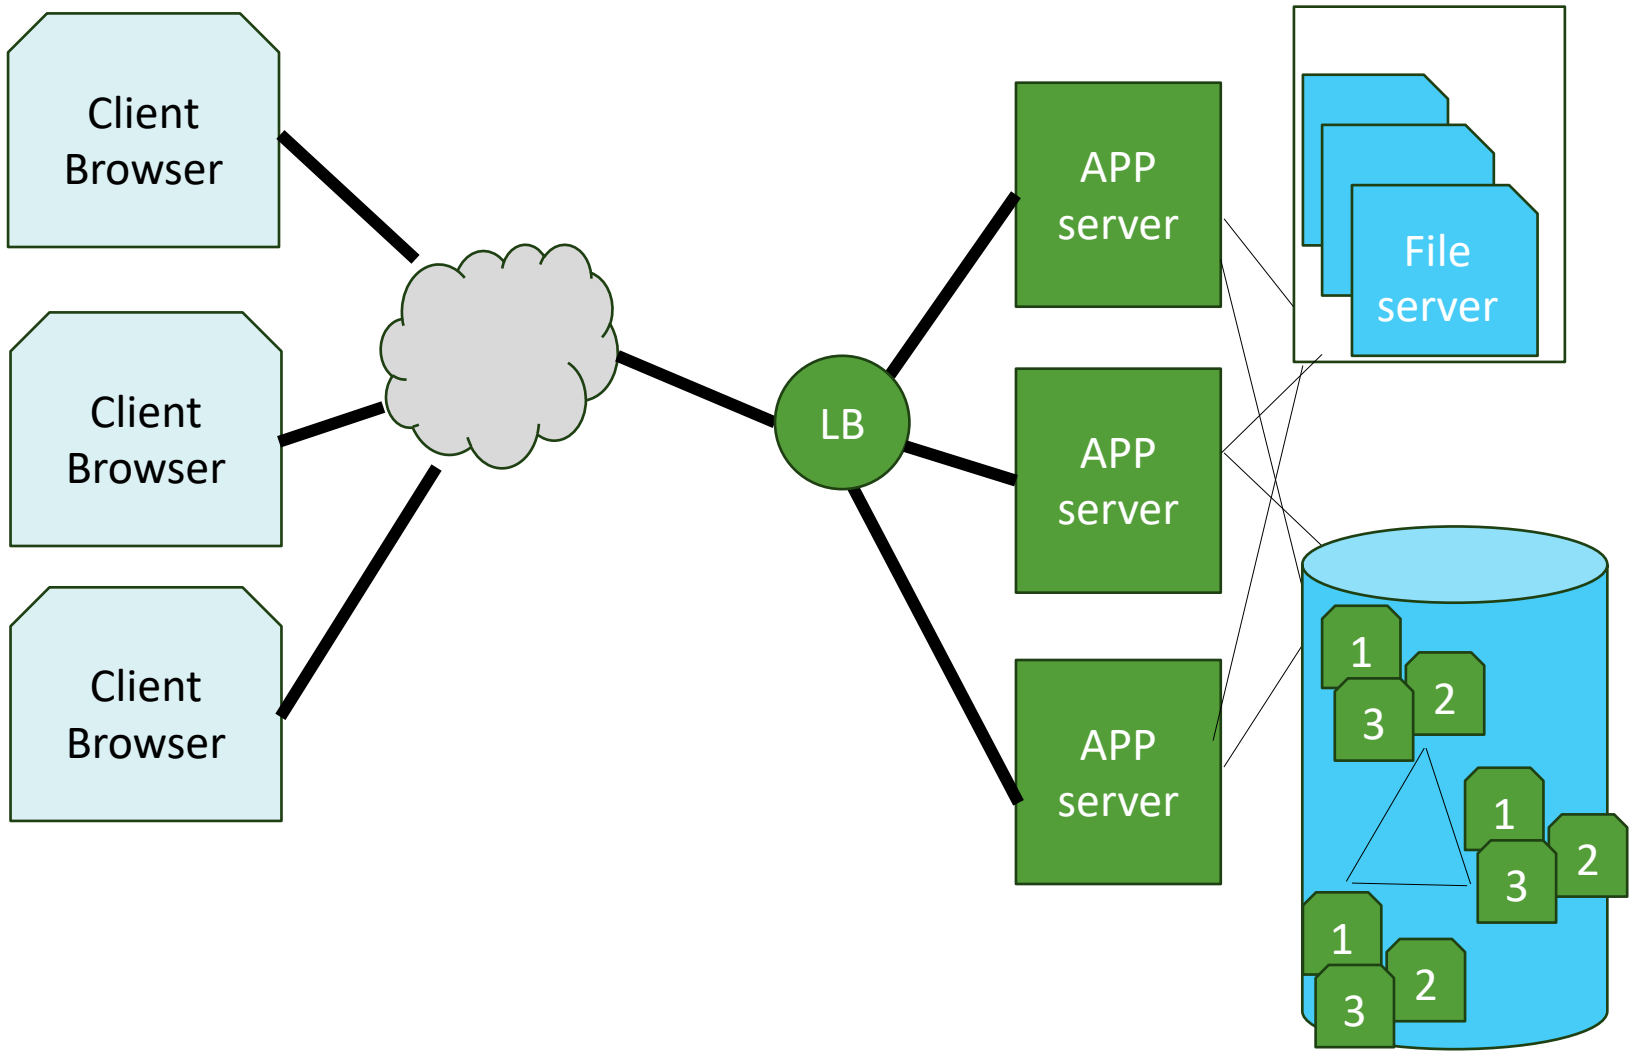

UNDERSTANDING DISTRIBUTED SYSTEMS GOALS

- Solution
  - Splices (RAIDs)
  - Logical Representation
- Solved problems
  - **Scalability**
  - **Availability**
  - **Fault Tolerance**
  - **Consistency**
  - **Transparency**
- Challenges
  - Performance, Complexity etc

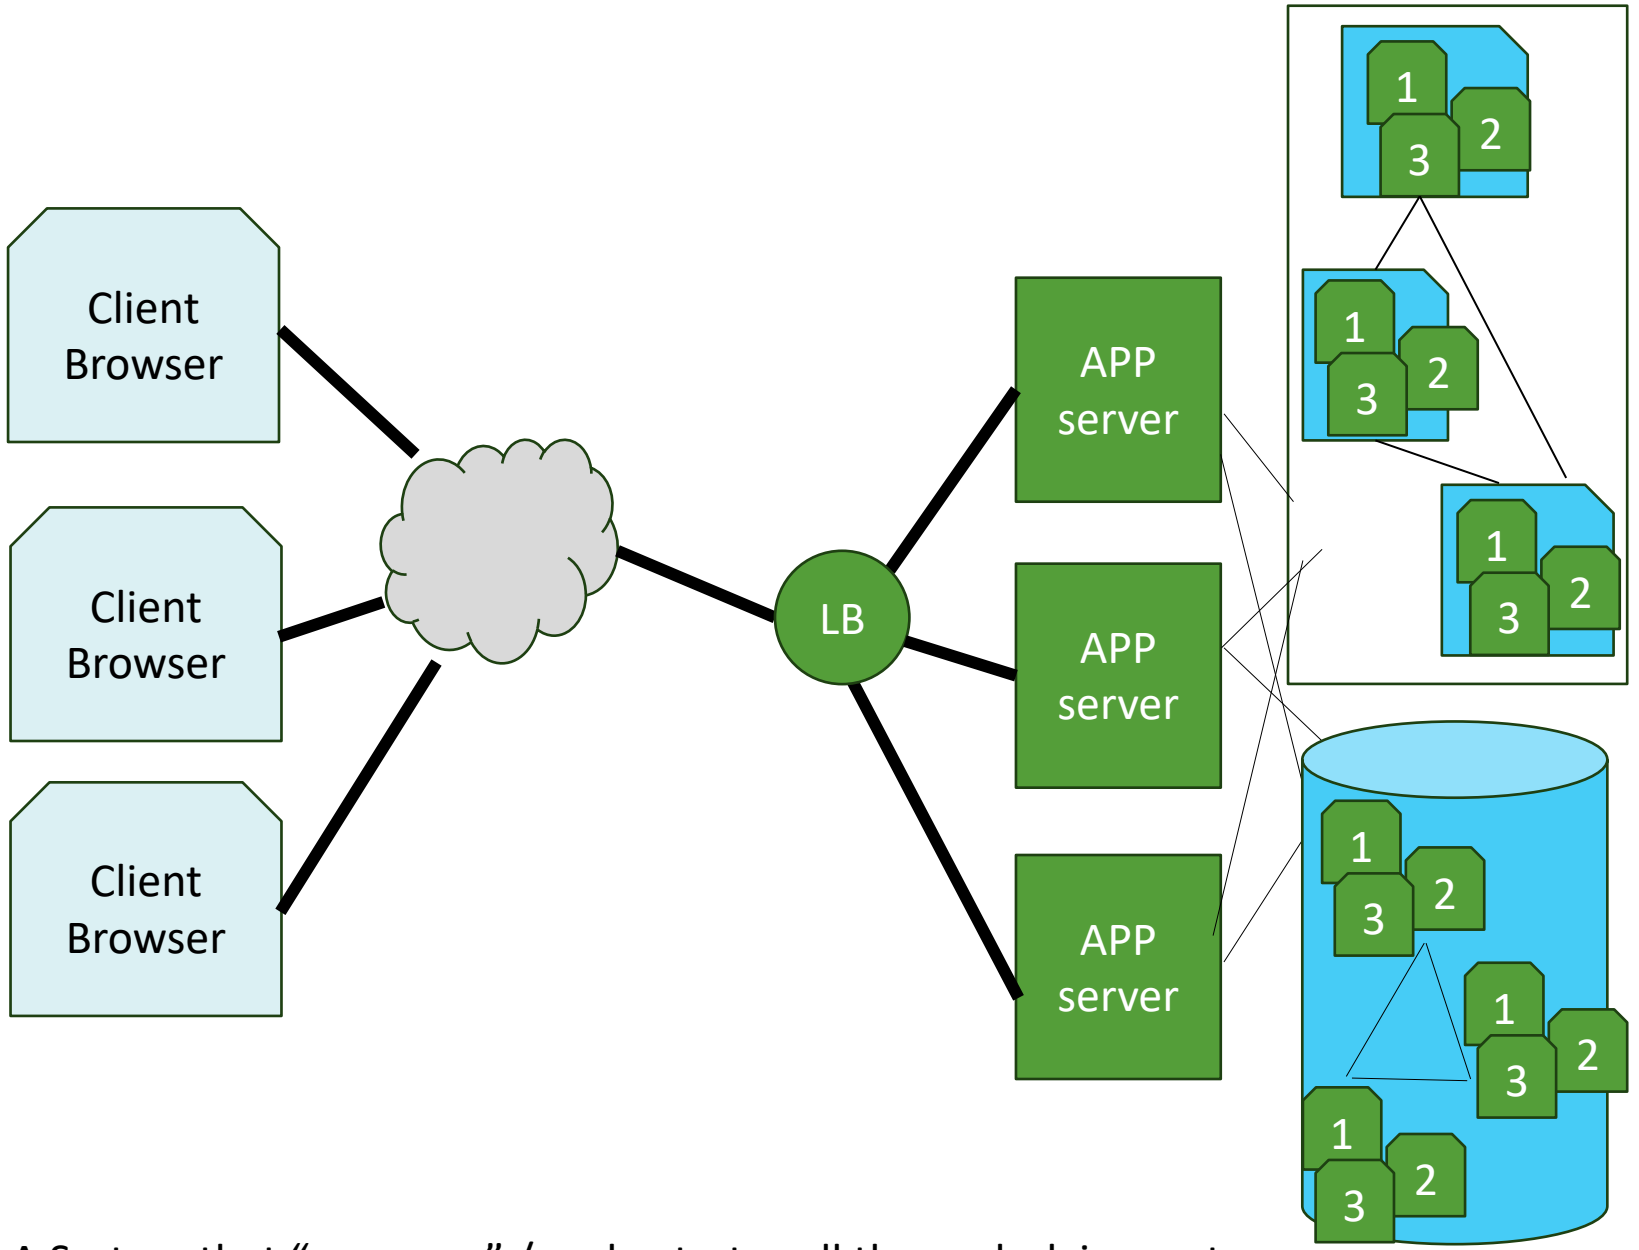

A System that “manages” / orchestrates all the underlying systems

UNDERSTANDING DISTRIBUTED SYSTEMS GOALS

- One size fits all?!!
  - No one solution
  - Many types of Dist Systems
- Known issues and pitfalls
  - **No global state (local decisions)**
  - **No global clock (decentralized)**
  - **Reliability, Security, Fault Tolerance, Latency, Cost**
- **Types of Dist Sys**
  - **High Performance Computing (HPC)/Cluster**
  - **Grid Systems**
  - **Cloud Systems**
  - **Transaction Processing Systems**

TYPES OF DISTRIBUTED SYSTEMS

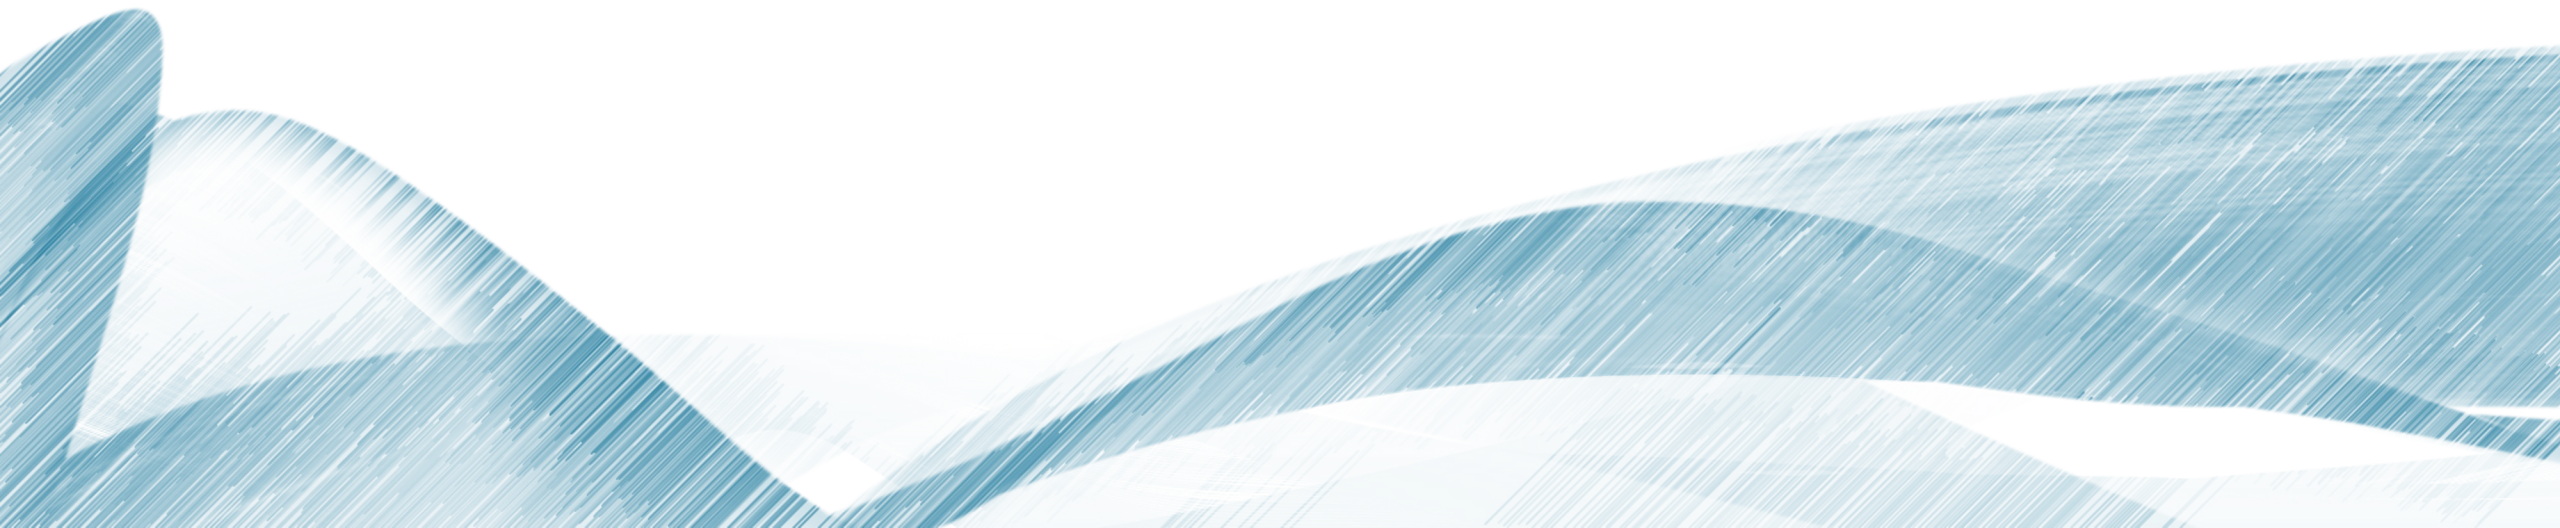

TYPES OF DISTRIBUTED SYSTEMS

- 1. High Performance Systems (Cluster)

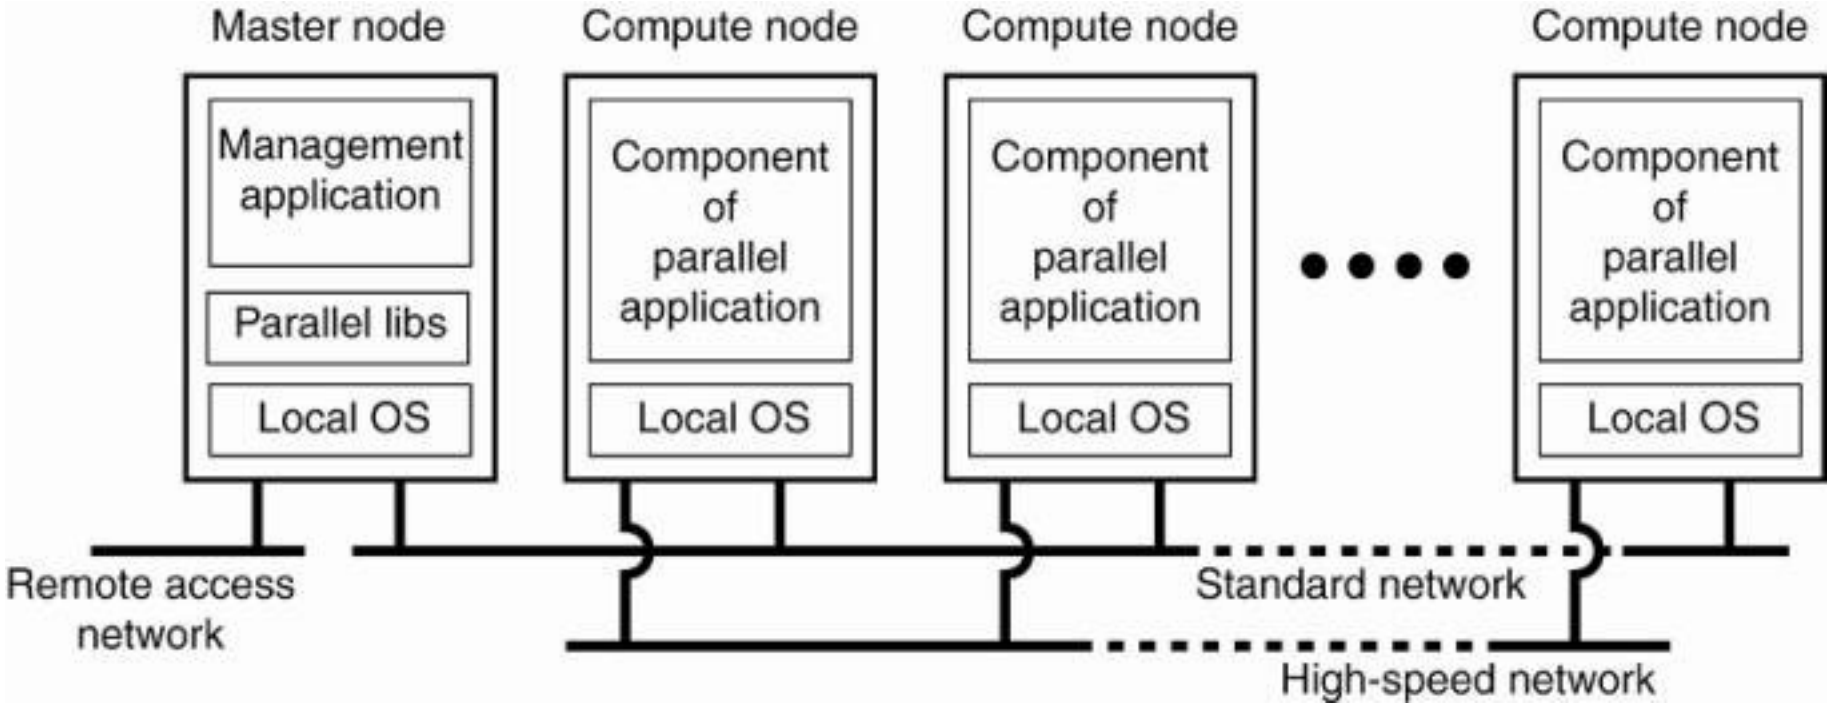

TYPES OF DISTRIBUTED SYSTEMS

- 1. High Performance Systems (Cluster)

- Similar computers
- High speed network
- Same OS on each node (e.g. Linux)
- 1 “Master” and several “slave” nodes
- Beowulf Cluster made of whitebox PCs

https://en.wikipedia.org/wiki/Beowulf_cluster

TYPES OF DISTRIBUTED SYSTEMS

- 1. High Performance Systems (Cluster)

- PSU RPI Cluster
- 40 Raspberry Pis
- High speed network
- Same OS on each node (Raspian)
- 1 “Master” and 39 “slave” nodes

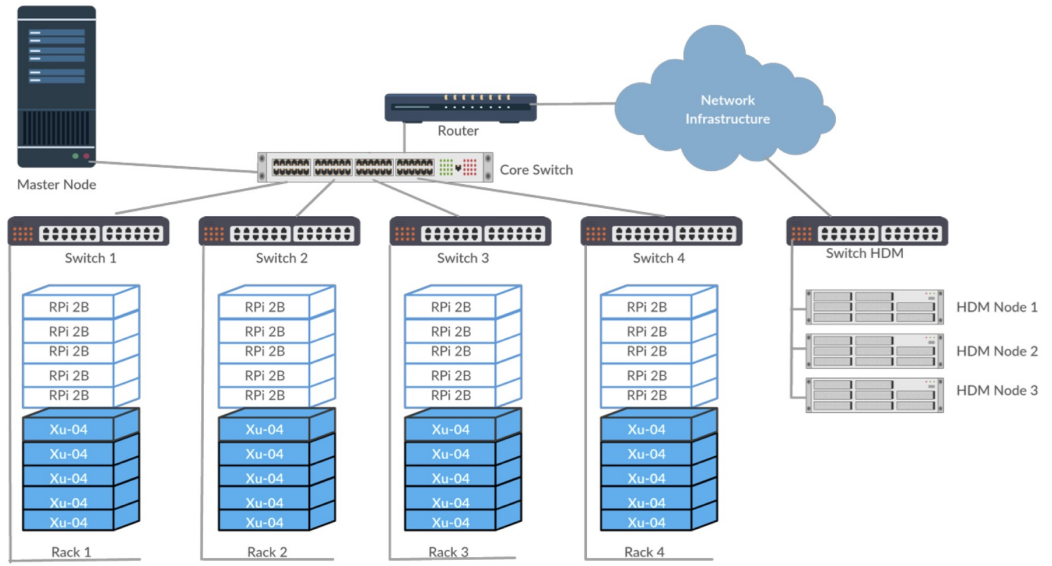

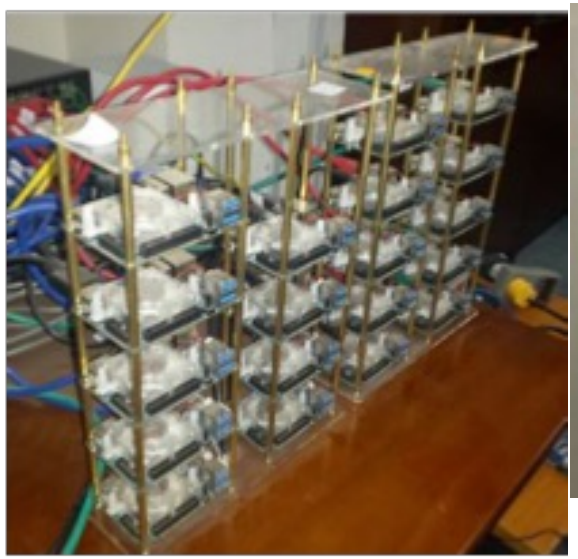

Qureshi, B.; Koubaa, A. On Energy Efficiency and Performance Evaluation of Single Board Computer Based Clusters: A Hadoop Case Study. *Electronics* **2019**, *8*, 182. https://doi.org/10.3390/electronics8020182

TYPES OF DISTRIBUTED SYSTEMS

- 2. Cloud Computing Systems: A data center hardware and software that the vendors use to offer the computing resources and services

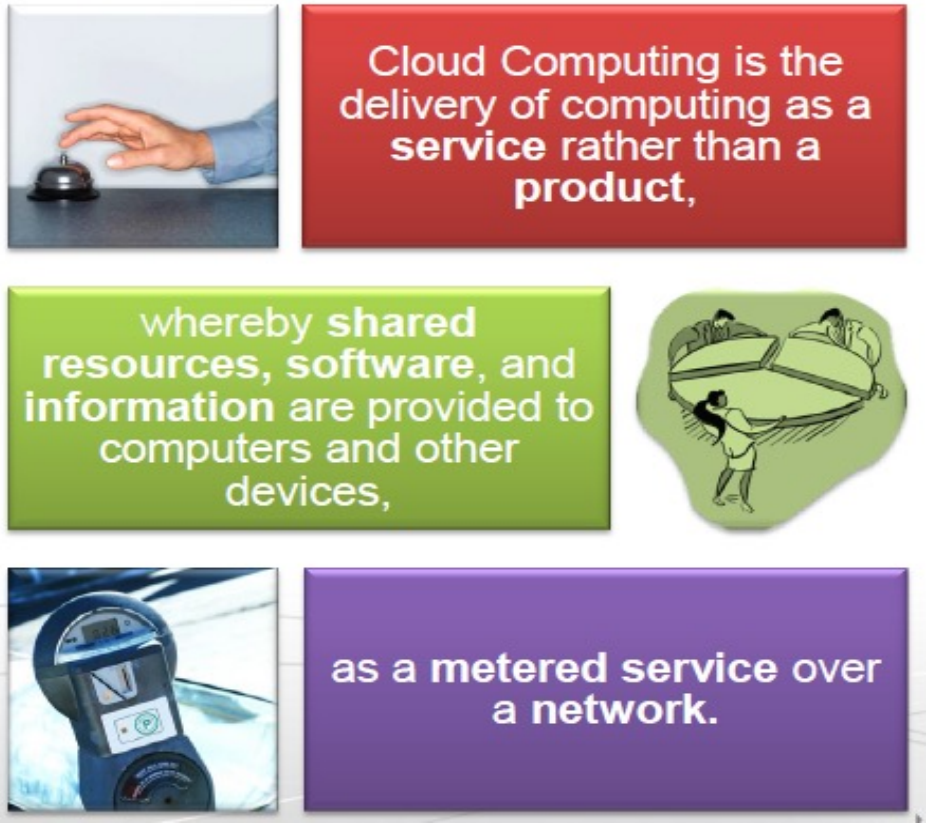

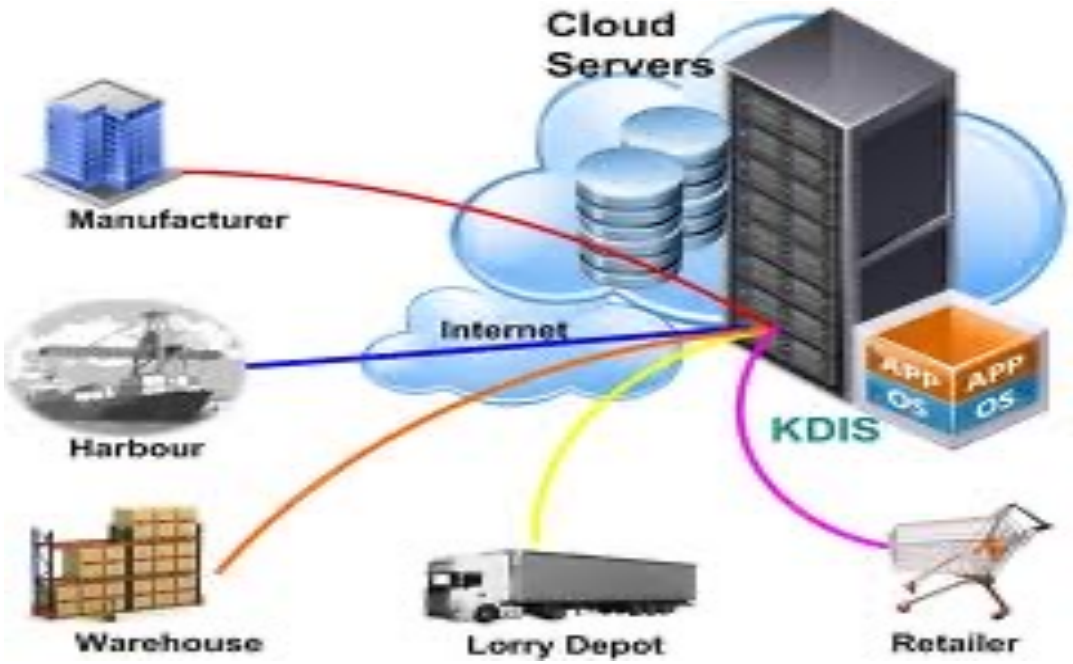

TYPES OF DISTRIBUTED SYSTEMS

Cloud computing means **selling "X as a service"**

IaaS: Infrastructure as a Service

- Selling virtualized hardware

PaaS: Platform as a service

- Access to a configurable platform/API

SaaS: Software as a service

- Software that runs on top of a cloud

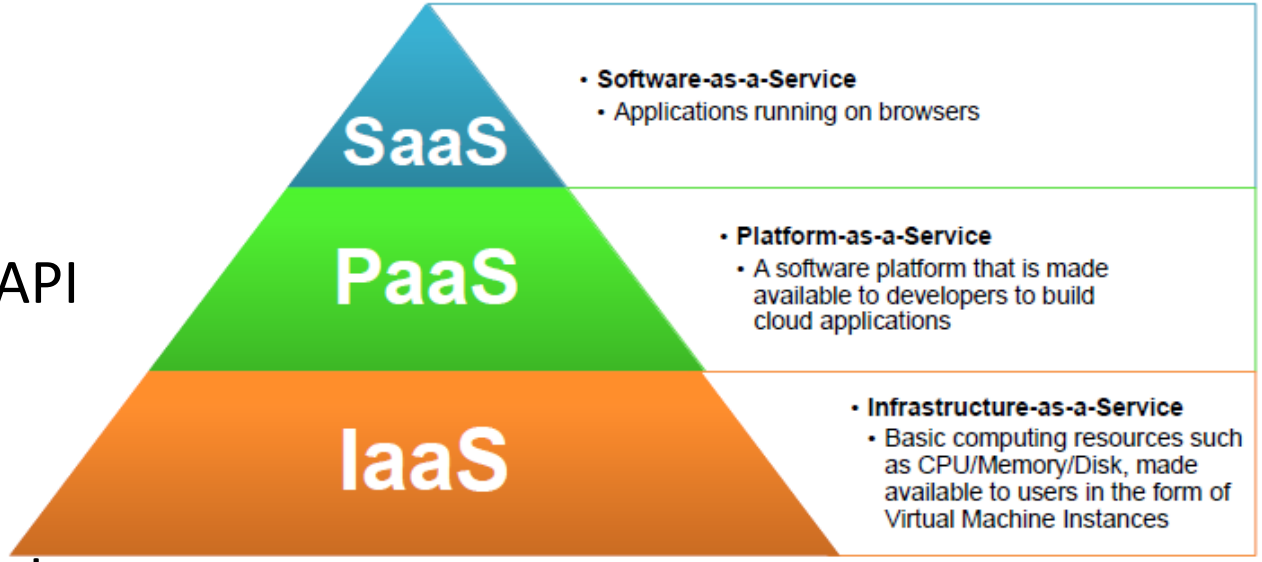

TYPES OF DISTRIBUTED SYSTEMS

SaaS

- You are most familiar with this!
- Software is delivered as a service over the Internet, eliminating the need to install and run the application on the customer's own computer
- This simplifies maintenance and support
- Examples: Gmail, YouTube, and Google Docs, among others

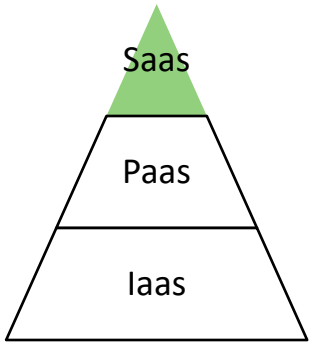

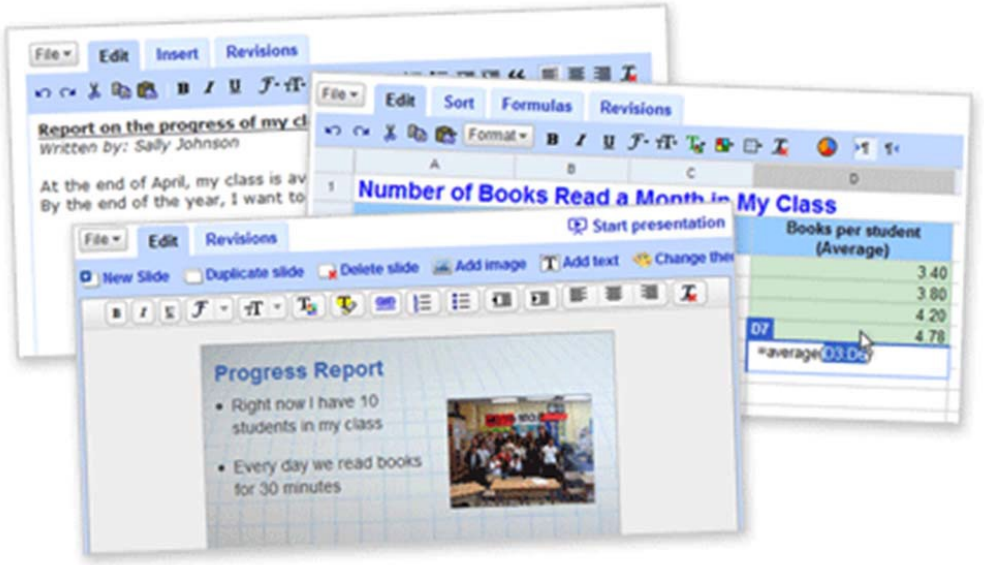

TYPES OF DISTRIBUTED SYSTEMS

PaaS

- The Cloud provider exposes a set of tools (a platform) which allows users to create SaaS applications
- The SaaS application runs on the provider's infrastructure
- The cloud provider manages the underlying hardware and requirements

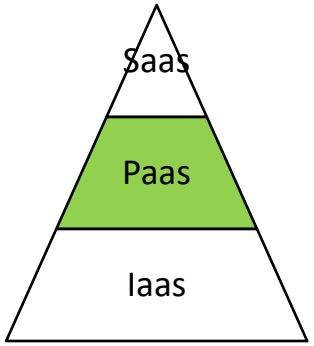

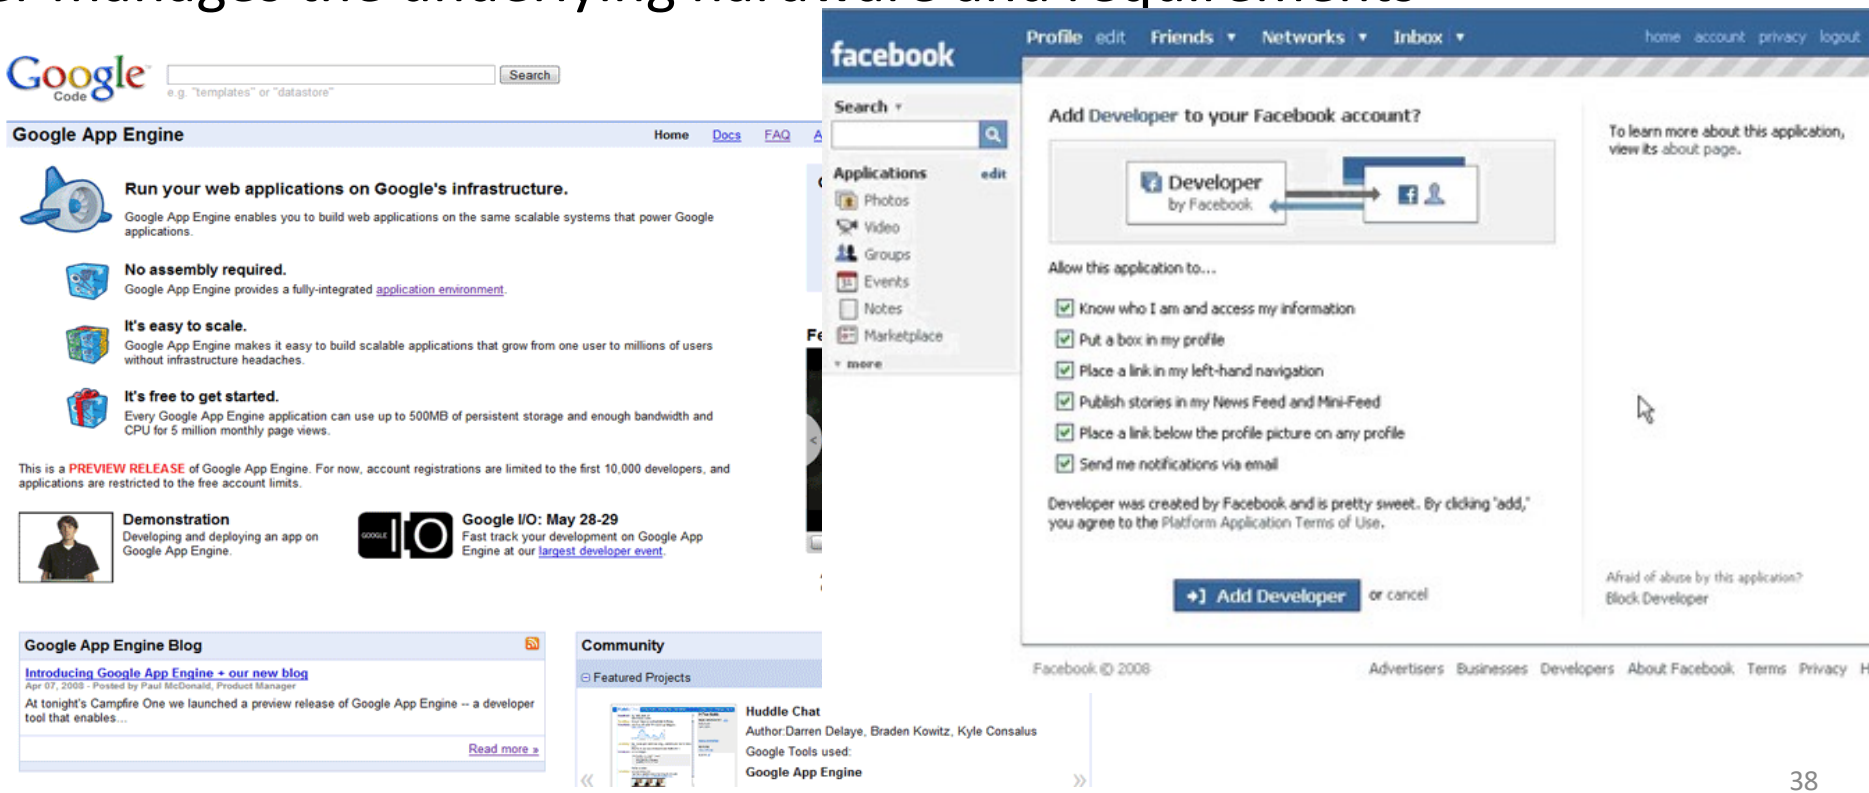

TYPES OF DISTRIBUTED SYSTEMS

IaaS

- The cloud provider leases to users Virtual Machine Instances (i.e., computer infrastructure) using the ***virtualization*** technology
- The user has access to a standard Operating System environment and can install and configure all the layers above it

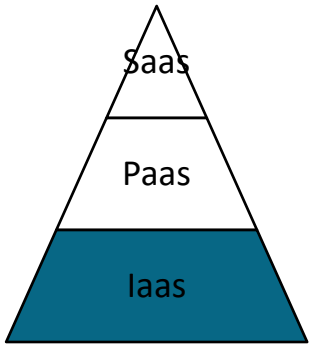

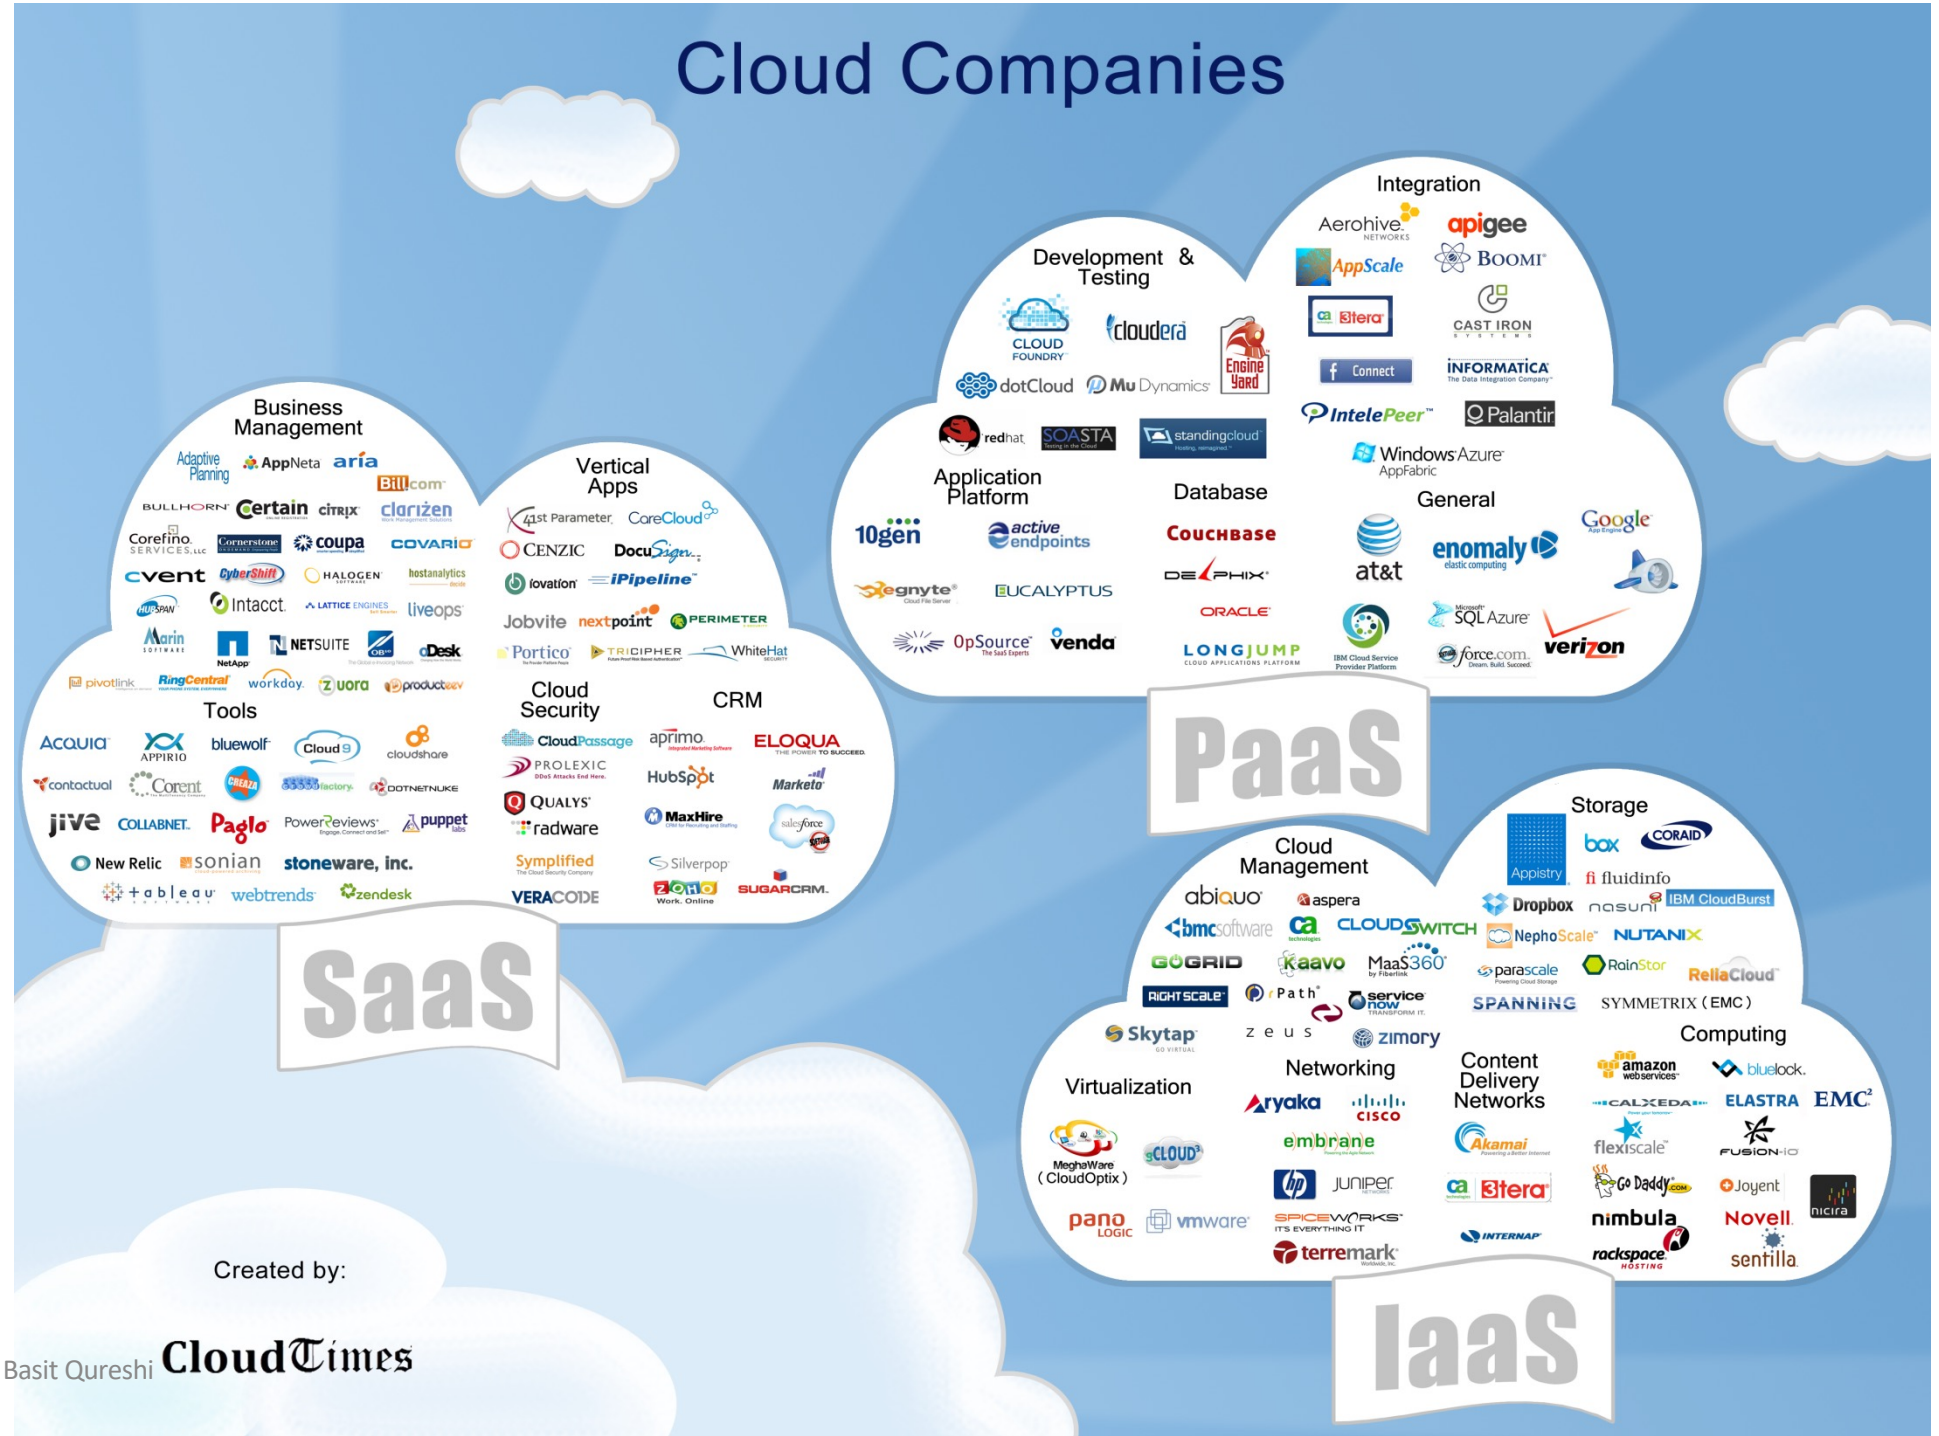

TYPES OF DISTRIBUTED SYSTEMS

- 3. Grid Computing Systems
  - Combines computer resources spread over different geographical locations to achieve a common goal.
  - All unused resources on multiple computers are pooled together and made available for a single task.
  - Perform large tasks or solve complex problems.

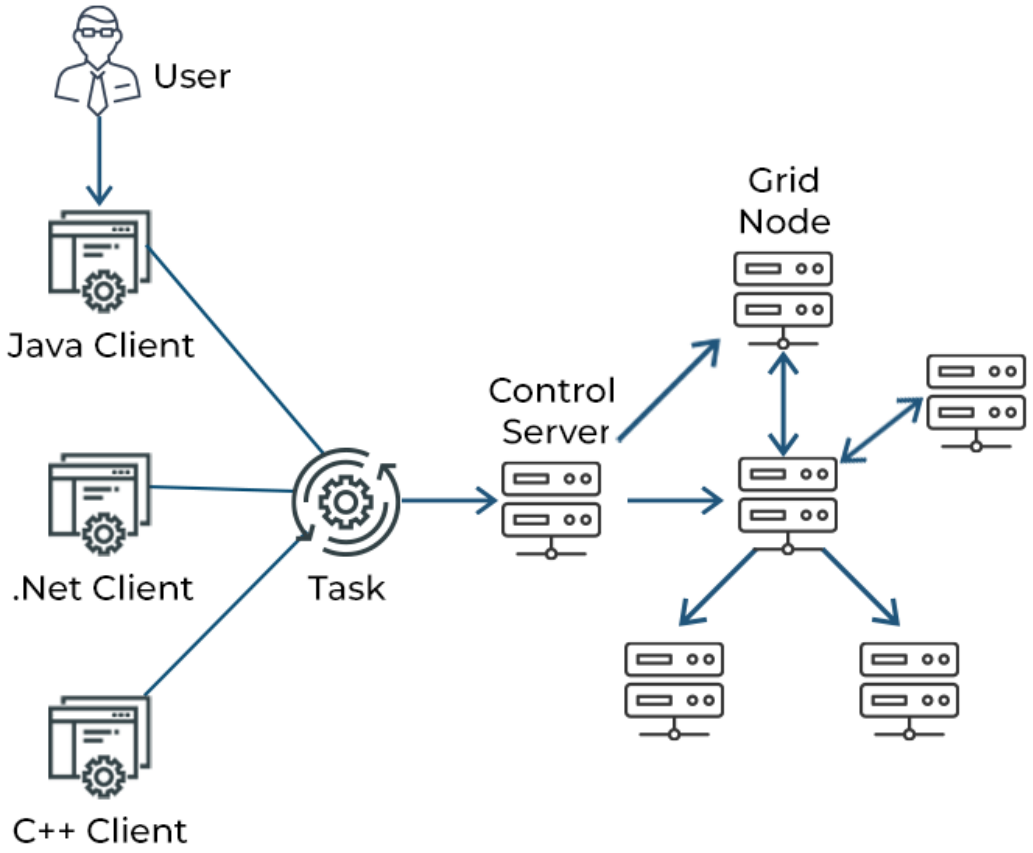

OSG

The Open Science Data Federation (OSDF)

https://osg-htc.org/services/osdf.html

TYPES OF DISTRIBUTED SYSTEMS

- 3. Grid Computing Systems
- 3 machine types:
  - Control node/server: A control node is a server or a group of servers that administers the entire network and maintains the record for resources in a network pool.
  - Provider/grid node: A provider or grid node is a computer that contributes its resources to the network resource pool.
  - User: A user refers to the computer that uses the resources on the network to complete the task.

https://www.spiceworks.com/tech/cloud/articles/what-is-grid-computing/

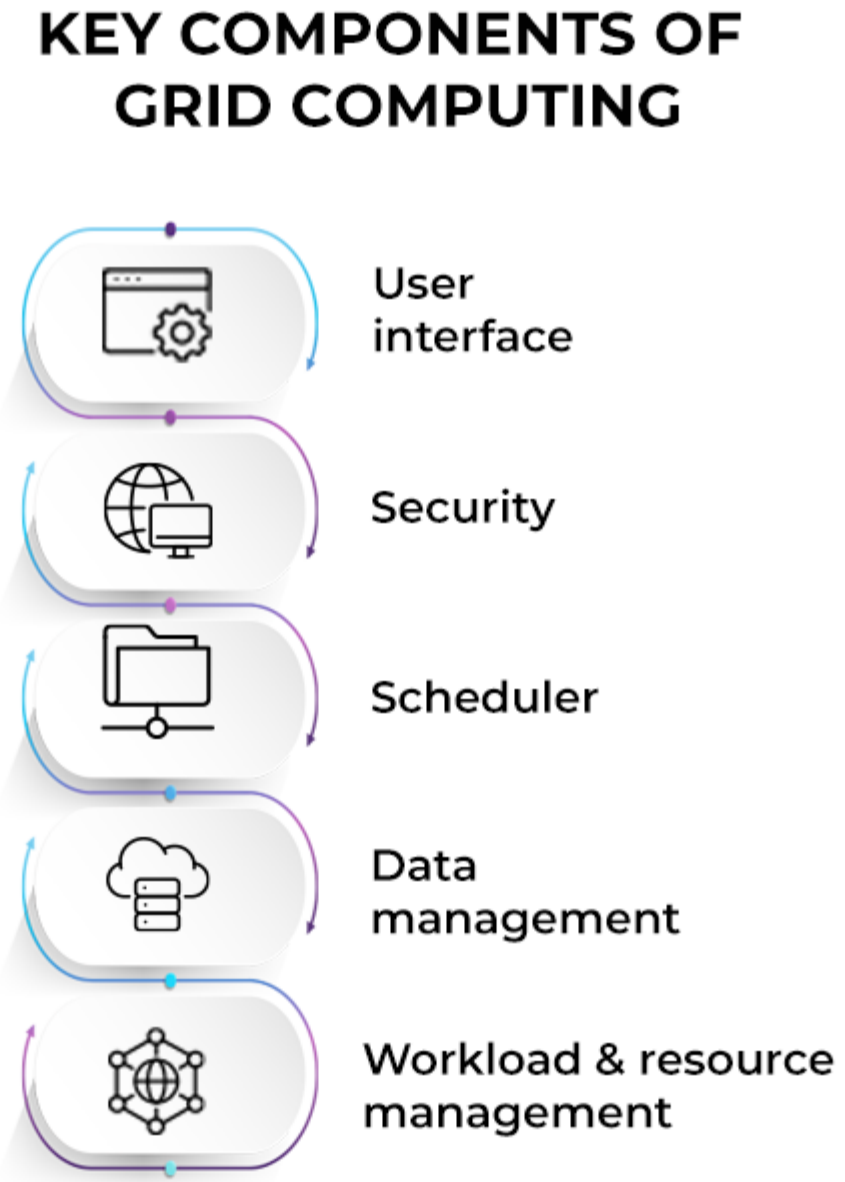

TYPES OF DISTRIBUTED SYSTEMS

- 3. Grid Computing Systems
- Examples:

European Grid Infrastructure (EGI) for research

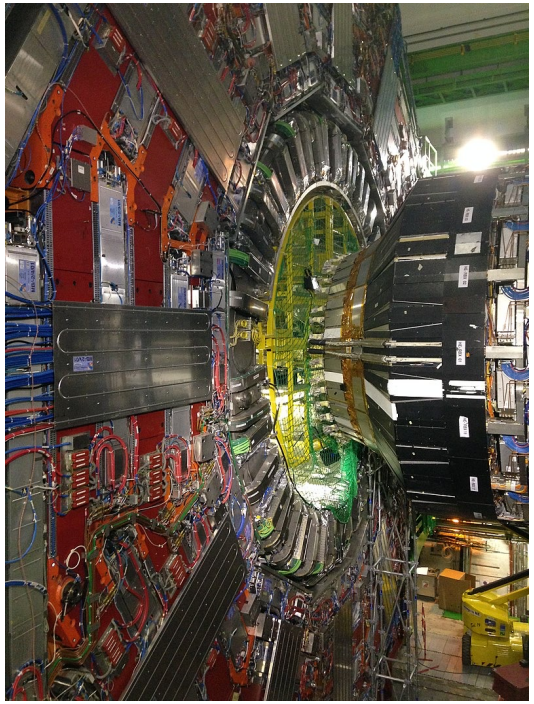

CMC detector for the Large Hadron Collider (CERN)

neuGRID is a web portal aimed to help neuroscientists do high-throughput imaging research

DIST SYS APPLICATIONS

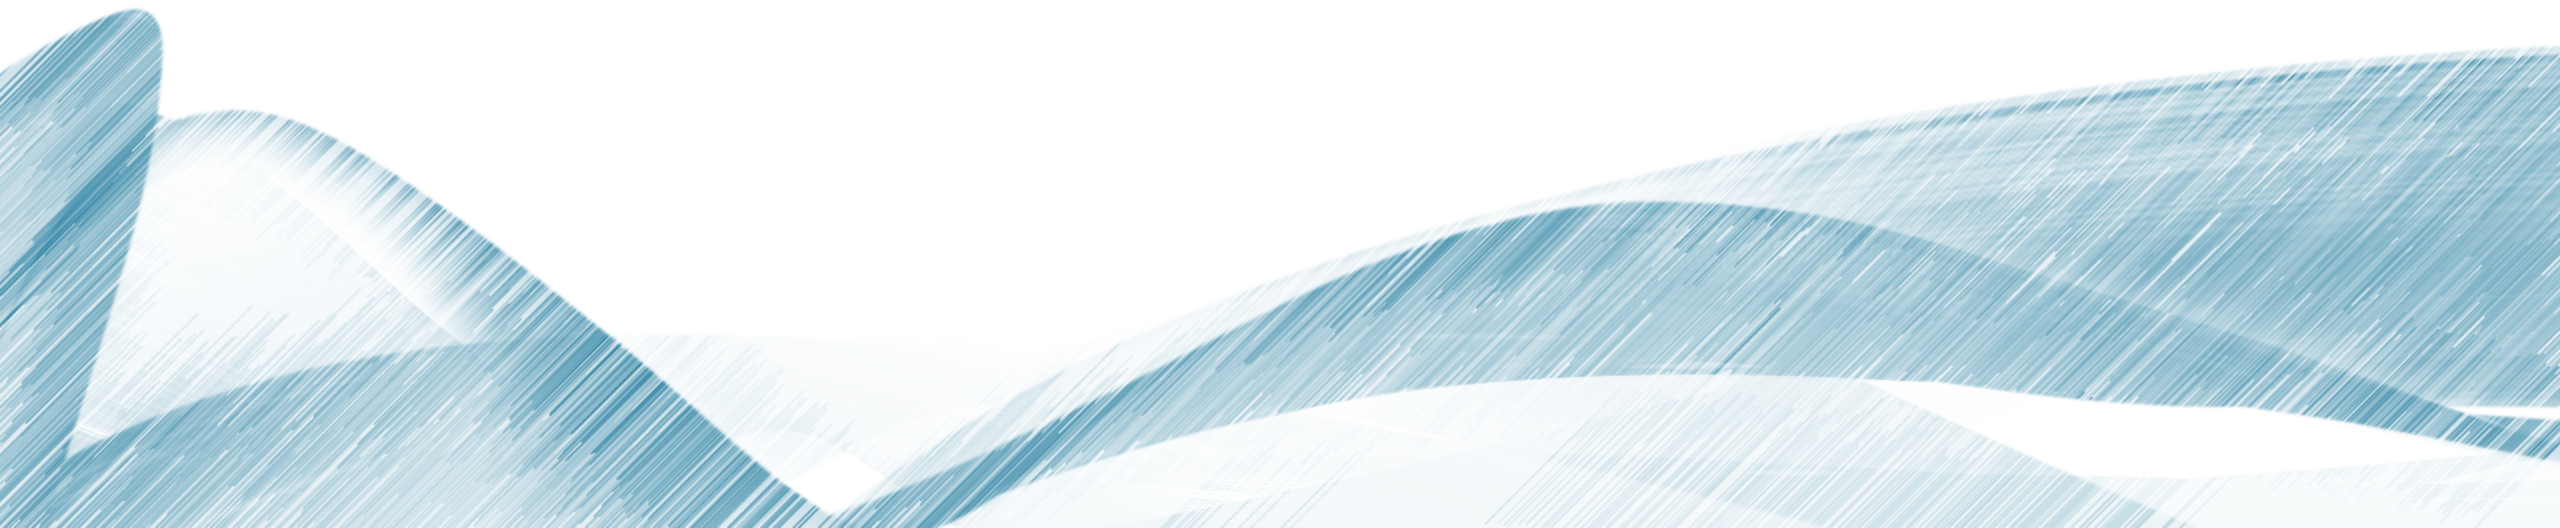

APPLICATIONS OF DISTRIBUTED SYSTEMS

- Transaction Processing Systems (TPS)

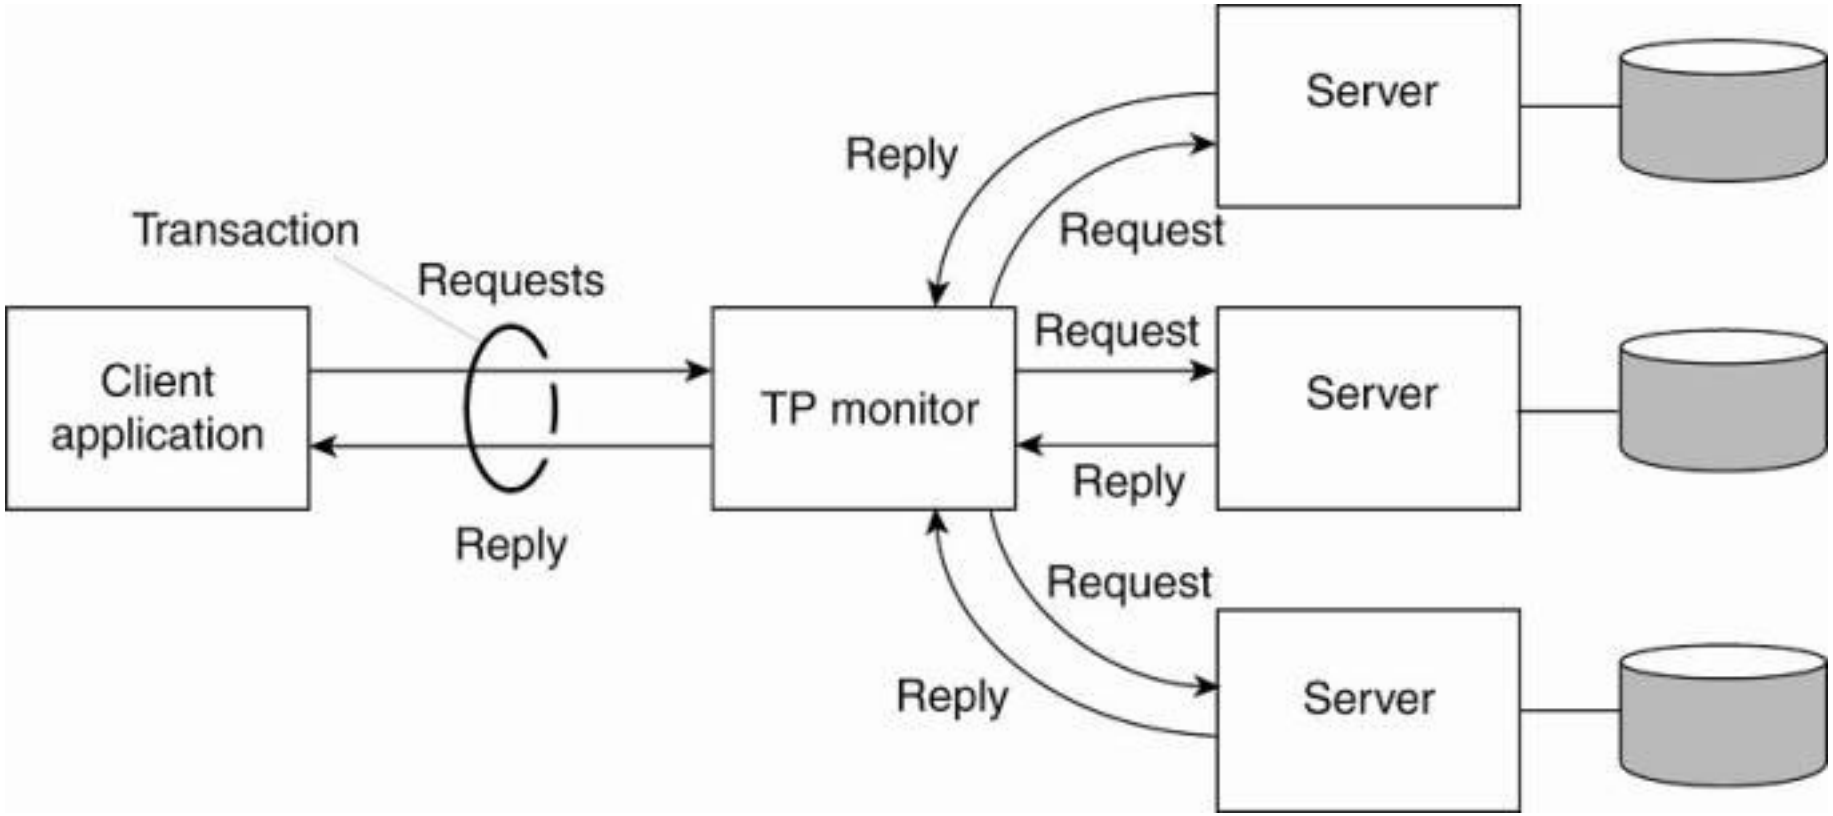

APPLICATIONS OF DISTRIBUTED SYSTEMS

- Peer to peer systems

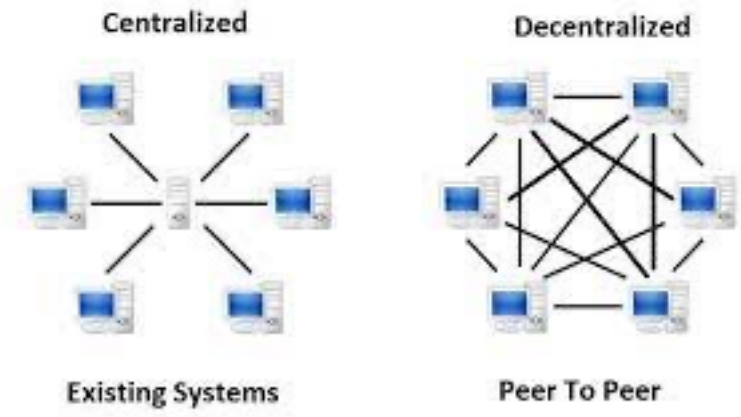

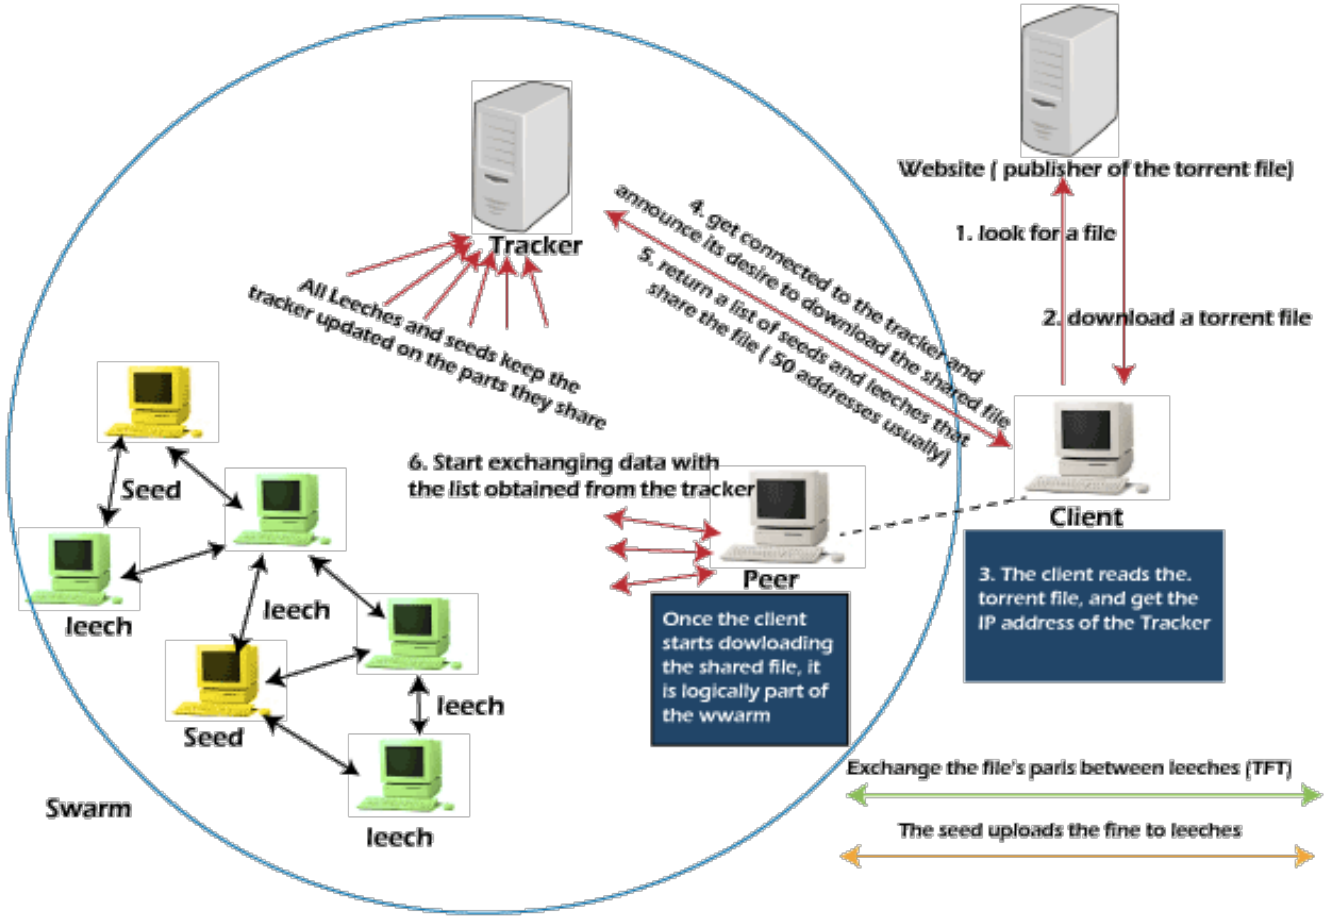

APPLICATIONS OF DISTRIBUTED SYSTEMS

- Blockchain

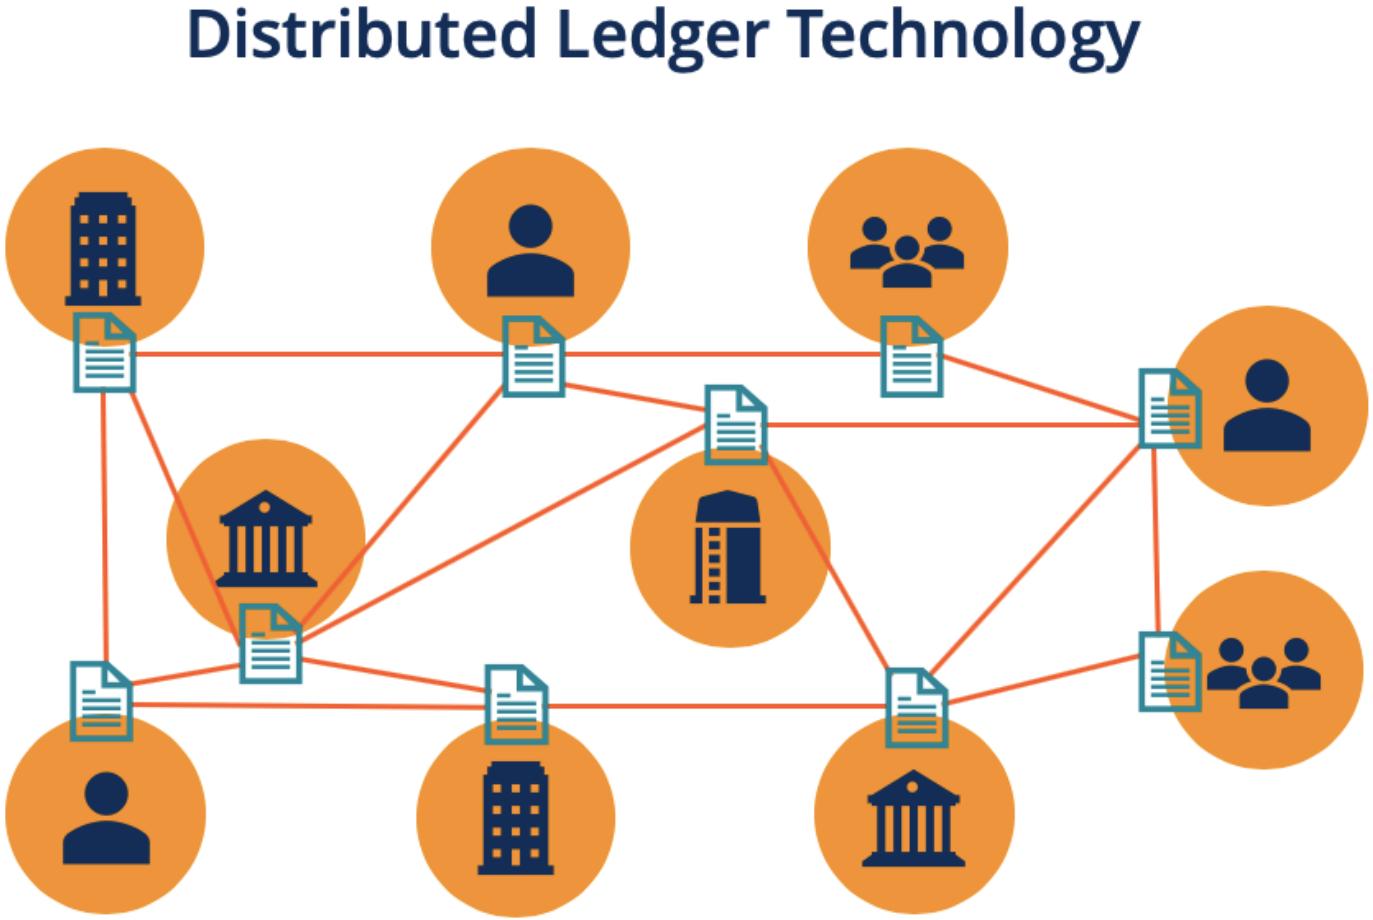

APPLICATIONS OF DISTRIBUTED SYSTEMS

- Smart Homes

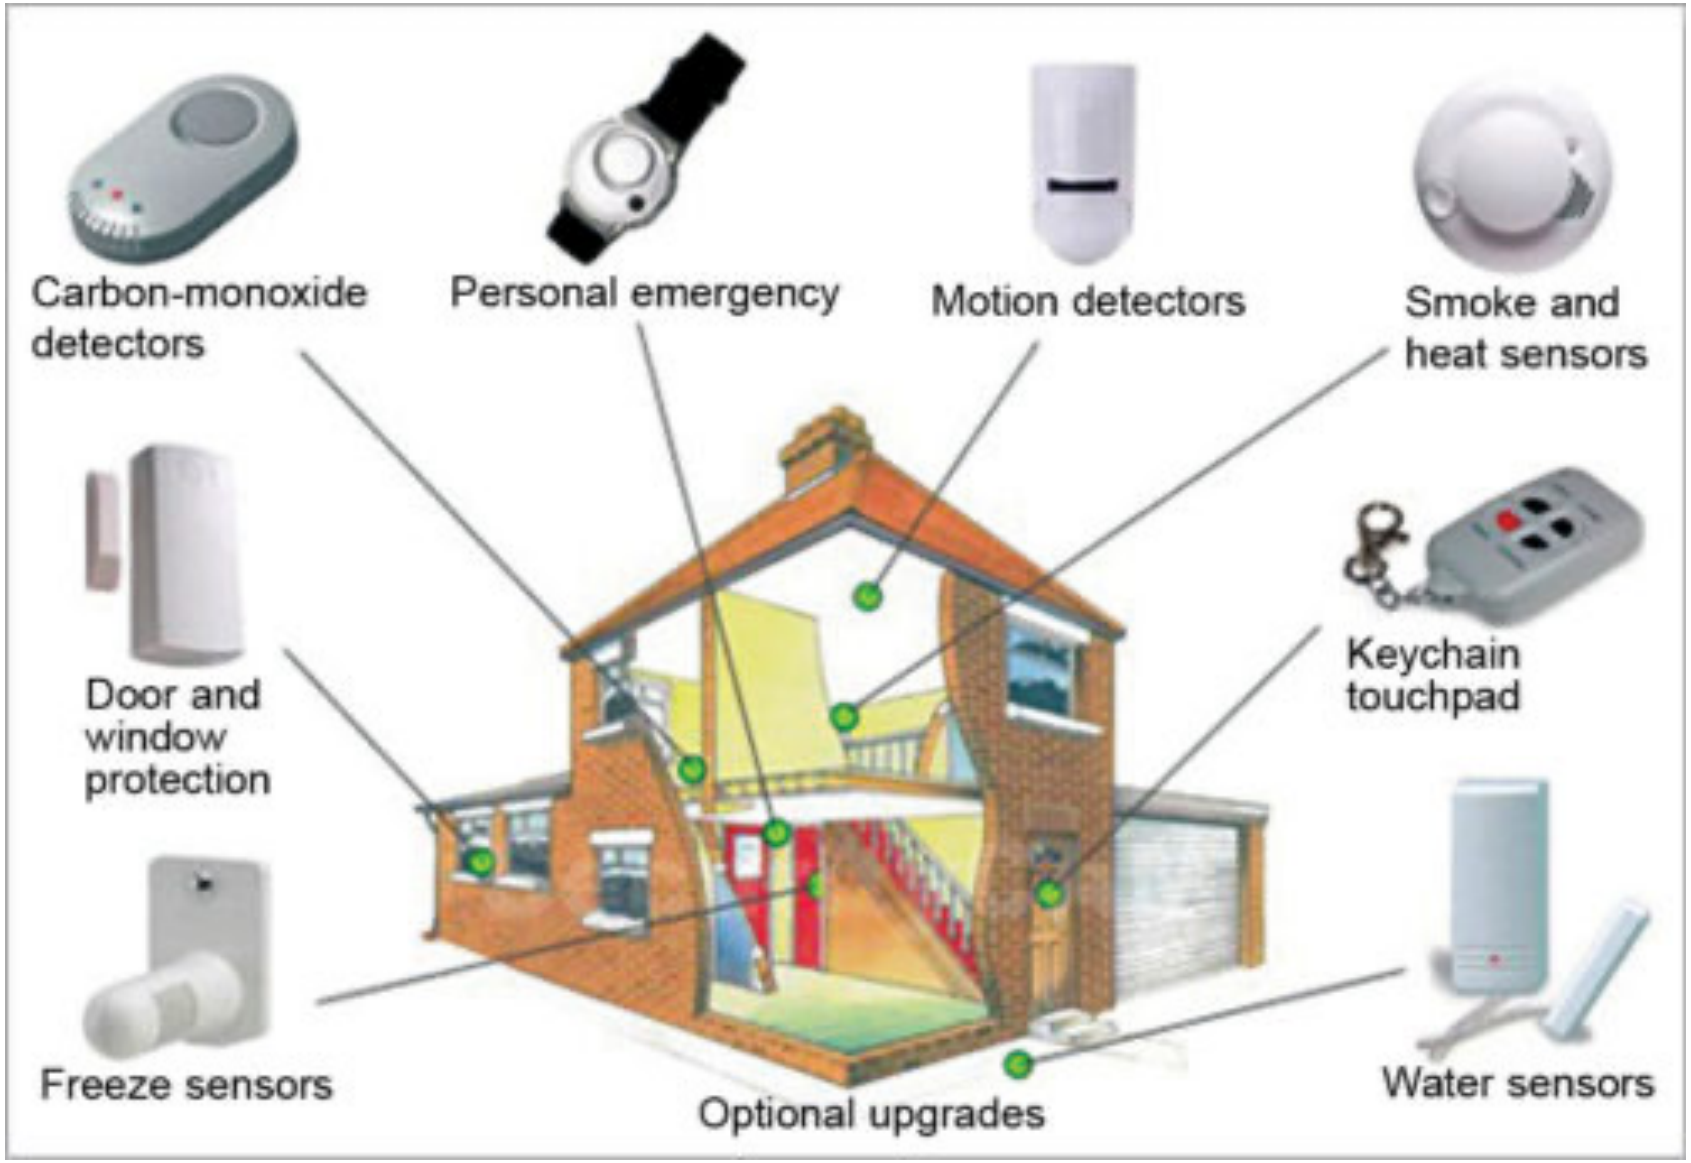

APPLICATIONS OF DISTRIBUTED SYSTEMS

- Healthcare
- Body sensors
- Patient care
- Hospital management

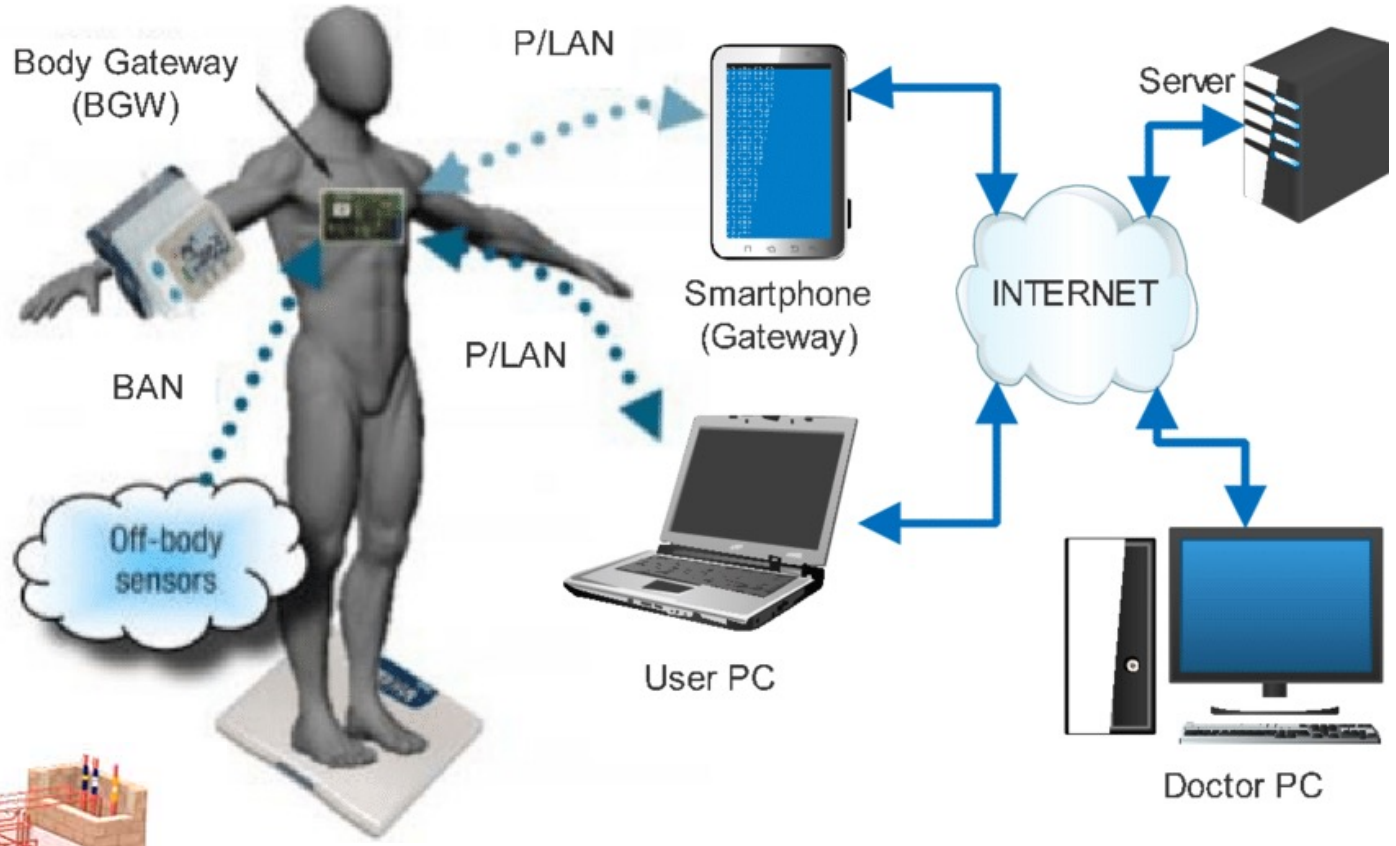

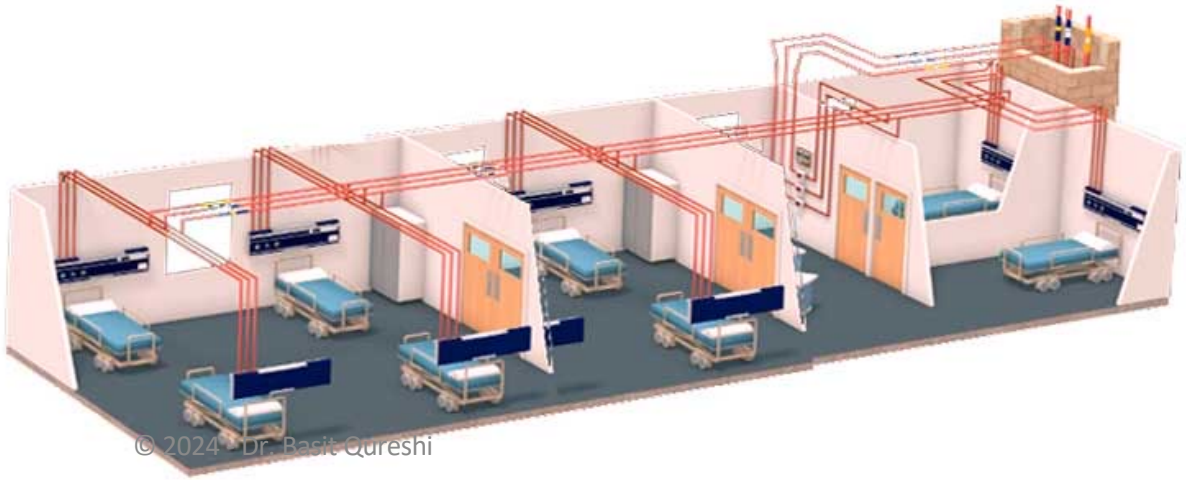

SUMMARY

- Distributed systems are composed of multiple computers connected by a network working together to achieve a goal/task
- Pretty much all systems nowadays are distributed systems
- Goals for Distributed systems:
  - Scalability (Scale up/down size/volume)
  - Consistency (Performance)
  - Reliability (Fault Tolerance/Failure)
  - Availability (No DNS)
  - Complexity (Transparency)
- Various types of distributed systems
- Various applications of distributed systems