Cloud Computing

Chapter 1 : Cloud Computing Fundamentals

By : Prof. Lamri SAYAD

2025

Course Content

- Introduction
  - History
  - Context
- What is a Cloud computing
  - Definition
  - Characteristics
- Basic Concepts
  - Service models
  - Deployment models
- Advantages and disadvantages

Introduction

Cloud computing is everywhere

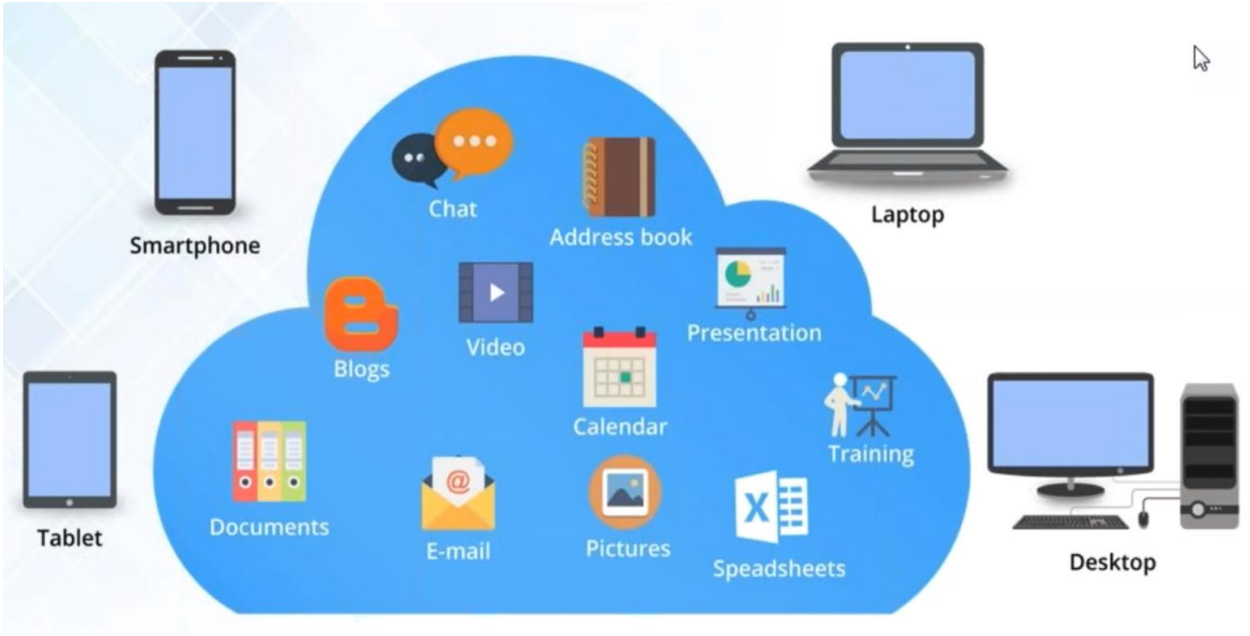

You are already a cloud user

Short history

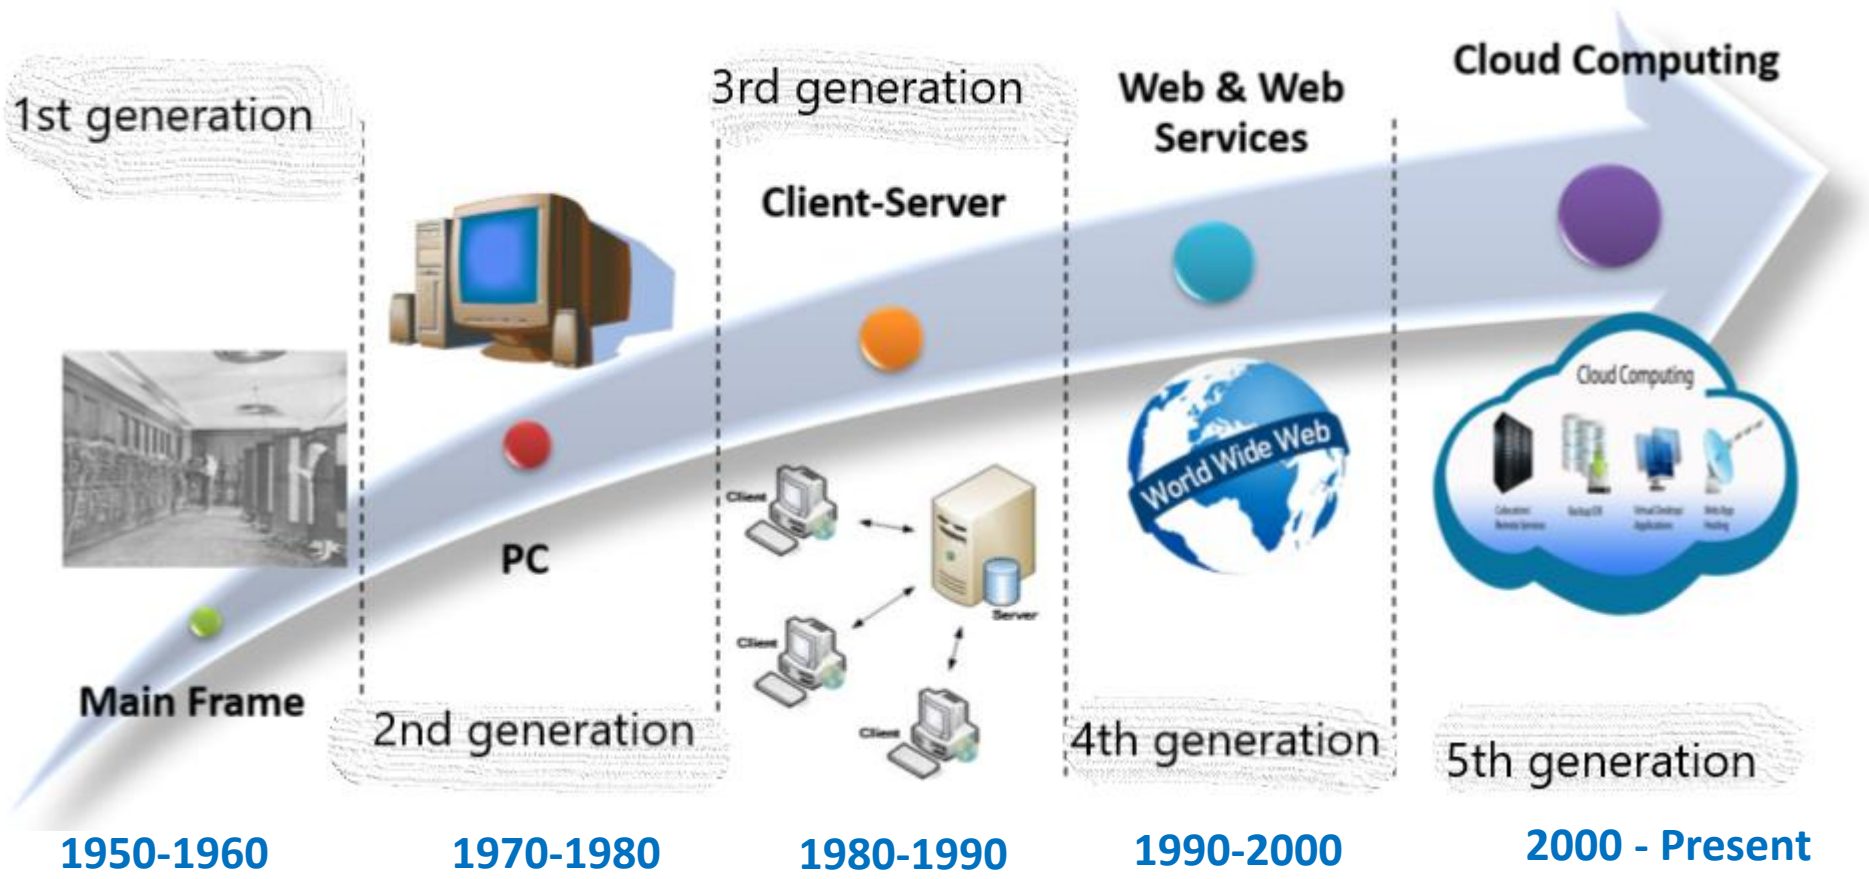

Short history

- **14th November 1996**, Huston (Texas), George Favaloro uses "Cloud" for the first time at the Compaq meeting
  - **Compaq bought a start-up company** (NetCentric) which supports the idea that **data** and **applications** will be in **Internet** and not in the PCs
- 2006 Amazon Web Service (AWS)
- 2008 Google Cloud Platform (GCP)
- 2010 Microsoft Azure
- 2020 Gaia-X: EU Cloud

Informal definitions

- The term “Cloud computing” is misleading. As a marketing **buzzword**, it’s used to suggest that something new and better is going on, when in fact there may be **nothing new about it**.

- **Cloud Computing** is a **business concept**, not a technological jargon dependent on "as a service model"

Informal definitions

Definition by NIST

Cloud computing is a **model** for enabling ubiquitous, convenient, **on-demand network access to a shared pool of configurable computing resource**s (e.g., networks, servers, storage, applications, and services) that can be **rapidly provisioned** and **released** with **minimal management** effort or service provider interaction. This cloud model is composed of:

- **5 essential characteristics**
- **3 service models**
- **4 deployment models**

Cloud at a glance

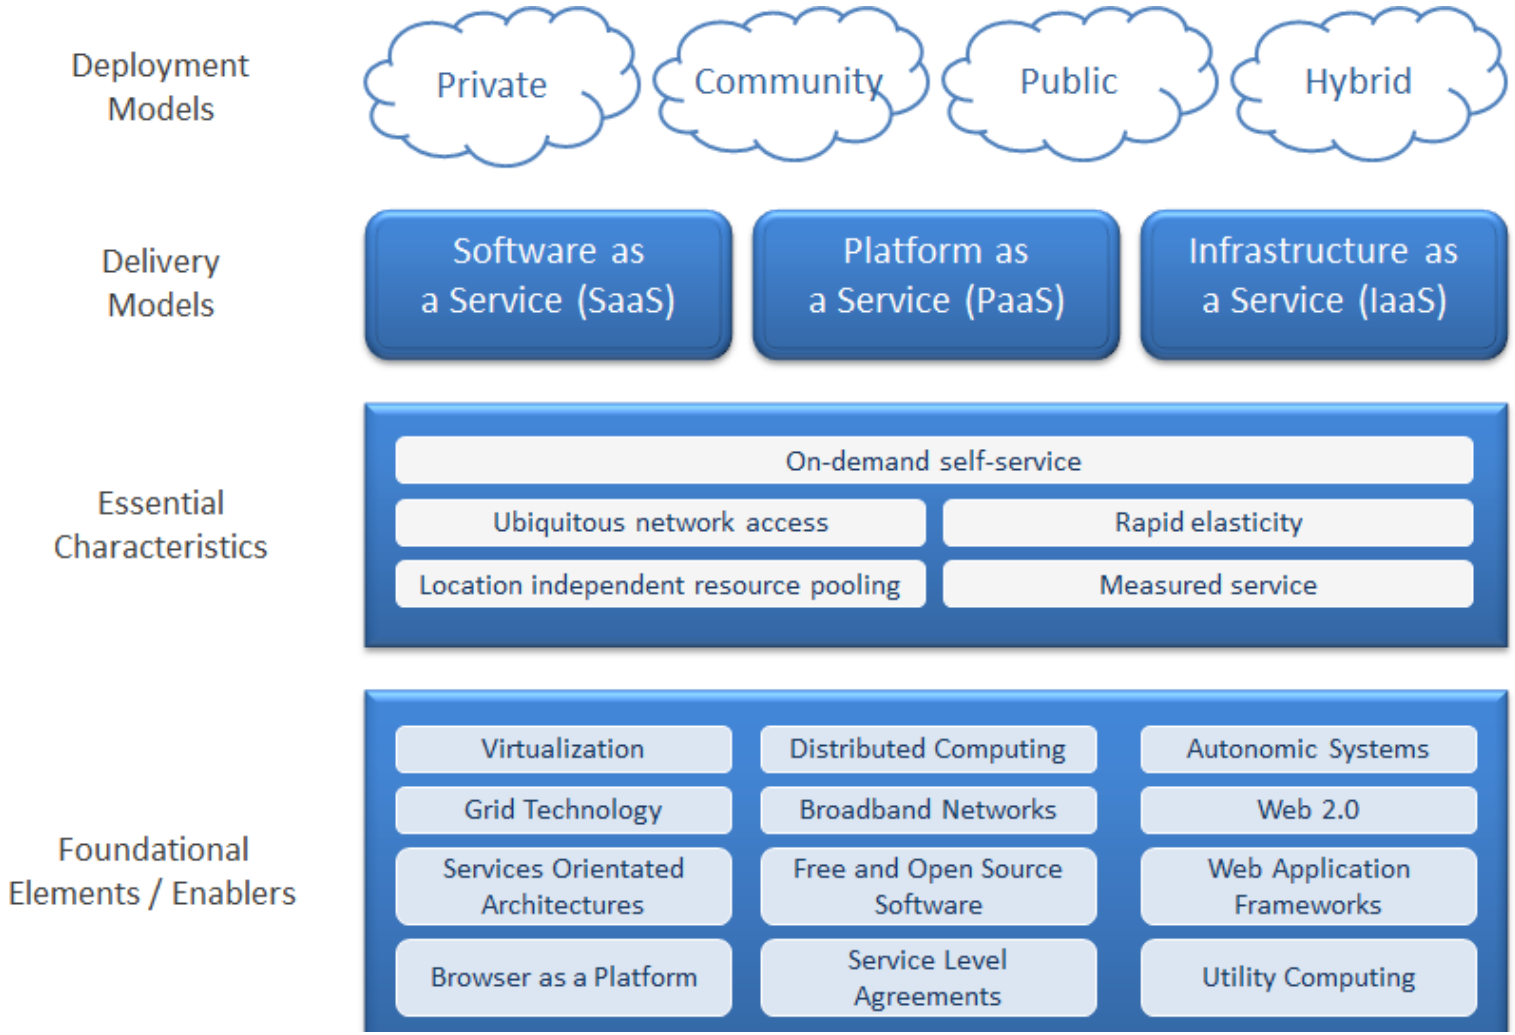

Essential characteristics

- **On-demand self-service:** consumer asks for resources (storage/processing/memory/network bandwidth) without human intervention
- **Broad network access:** resource available over internet
- **Resource pooling:** computing resources are pooled to serve multiple consumers and are dynamically assigned and reassigned
- **Rapid elasticity:** Capabilities can be elastically provisioned and released
- **Measured services:** Control and optimization of the resource usage by leveraging a **metering** capability. Enables pay-as-you-go model.

Service models

Service models

Service Models are the reference models on which the Cloud Computing is based.

These can be categorized into three basic service models as listed below:

1. Infrastructure as a Service (IaaS)
2. Platform as a Service (PaaS)
3. Software as a Service (SaaS)

Service Models

- **Software as a Service (SaaS):** The customer just asks for a service
- **Platform as a Service (PaaS):** The customer can deploy applications by using libraries/services/tools provided by the cloud infrastructure
- **Infrastructure as a service (IaaS):** The customer has access to processing, storage, networks and other fundamental computing resources.

Cloud service models

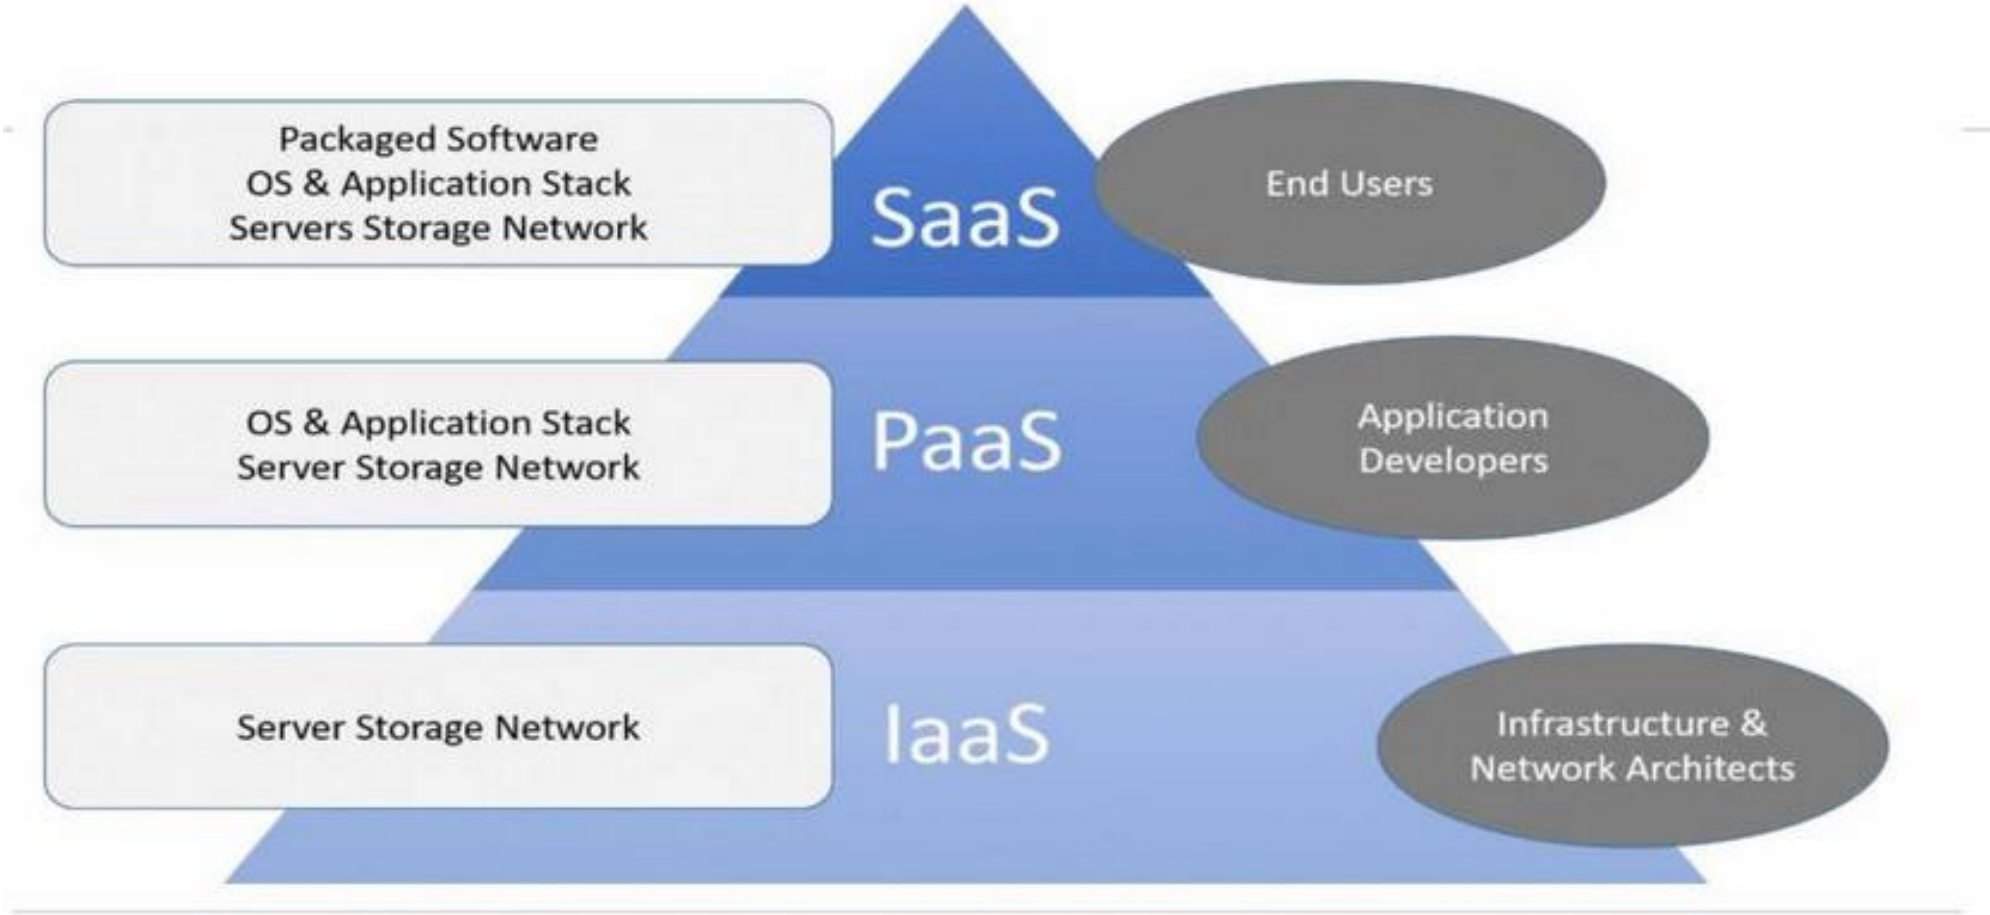

Infrastructure as a Service (IaaS) characteristics

- IaaS provides **basic storage** and **computing capabilities** as standardized services over the network
- Servers, storage systems, networking equipment, data centre space etc are pooled and made available to handle workloads
- The customer would typically deploy his **own software** on the infrastructure.
- IaaS examples: Amazon EC2, OpenStack, ...

Best IaaS providers of 2019

Google
Cloud Platform

IBM **Cloud**

IaaS providers comparison

	Windows Azure	amazon web services EC2	rackspace	DigitalOcean	linode.com	SOFTLAYER	Google Compute Engine
Low-end	1core/.75gb /20gb	1core/1gb/EBS	1core/1gb/20gb	1core/512mb /20gb	1core/1gb/24gb	1core/1gb/25gb	1 core/.6gb
High-end	16core/112gb/ 382gb	16core/117gb/24x2048	32core/120gb/1.2tb	20core/64gb/ 640gb	20core/96gb /1.9TB	custom	16core/104gb
Pricing	$.018 - $4.9/hr	$.013 - $4.6/hr	$.052 - $6.24/hr*	$.007 - $.952/hr	$.015 - $1.44/hr	$.04 - $x.xx/hr	$.012 - $1.184/hr
Bandwidth	5GB incl., then ~$.08/GB	15GB incl.	~$.10/GB	1-9TB Incl.	2-20TB Incl.	5TB incl.	~$.01/GB
Network Out	up to 10GB	**	200MB-10GB	1GB	125MB-10GB	up to 10GB	**
Bare Metal	No	No	Yes	No	No	Yes	No
SLA	99.95%	99.95%	99.9%	99.99%	99.9%	1	99.95%

https://www.scriptrock.com/articles/cloud-service-provider-roundup-the-best-of-the-best

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

Platform as a Service (PaaS) characteristics

- A **layer of software**, or development environment is encapsulated and offered as a service
- PaaS provides a runtime environment for applications, development and deployment tools, etc.
- PaaS provides all the facilities required to support the life cycle for building and delivering applications and services from internet.
- The customer has the freedom to build his **own applications**, which run on the provider's infrastructure
- **PaaS providers** offer a predefined **combination of OS and application servers**, such as **LAMP** platform, **authentication** layer, **storage** layer, etc.
- PaaS example: **Google's App Engine**, Force.com

Top PaaS providers 2019

AWS Lambda

Morpheus

The main PaaS providers in the market

Service	Editeur	Langages	IDE	Bases de données
Amazon Elastic Beanstalk	Amazon Web Services	Containers Docker, Java (Tomcat), .NET, Node.js, PHP, Python, Ruby (Passenger)	Extensions AWS Toolkit pour Eclipse et Microsoft Visual Studio	Amazon RDS, DynamoDB, SimpleDB, SQL Server, Oracle, IDBM DB2, Informix
Google App Engine	Google	Java, Python, PHP 5.4, Go	Google Plugin for Eclipse pour Java	Cloud SQL
Azure Cloud Services	Microsoft	Java, Node.js, PHP, Python, .NET, Ruby	WebMatrix, Visual Studio + Azure SDK	Data Services Microsoft : Blobs, base SQL, SQL Server, MongoDB
Salesforce1 platform	Salesforce	Langage Apex sur Force.com, Java, Ruby, Node.js, Python sur Heroku1	Outil web pour Force.com, client Heroku toolbelt pour Windows, MacOS X, Linux Debian/Ubuntu	Force.com Database, Heroku Postgres + les les SGBD de l'add-on marketplace Heroku, dont MongoDB, Postgres, MySQLÂÂ
Cloud Foundry	Cloud Foundry Community	Java, Grails, Play, Spring, Node.js, Ruby on Rails, Sinatra, Go, Erlang, etc.	Plugin Eclipse	MySQL, PostgreSQL, MongoDB

Source : JDN

PaaS - Example

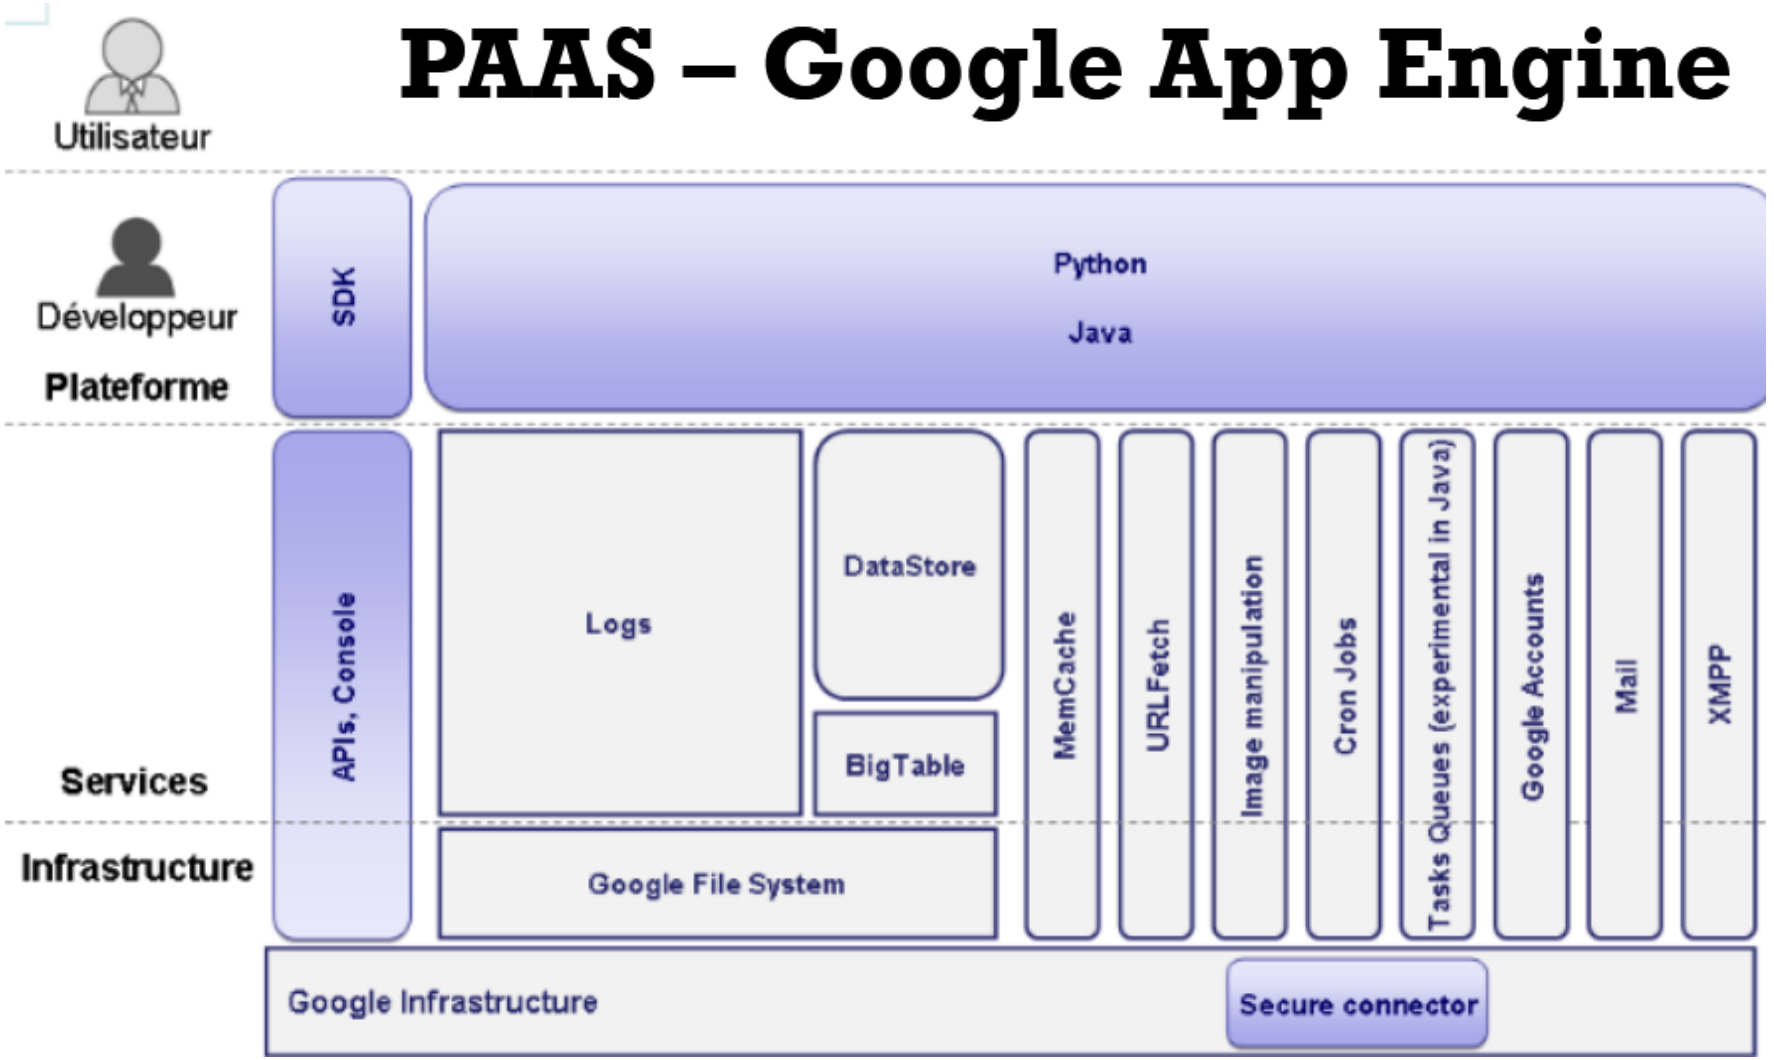

PAAS – Google App Engine

	FREE *LIMIT PER DAY*	PRICE *ABOVE FREE LIMIT*
Instances	28 instance hours	$0.05 / instance / hour
Cloud Datastore *(NoSQL)*	• 50k read/write/small • 1 GB storage	• $0.06 / 100k read or write ops • Small operations free* • $0.18 / GB / month
Network Traffic *(Outgoing)*	1 GB	$0.12 / GB
Network Traffic *(Incoming)*	1 GB	FREE
Cloud Storage	5 GB	$0.026 / GB / month
Memcache	• Free Usage of Shared Pool • No free quota for Dedicated Pool	• Free Usage of Shared Pool • Dedicated Pool: $0.06 / GB / hour
Search	• 1000 basic operations • 0.01 GB indexing documents • 0.25 GB document storage • 100 searches	• 0.50 / 10k searches • 2.00 / GB indexing documents • 0.18 / GB / month Storage
Email API	100 recipients	Contact Sales
Logs API	100 MB	$0.12 per GB
Task Queue	5 GB	$0.026 / GB / month
Logs Storage	1 GB	$0.026 / GB / month
SSL Virtual IPs	-	$39 / virtual IP / month
Bundled Services	Cron, Image Manipulation, Pagespeed, SNI SSL Certificates, Socket API, Task Queue API, URLFetch, Users API	

Software as a Service (SaaS)

- A **complete application** is offered to the customer, **as a service** on demand
- A single instance of the service runs on the cloud and multiple end users are serviced
- **Customers' side**: no need for **upfront investment** in servers or software licenses
- **Providers' side**: the costs are lowered, since only a **single application** needs to be **hosted** and **maintained**
- SaaS examples: any **webmail**, any **social network**, Dropbox, GDocs, Salesforce, **Netflix**, ...

Top SaaS services

Microsoft Office 365

Google Apps

salesforce

Examples

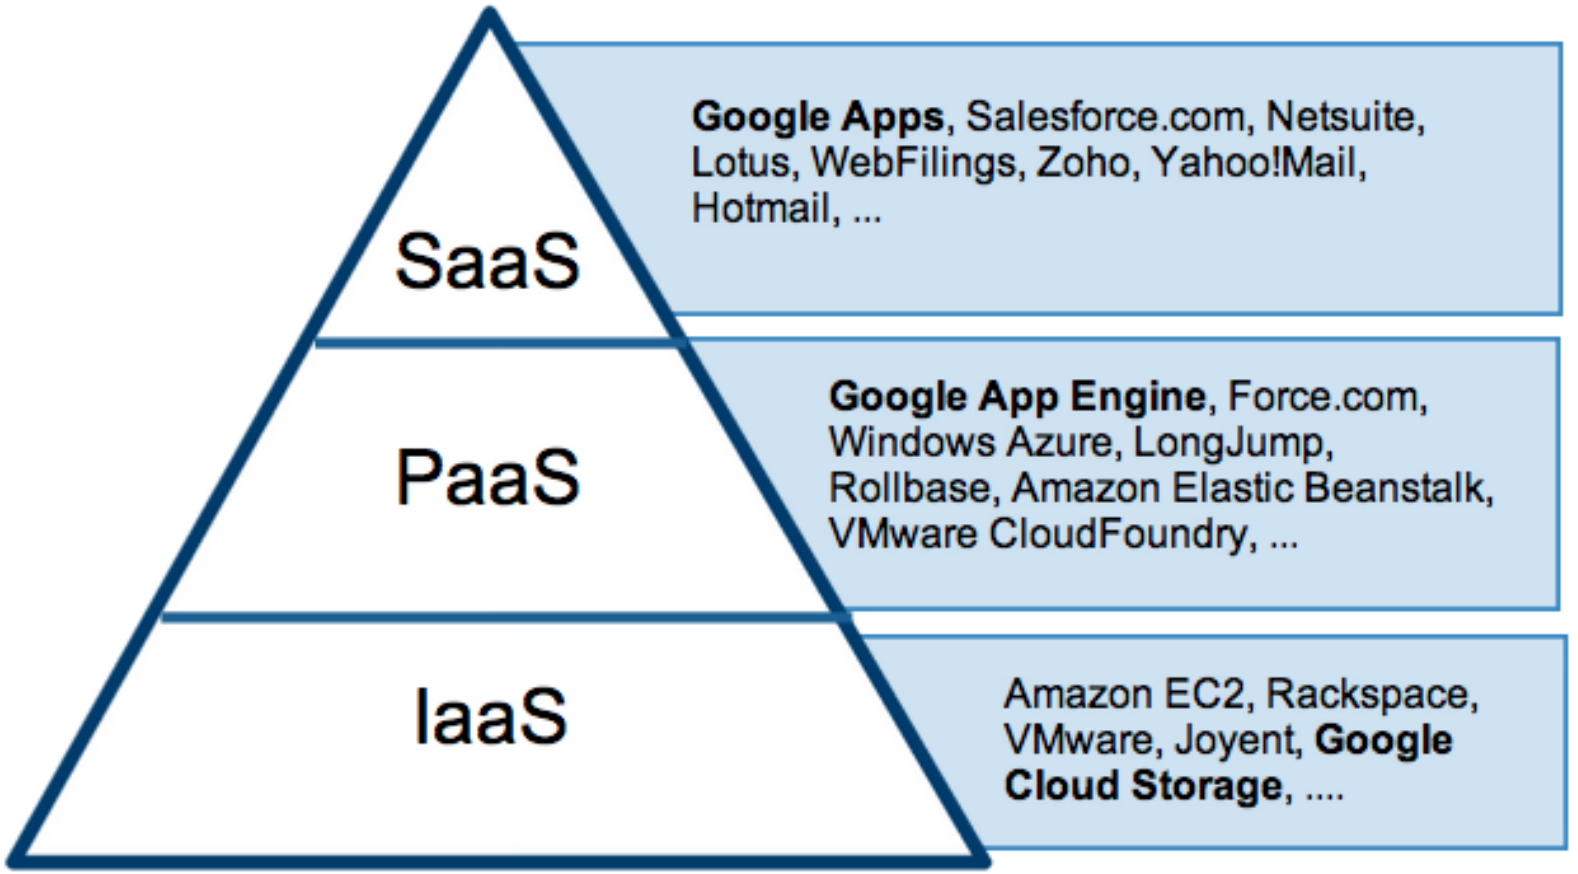
SaaS
PaaS
IaaS
Google Apps, Salesforce.com, Netsuite, Lotus, WebFilings, Zoho, Yahoo!Mail, Hotmail, ...
Google App Engine, Force.com, Windows Azure, LongJump, Rollbase, Amazon Elastic Beanstalk, VMware CloudFoundry, ...
Amazon EC2, Rackspace, VMware, Joyent, **Google Cloud Storage**, ...

Pizza as a Service

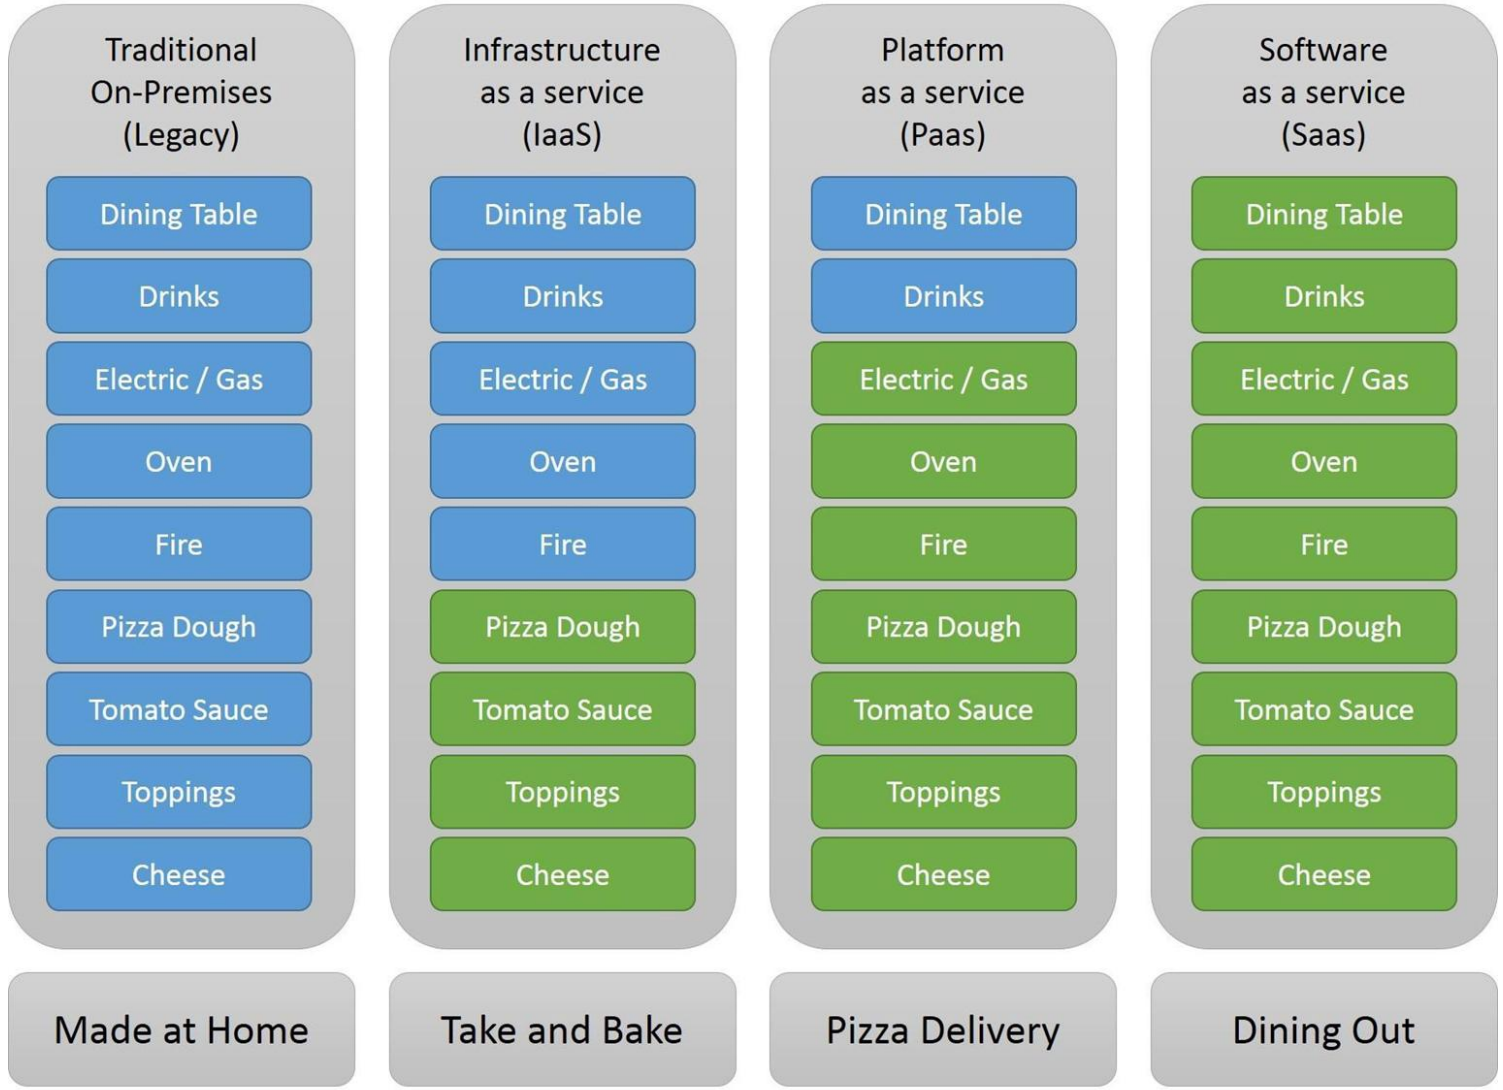

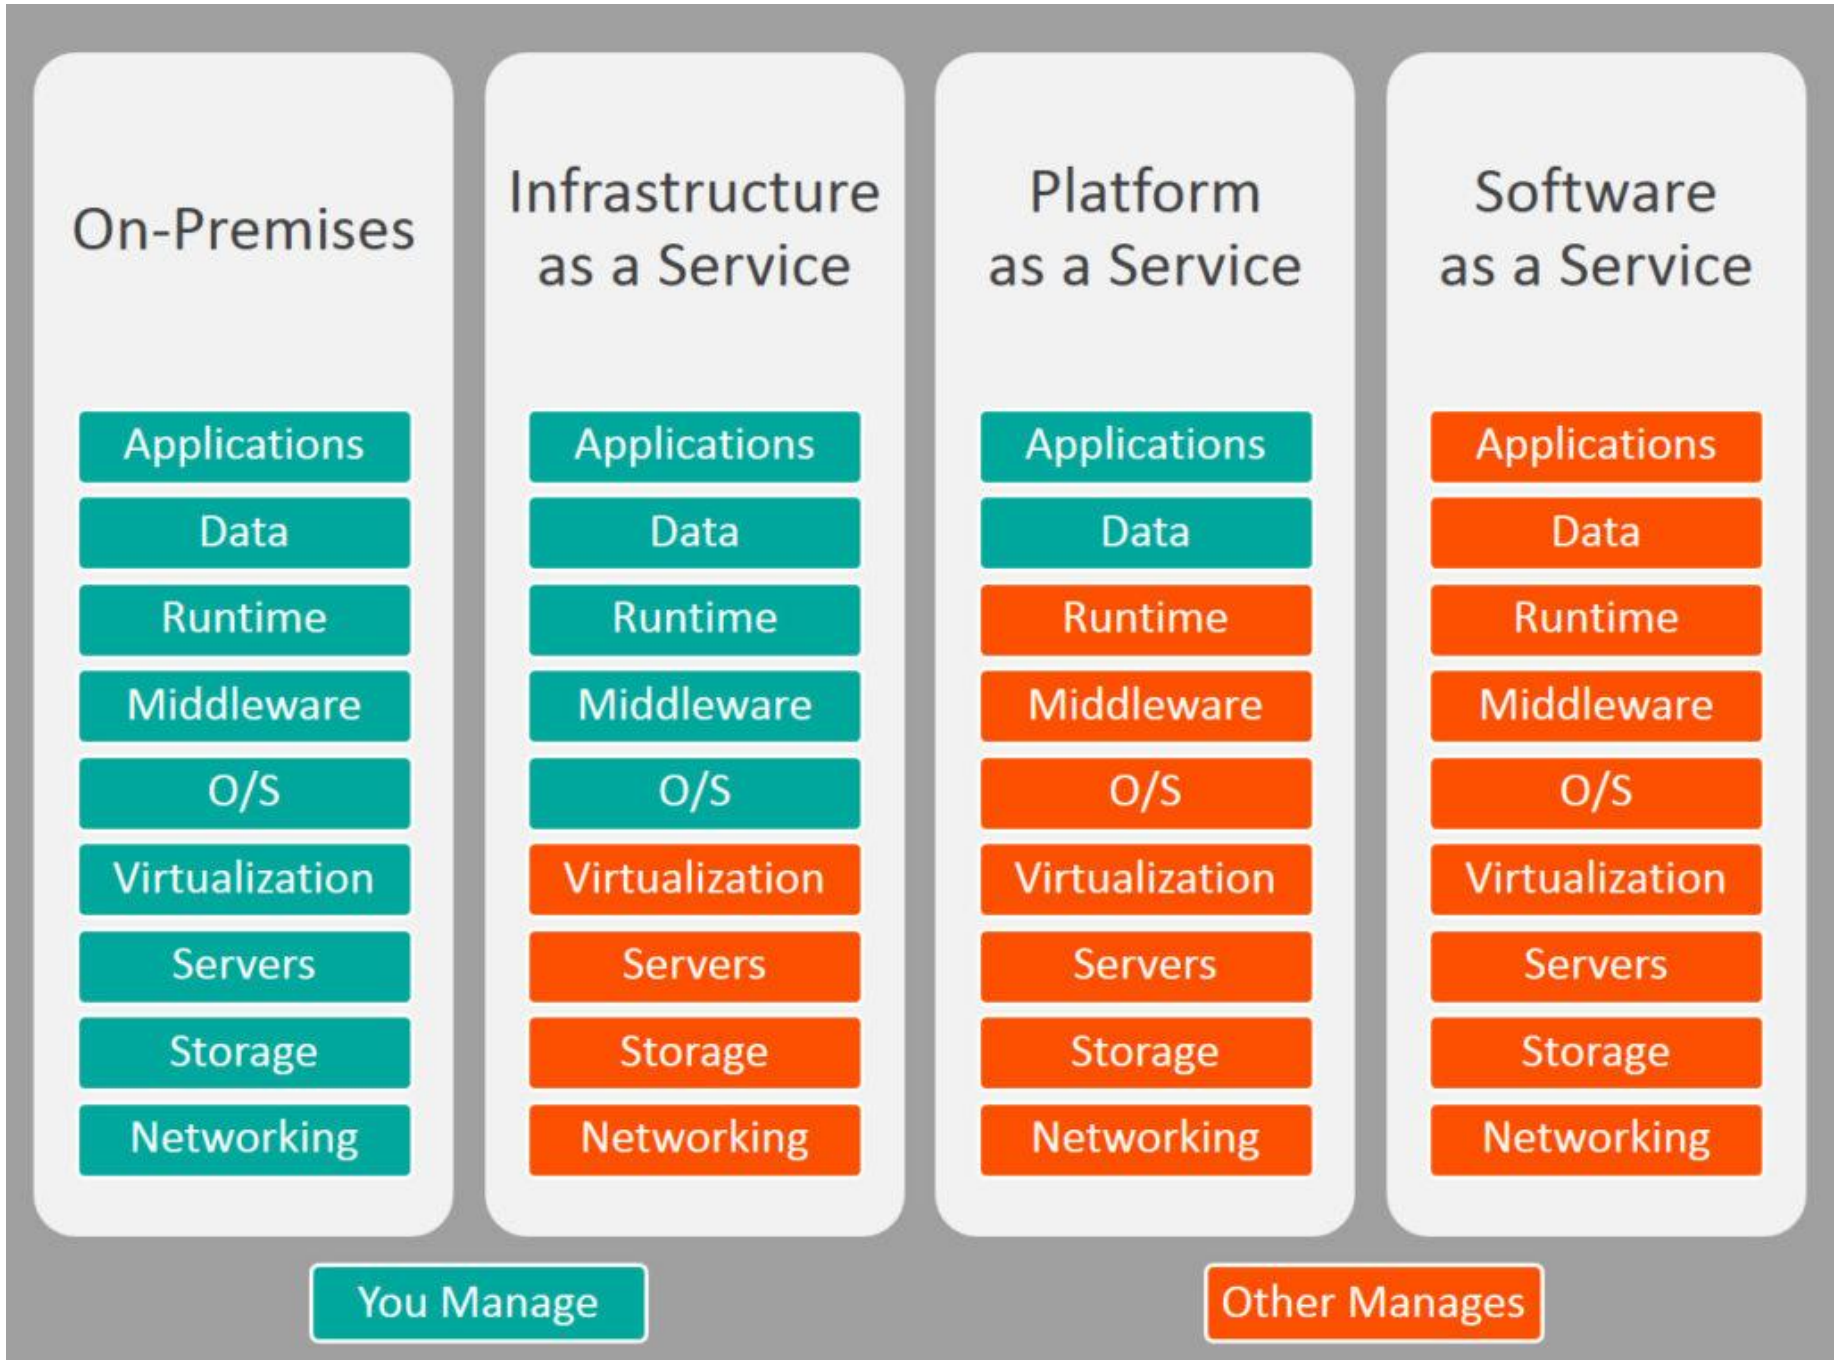
On-Premises
Applications
Data
Runtime
Middleware
O/S
Virtualization
Servers
Storage
Networking
Infrastructure as a Service
Applications
Data
Runtime
Middleware
O/S
Virtualization
Servers
Storage
Networking
Platform as a Service
Applications
Data
Runtime
Middleware
O/S
Virtualization
Servers
Storage
Networking
Software as a Service
Applications
Data
Runtime
Middleware
O/S
Virtualization
Servers
Storage
Networking
You Manage
Other Manages

Deployment models

Deployment models

Deployment models define the type of access to the cloud, i.e., how the cloud is located?

Cloud can have any of the four types of access: Public, Private, Hybrid and Community.

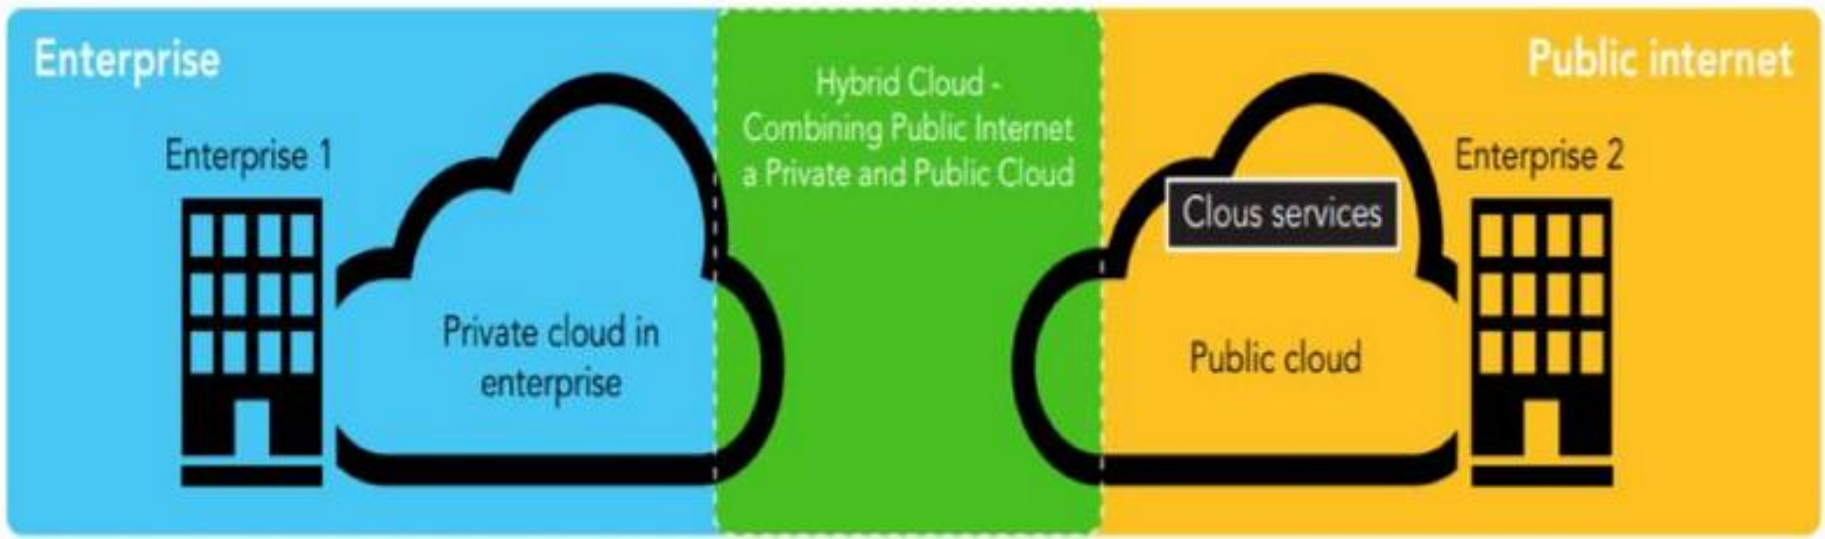

Deployment models

- **Private cloud:** The cloud computing is provisioned for exclusive use by a single organization compromising multiple consumers
- **Community cloud:** The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns
- **Public cloud:** The cloud infrastructure is provisioned for open use by the general public.
- **Hybrid cloud:** The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public)

Public cloud

- **Owned** and **operated** by **third parties**
- Give to each individual client an attractive **low-cost**, "**Pay-as-you-go**" model
- **All customers share the same infrastructure pool** with
  - limited configuration
  - security protections
  - availability variances
- May be **larger than an enterprises** cloud
- **Scale seamlessly on demand**

Name	VCPUS	RAM	Total Disk	Root Disk
m1.tiny	1	512 MB	1 GB	1 GB
m1.small	1	2 GB	20 GB	20 GB
m1.medium	2	4 GB	40 GB	40 GB
m1.large	4	8 GB	40 GB	40 GB
m1.xlarge	8	16 GB	40 GB	40 GB
m1.xxlarge	16	32 GB	40 GB	40 GB

Public cloud

Private cloud

- Are **built exclusively** for a **single enterprise**
- Aim to address concerns on **data security** and offer **greater control**
- There are two options:
  - **On-premise** Private Cloud
  - **Externally hosted** Private Cloud
    - Hosted externally with a cloud provider
    - The provider facilitates an **exclusive cloud environment** with **full guarantee of privacy**
    - **Best suited for enterprises** that **don't prefer a public cloud** due to sharing of physical resources

Hybrid cloud

- **Combine** both **public** and **private** cloud models
- Service providers **can utilize 3rd party Cloud Providers** in a **full** or **partial manner** thus increasing the flexibility of computing
- The ability to augment a private cloud with the resources of a public cloud can be used to **manage any unexpected surges in workload**

Hybrid Cloud

A cloud computing service that is composed of some combination of private, public and community cloud services from different service providers for capacity or capability.

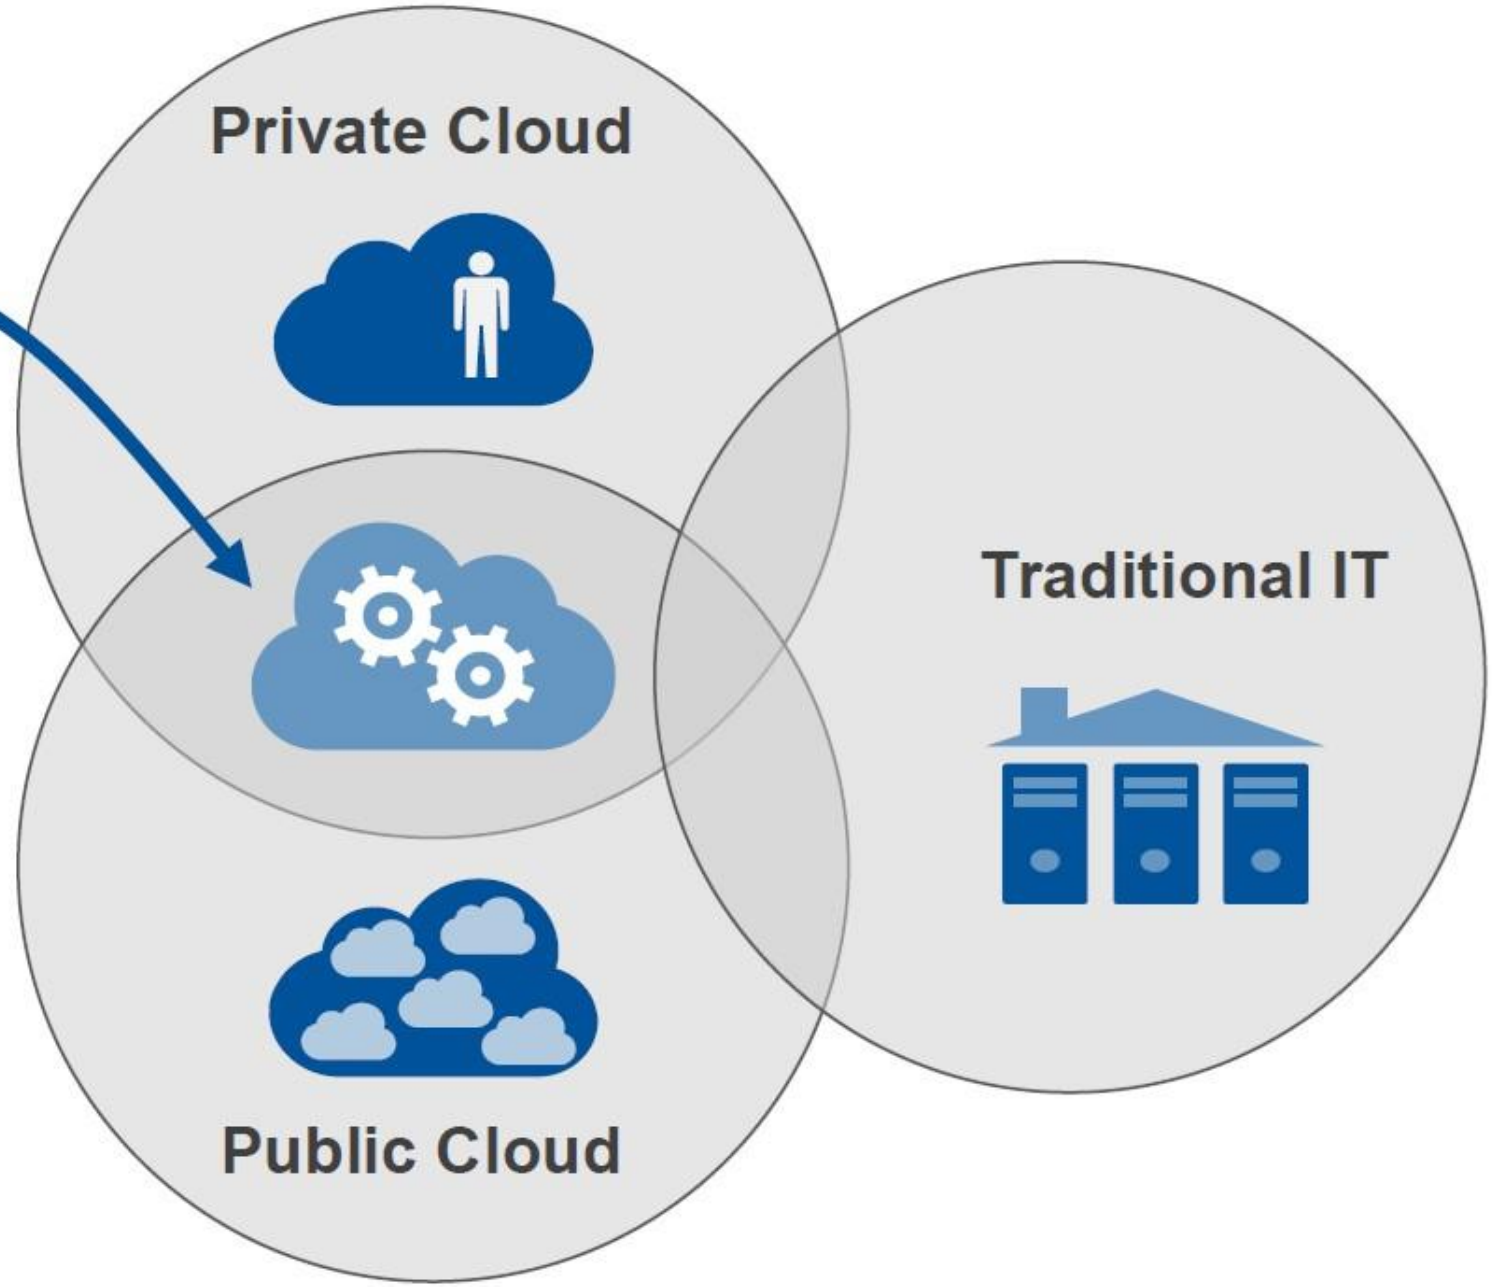

Where we are

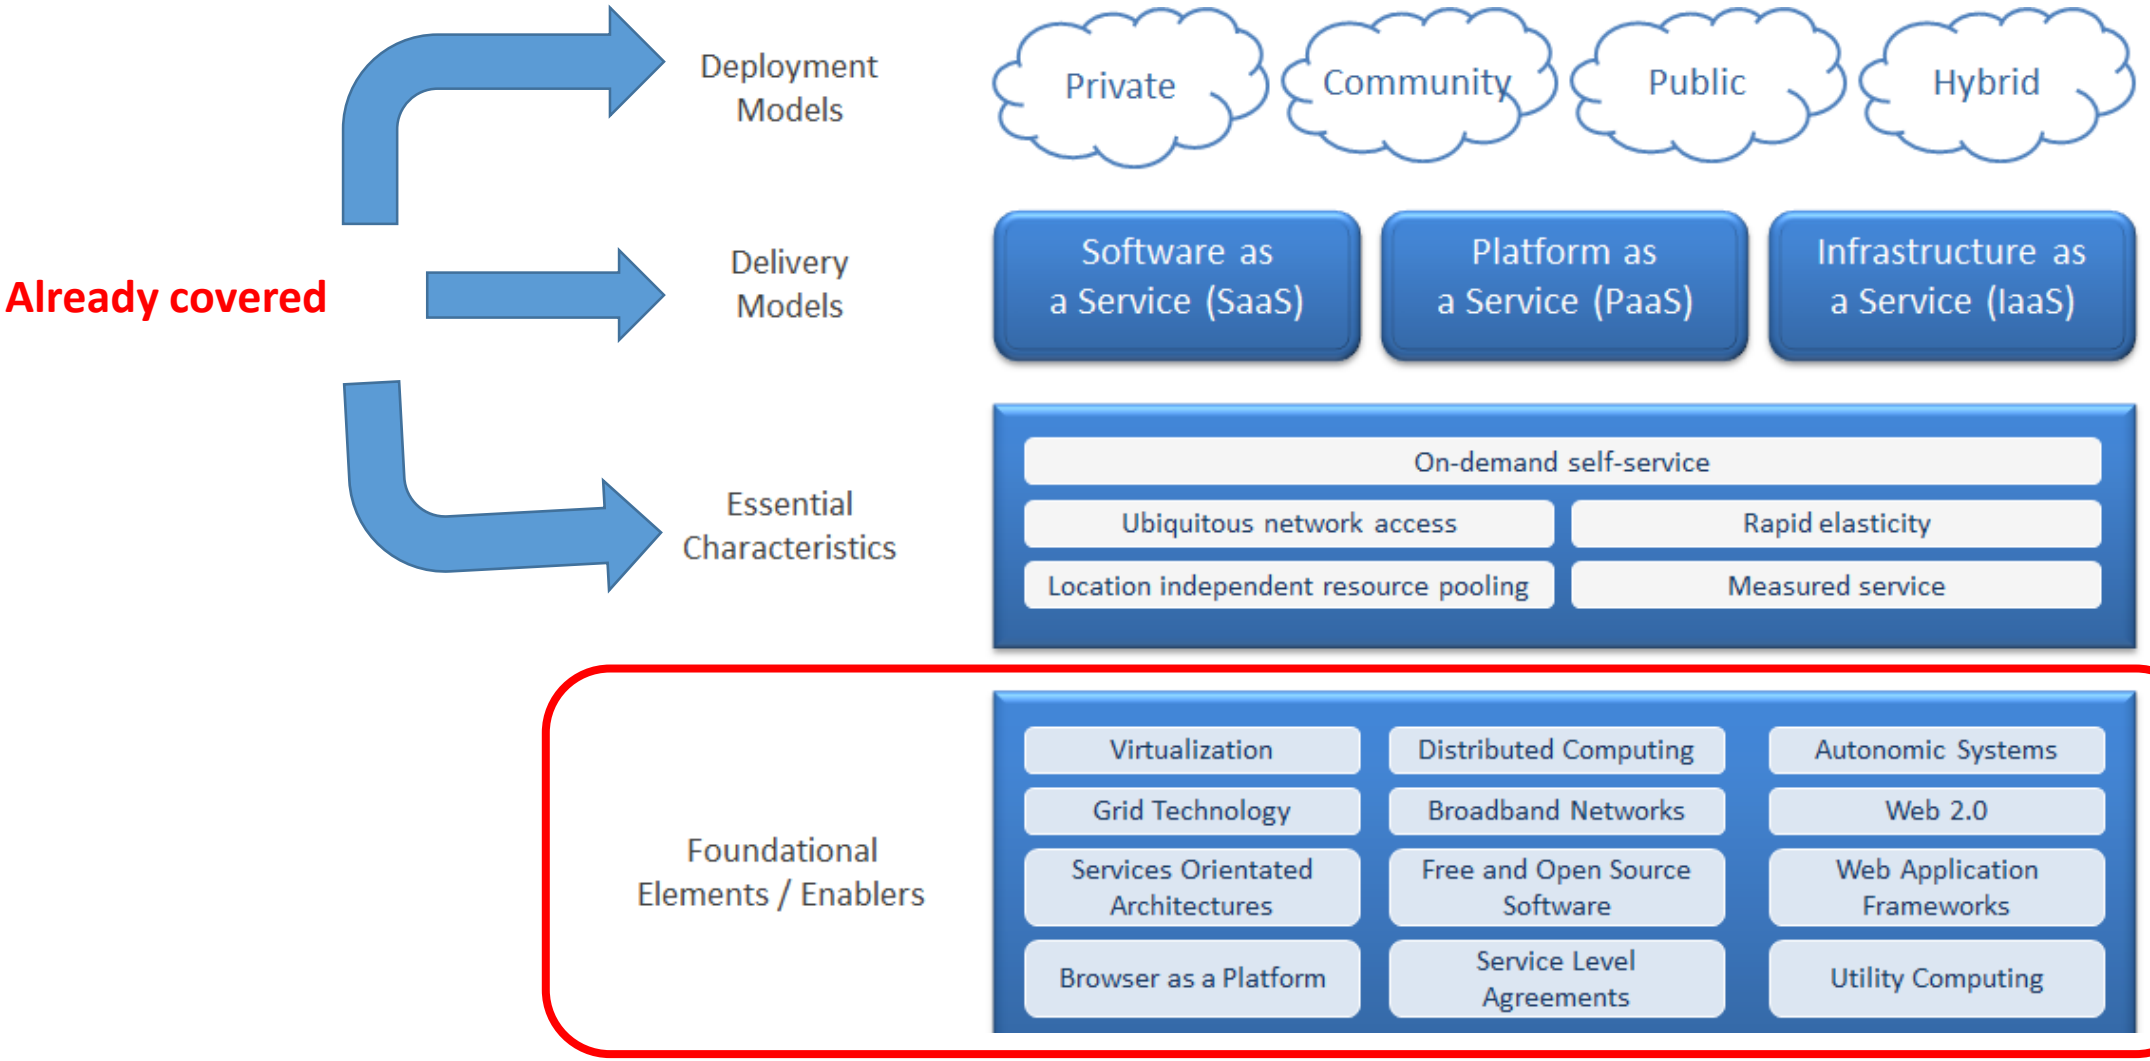
Already covered
Deployment Models
Private
Community
Public
Hybrid
Delivery Models
Software as a Service (SaaS)
Platform as a Service (PaaS)
Infrastructure as a Service (IaaS)
Essential Characteristics
On-demand self-service
Ubiquitous network access
Rapid elasticity
Location independent resource pooling
Measured service
Foundational Elements / Enablers
Virtualization
Distributed Computing
Autonomic Systems
Grid Technology
Broadband Networks
Web 2.0
Services Orientated Architectures
Free and Open Source Software
Web Application Frameworks
Browser as a Platform
Service Level Agreements
Utility Computing

Cloud Ecosystem

Cloud Ecosystem

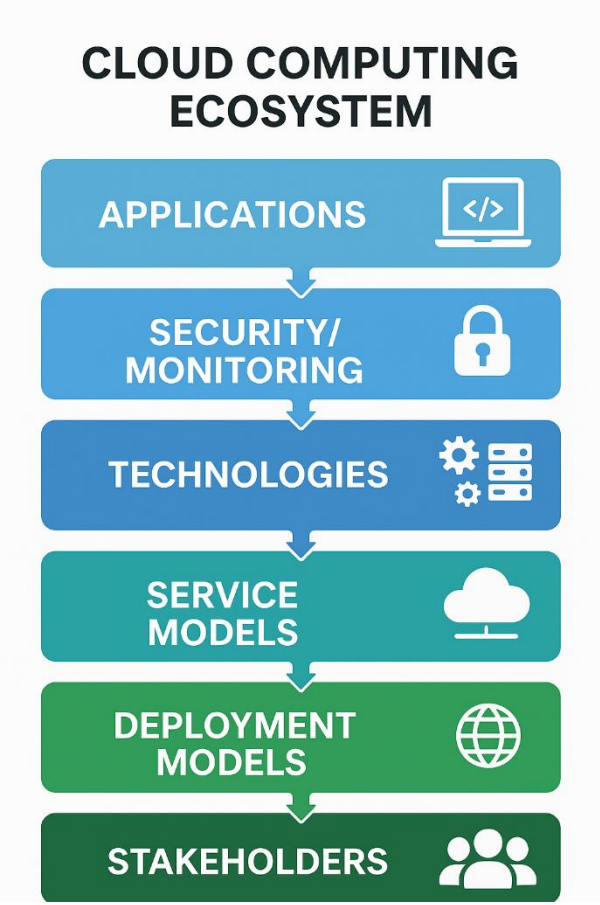

AI/ML, IoT, Big Data …

Virtualization, containers …

SaaS, PaaS, IaaS

Public, Private, Hybrid

Providers, Clients …

Cloud Ecosystem :
Stakeholders

- **Ecosystem : complete environment of actors, layers, and technologies** that work together to deliver cloud-based services.

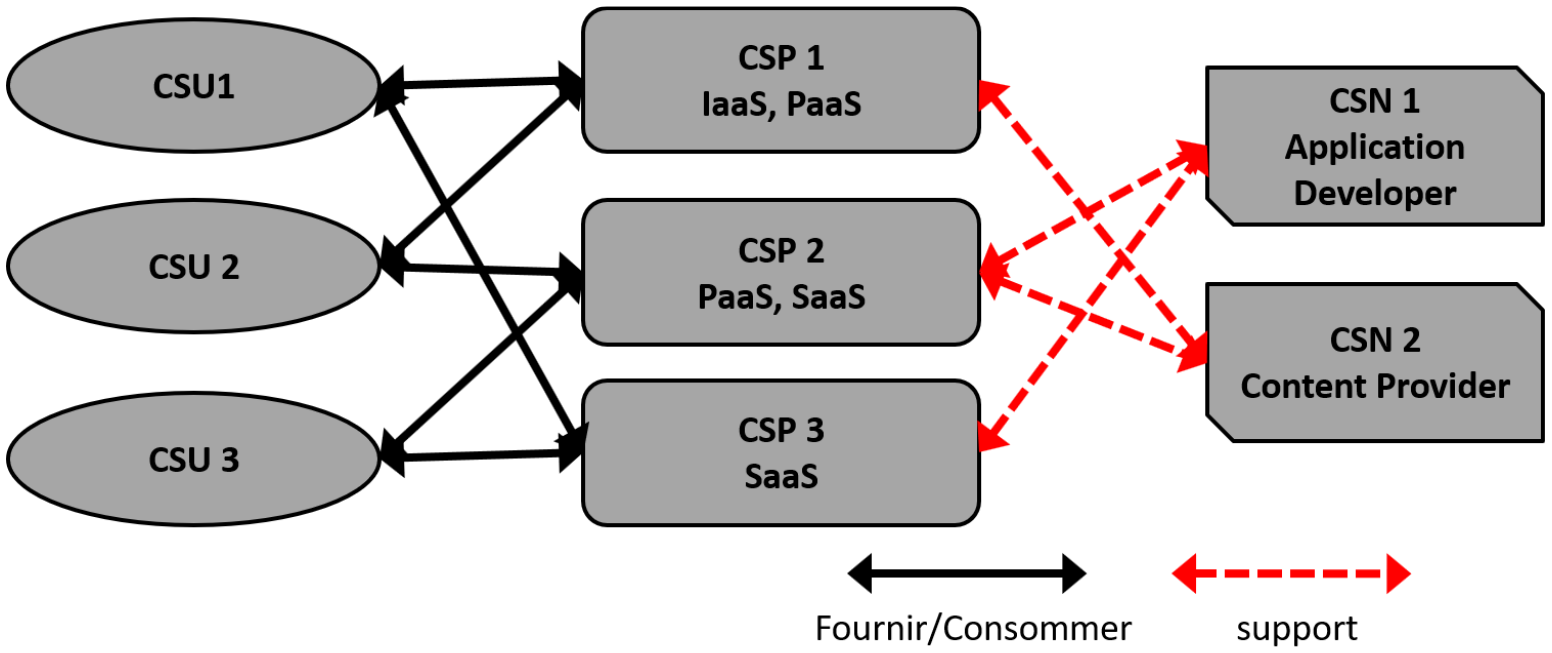

Cloud Ecosystem :
Stakeholders

- **Cloud Service Users (CSU)** : Consume Cloud services like individuals, startups, companies, governments.
  - They may consume SaaS apps (Gmail, Office 365), deploy apps on PaaS, or rent infrastures on IaaS

- **Cloud Service Providers (CSP)** : provide and maintain cloud services at different levels (IaaS, PaaS, SaaS).
  - Commercial providers : AWS, Microsoft Azure, Google Cloud, IBM Cloud, Oracle Cloud.
  - OpenSource providers : Eucalyptus, OpenNebula , Nimbus, Xen Cloud Platform

- **Cloud Brokers :**
  - Middle entities that mediate between cloud providers and consumers.
  - They aggregate services from multiple providers, compare offerings, or add value

- **Cloud Service Partners (CSN**) : persons or organizations that support CSPs (Ex : application developer).

Cloud Ecosystem :
Enabling technologies

- Virtualization (VMware, Hyper-V, KVM).
- Containerization & Orchestration (Docker, Kubernetes).
- Networking (SDN, CDN, VPN).
- Storage (object storage like Amazon S3, block storage, distributed file systems).
- APIs & Middleware (REST, GraphQL, gRPC).
- Automation & DevOps (CI/CD pipelines, Infrastructure as Code, Ansible, Terraform).

Cloud Ecosystem :
Applications & Use Cases

- Enterprise apps (ERP, CRM).
- Big data & analytics (Hadoop, Spark on the cloud).
- AI & Machine Learning (AWS SageMaker, Google Vertex AI).
- IoT platforms (AWS IoT Core, Azure IoT Hub).
- Edge computing & 5G integration.

Cloud Architecture

Cloud Architecture

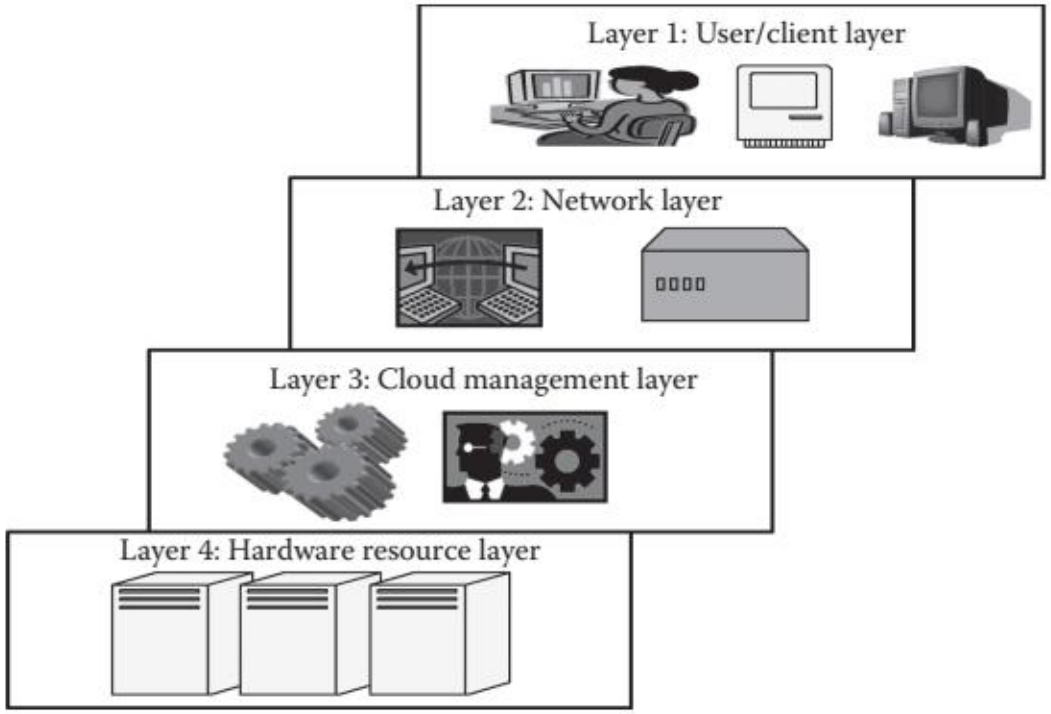

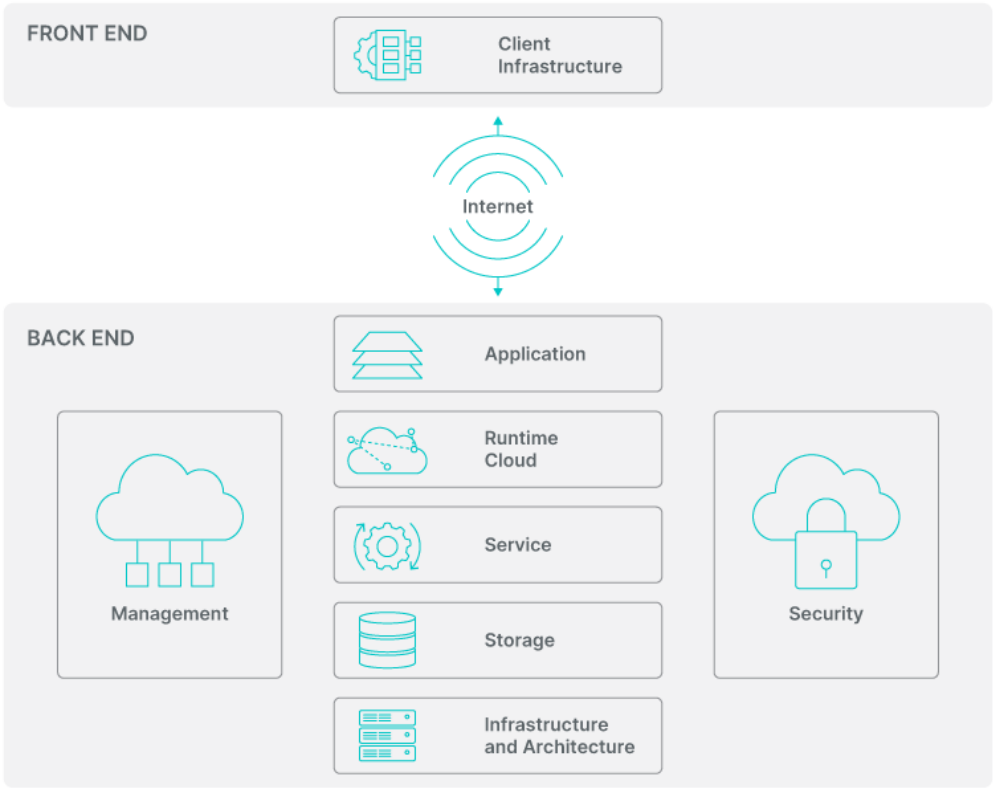

Layer 1 : User/Client Layer

- The interface between the cloud and the user.
  - Devices: PC, laptop, thin client, smartphone.
  - Applications: Browsers, SaaS apps, client-side apps.

- 2 types of clients:
  - Thin client : lightweight device (computer) that relies on a central server for most of its processing power, storage, and applications (ex : web browser, terminal).
  - Thick client (also called fat client) : is a computer that can operate independently of a server because it has its own **processing power, storage, and applications** installed locally.

Layer 2 : Network Layer

- Provide connectivity between users and the Cloud.
- No Network ⟶ No Cloud
- Includes the Internet, intranet, VPNs, and communication protocols (HTTP, HTTPS, REST APIs).
  - Public clouds use internet as a network layer
  - Private clouds use Local Area Networks (LANs) as Network layer
- Ensures **data transfer, security (TLS/SSL, firewalls), QoS, and latency control**.
- The QoS of the network is generally not considered in the SLA of the Cloud.

Layer 2 : Network Layer
SLA

- Service Level Agreement (SLA):
  - Commitment between the Provider and the customer (end user)
  - Defines the expected level of quality of service.
  - It includes some metrics like :
    - Availability
    - Bandwidth
    - Response time
    - security

Layer 3 : Management Layer

- It plays the role of an interface between the datacenter and the customer.
- Responsible for resource allocation, virtualization, orchestration, and monitoring.
- Contains middleware, APIs, hypervisors, cloud operating systems, and management consoles.
- Provides elasticity (scaling up/down), load balancing, and billing.
- It monitors compliance with the SLA. SLA violations result in penalties for the cloud provider

Layer 4 : Hardware resource Layer

- **The foundation**: physical resources inside datacenters.
- Includes servers, storage systems, processors, and networking devices.
- Virtualization transforms these into scalable resources for IaaS.
- The layer the most concerned by the SLA
  - Ex : delay should comply with the SLA
- To comply with SLA,
  - Datacenter network should support very high bandwidth.
  - Data transfer algorithms should be very efficient
- SLA violations result in penalities for the cloud provider.

Benifits and Limits

Any ideas?

Benefits

Some of the typical benefits are listed below:

- Cost savings (pay as you go)
- Increased storage
- Scalability/Flexibility
- Time saving
- Redundant storage
- Faster recovery
- 24/7 access to data
- Works on any device

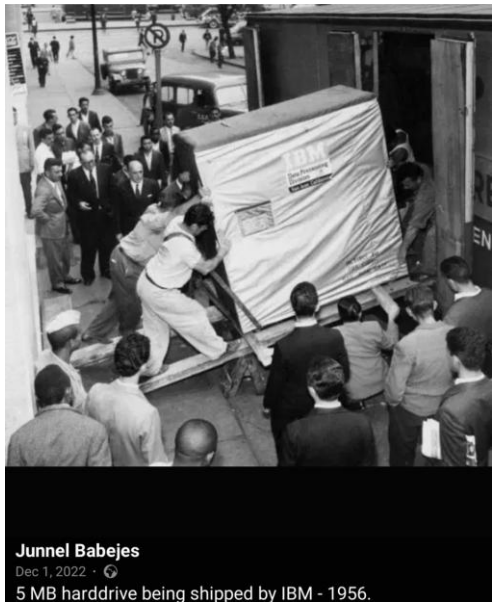

1966: 5MB of data:

- 62500 punch cards
- 4 days to load them

Cost savings

- The billing model is **pay as you go**
- The **infrastructure** is **not purchased** thus lowering maintenance
- **Initial expense** and recurring expenses **are much lower** than traditional computing

Increase storage

- With the massive infrastructure that is offered by Cloud providers today, **storage** and **maintenance** of **large volumes** of data is achievable

- **Sudden workload spikes** are also **managed effectively** and **efficiently**, since the cloud can scale dynamically

Flexibility

- With enterprises having to adapt, even more rapidly, to **changing business conditions**, speed to deliver is critical

- Cloud computing stresses on getting applications to **market very quickly**, by using the most **appropriate building blocks necessary for deployment**

Flexibility

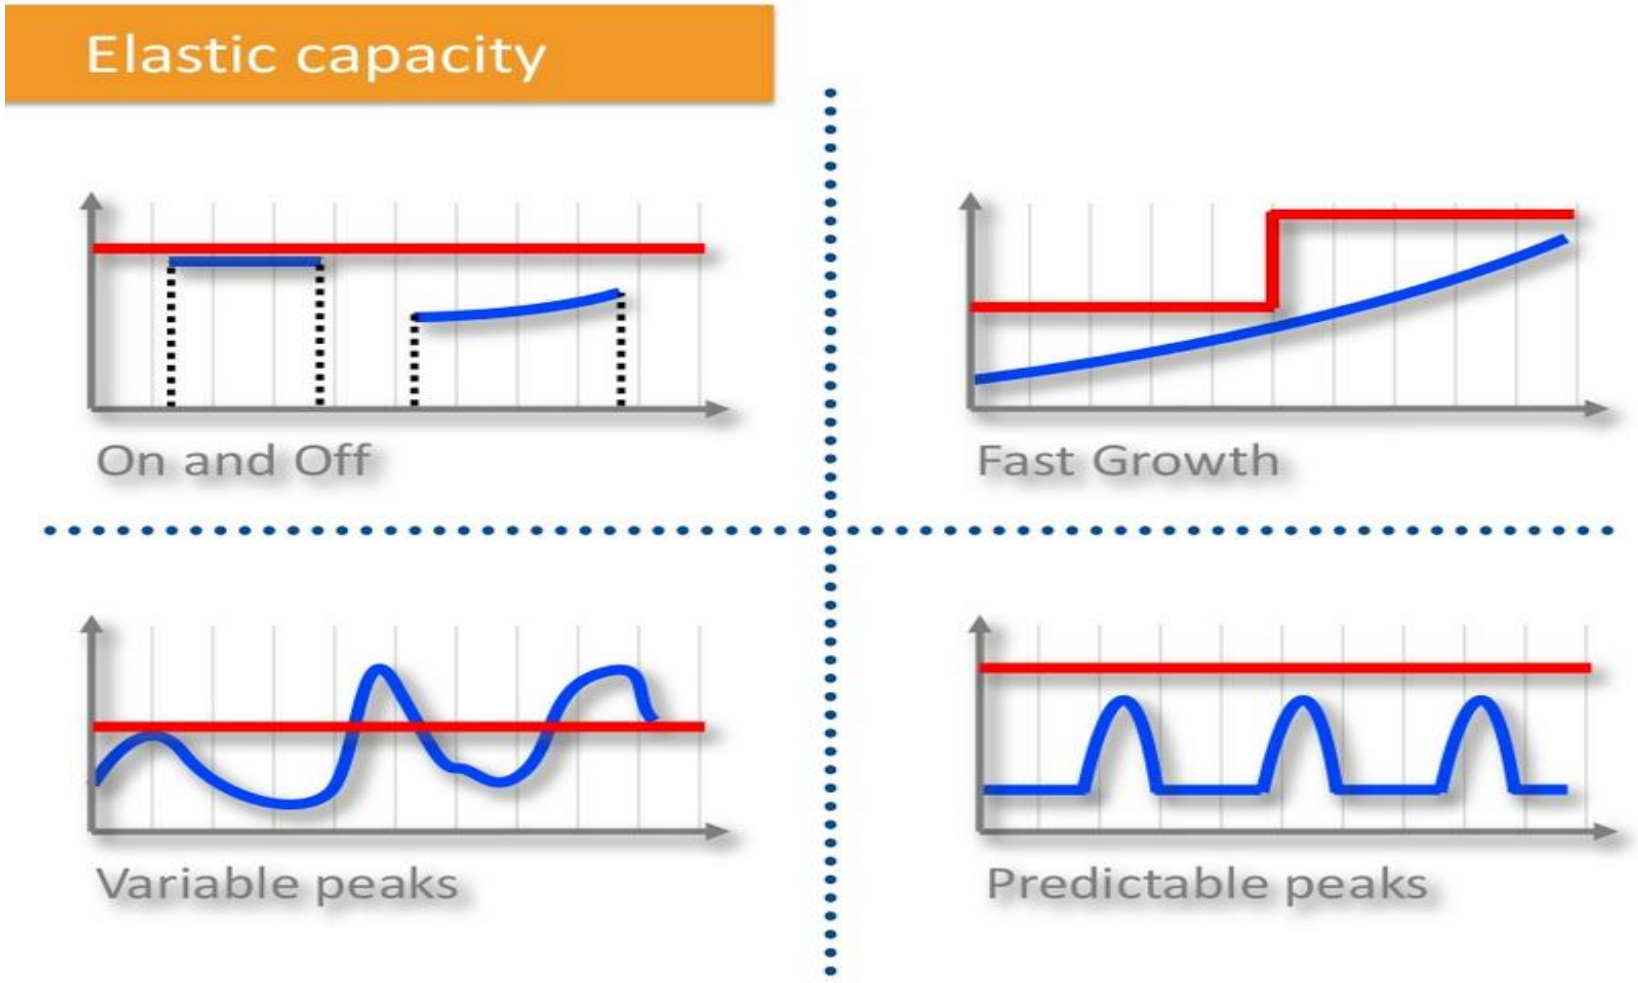

Flexibility (cont'd)

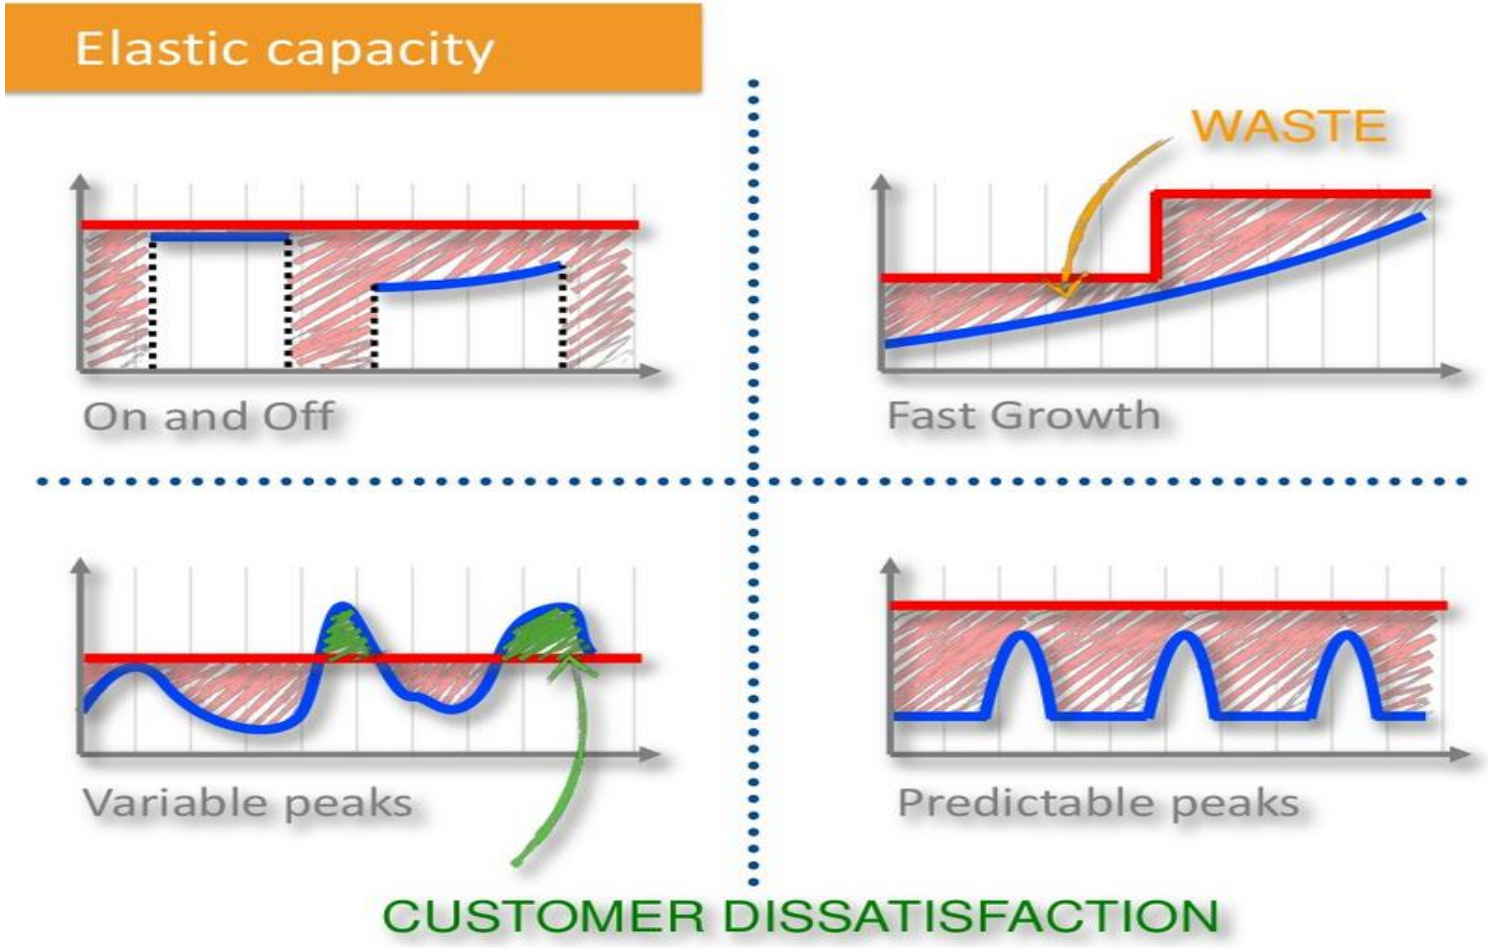

Limits & disadvantages

- Despite its growing influence, concerns regarding cloud computing still remain.

- Some common limits are:
  - Data Protection (Security and privacy)
  - Data Recovery and Availability
  - Management Capabilities
  - Regulatory and Compliance Restrictions
  - Requires a constant internet connection

Data protection

- **Data Security and privacy** are crucial elements that warrant scrutiny
- **Enterprises fear losing data** to competition and the **data confidentiality** of consumers
- Moreover, the actual storage location is not disclosed, adding onto the security concerns of enterprises
- **Service providers** are responsible for **maintaining data security** and **enterprises would have to rely on them**

Cloud computing = Your data in the computer of someone else

Solution : Use Encryption to secure data

Data recovery and availability

- All business applications have **Service Level Agreements** (SLA) that are stringently followed

- Cloud providers must serve
  - Appropriate clustering and Failover
  - Data Replication
  - System monitoring
  - Maintenance
  - Disaster recovery
  - Capacity and performance management

Management capabilities

- Despite there being multiple cloud providers, the **management of platform** and infrastructure is still a **“working in progress” phase**

- Features like „auto-scaling“ for example, are a crucial requirement for many enterprises

- There is huge potential to improve on the scalability and load balancing features provided today

Regulatory and compliance restrictions

- In some of the European countries, Government regulations **do not allow customer's personal information** and other sensitive information to be physically **located outside the state or country**
- In order to meet such requirements, **cloud providers** need to setup a **data center** or a storage site exclusively **within the country** to comply with regulations.
- Having such an infrastructure may not always be feasible and is a big challenge for cloud providers

Cloud limits study

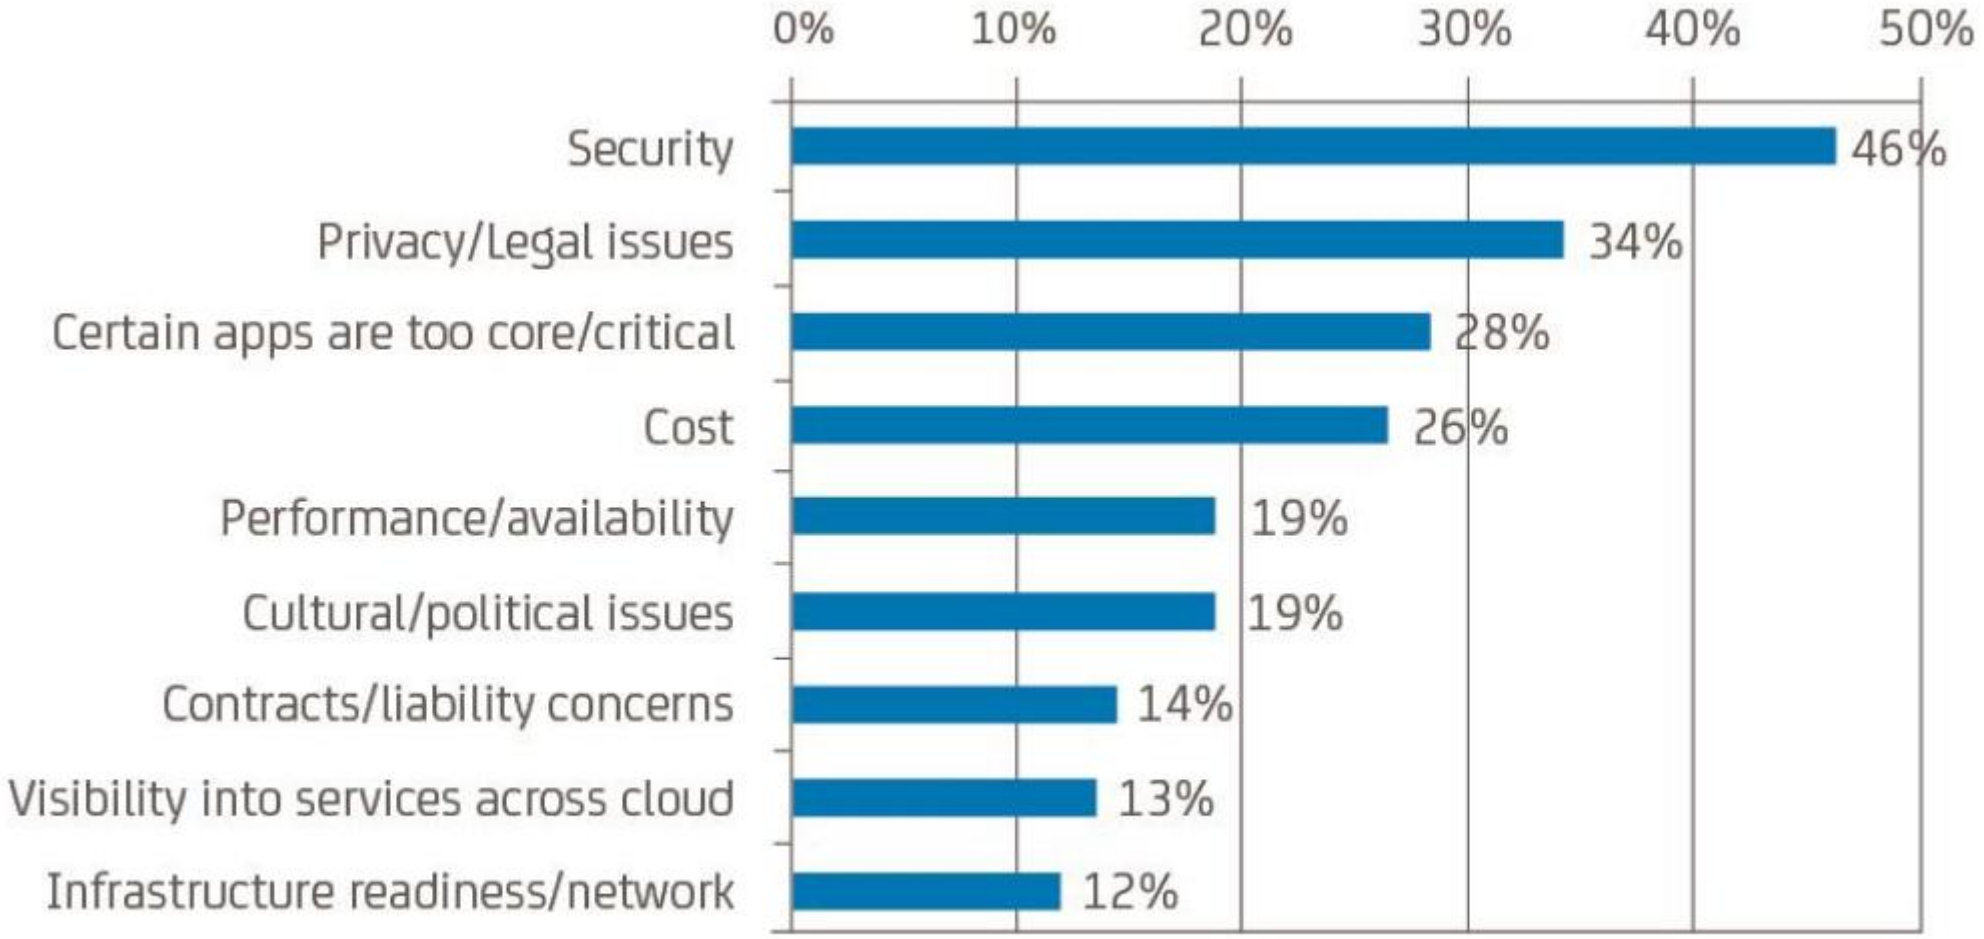

http://cloud-computing.tmcnet.com/features/articles/351722-cloud-adoption-exceeding-expectations-new-global-study-finds.htm

Cloud Computing Applications

Applications

1. **Business & Enterprise Applications**
   - **ERP :** SAP S/4HANA Cloud, Oracle NetSuite.
   - **CRM :** Salesforce, Zoho CRM.
   - **Collaboration tools**: Google Workspace, Microsoft 365,
2. **Storage, Backup & Disaster Recovery**
   - **Store and back up :** Google Drive, Dropbox, OneDrive, Amazon S3, Backblaze.
   - Disaster recovery services: AWS Backup, Azure Site Recovery.
3. **Software as a Service (SaaS)**
   - Gmail, Zoom, Canva, Shopify.
4. Infrastructure & Development
   - IaaS: Renting VMs, storage, networking (AWS EC2, Azure VMs).
   - PaaS: Platforms for app development (AWS Elastic Beanstalk, Google App Engine).
   - DevOps & CI/CD: GitHub Actions, GitLab CI, Jenkins on cloud.

Applications

5. Machine Learning & Artificial Intelligence
   - Training and deploying ML/AI models in the cloud.
   - Services: AWS SageMaker, Google AI Platform, Azure ML Studio.
   - Applications: Chatbots, image recognition, predictive analytics.
6. Healthcare & Life Sciences
   - Cloud-hosted Electronic Health Records (EHRs).
   - Genomics data analysis.
   - Telemedicine platforms.
7. Education
   - E-learning platforms: Google Classroom, Moodle Cloud, Coursera.
8. Government & Smart Cities
   - Smart traffic, smart grids, public safety monitoring.

The Cloud Computing market

Top 3 Cloud providers

- Amazone Web Service (AWS)
- Microsoft Azure
- Google Cloud (GCP)

The marcket
Worldwide

Cloud Provider	Market share Q2-2022
Amazon Web Services	34%
Azure	21%
Google Cloud	10%
Alibaba Cloud	5%
IBM Cloud	4%
…	

The marcket
in china

Cloud Provider	Market share Q1-2022
Alibaba Cloud	36.7%
Huawei Cloud	18%
Tencent Cloud	15.7%
Baidu AI Cloud	8.4%
Autres fournisseurs (AWS, Azure, etc.)	21.2%

In Algeria !!!

- **Data Centers & Infrastructure**
  - Two major **national data centers** are under construction: one in Mohammadia (≈ 80% complete) and one in Blida (≈ 50%).
  - **AI-oriented high-performance computing data center** is being built in Oran, for research, startups, and institutions.
- **Cloud Platforms**
  - CERIST has launched its own cloud platform (for data hosting, processing …).
  - Some local cloud service providers (LSPs) and companies are offering cloud or hosting services inside the country.
  - Some big organizations, including government bodies, are investing in private clouds or cloud-like infrastructure (private data centers). Sonatrach is one such example.
- **Regulations & Data Sovereignty**
  - The law **18-07** (on protection of personal data) is in effect (since August 2023). This raises requirements for storing or processing certain data within Algeria and ensuring data protection.

What’s next

What's next? Fog Computing

Fog computing is a **decentralized computing infrastructure** in which data, compute, storage and applications are located somewhere **between** the **data source** and the **cloud**. Like edge computing, fog computing brings the advantages and power of the cloud closer to where data is created and acted upon.

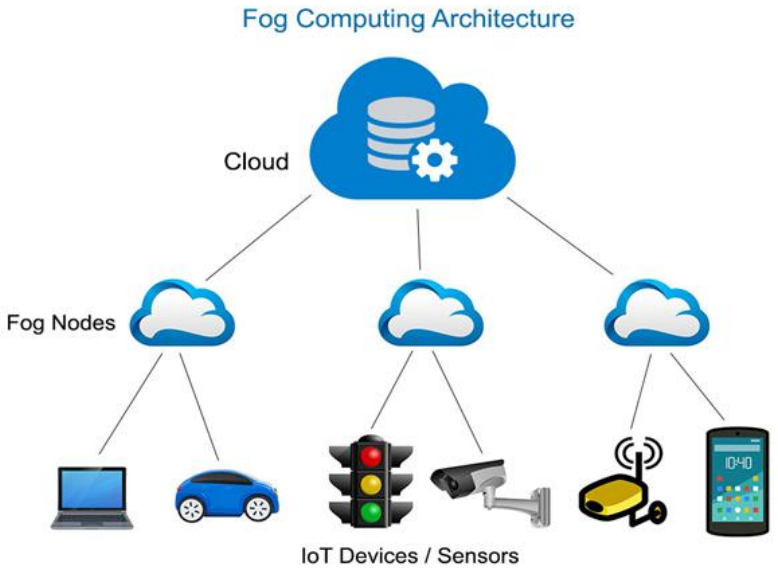

What's next? FemtoClouds

Online user-driven task scheduling for FemtoClouds

Cosimo Anglano, Massimo Canonico and Marco Guazzone
{cosimo.anglano,massimo.canonico,marco.guazzone}@uniupo.it
Computer Science Institute, DiSIT, University of Piemonte Orientale, Italy

***Abstract*—In Fog Computing, FemtoClouds are emerging computing systems consisting of a set of heterogeneous mobile devices whose users allow to run tasks offloaded by other users. FemtoClouds are well suited to run Bag-of-Tasks (BoTs) applications, but they need effective scheduling algorithms that are able to deal with collections of independently-owned, heterogeneous devices that can suddenly leave the system. In this paper, we present UDFS, an online scheduling algorithm that, by combining knowledge-free task and device selection policies with suitable heterogeneity and volatility tolerance mechanisms, can effectively schedule a stream of BoT applications on FemtoClouds. We evaluate the ability of UDFS to achieve its design goals and to perform better than existing scheduling alternatives, by carrying out a thorough simulation study for a large set of realistic scenarios. Our results indeed show that UDFS can effectively schedule a stream of BoT applications on FemtoClouds, and it can do so more effectively than existing scheduling alternatives.**

***Index Terms*—FemtoCloud, Scheduling policy, Fog computing.**

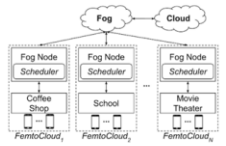

Fig. 1. FemtoCloud system architecture.

Figure 1 shows the architecture of a typical FemtoCloud system, consisting of a set of distinct FemtoClouds where each FemtoCloud is coordinated by a *Fog Node*, an always-present/always-on machine which is in charge of receiving offloaded tasks from a population of users, of dispatching them on the devices it coordinates, of receiving the results

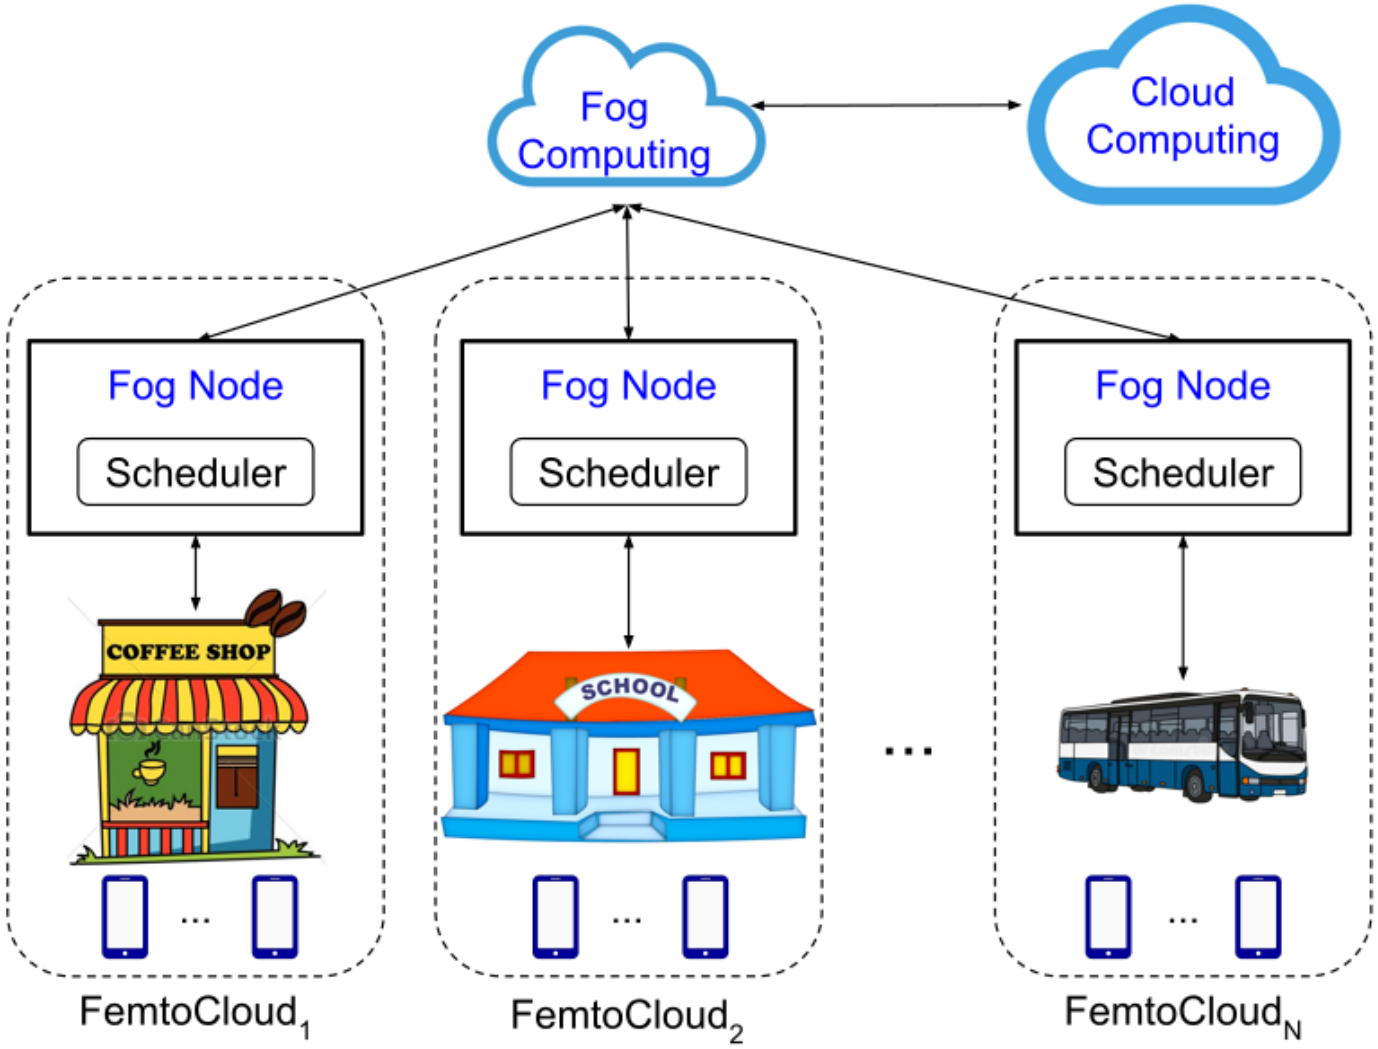

Paper published in the proc. of the 4th International Conference on Fog and Mobile Edge Computing (FMEC), 2019.

Any questions?

lamri.sayad@univ-msila.dz