unimc
UNIVERSITÀ DI MACERATA
l'umanesimo che innova

vision robotics
artificial
intelligence

Laboratorio di Cloud Computing

Lecture 3 - Introduction to Cloud Computing

Docente: Riccardo Rosati

riccardo.rosati@unimc.it

Today's Lesson

- Previous lessons recap
- Cloud Computing introduction
- Real examples of Cloud Computing technologies
- Question Time

What is IoT ?

- The *Internet of Things (IoT)* refers to the ever-growing network of physical objects that feature an IP address for internet connectivity, and the communication that occurs between these objects and other Internet-enabled devices and systems.

- In simple words, *Internet of Things (IoT)* is an ecosystem of connected physical objects that are accessible through the internet.

- It is also referred to as *Machine-to-Machine (M2M*), *Skynet* or *Internet of Everything*.

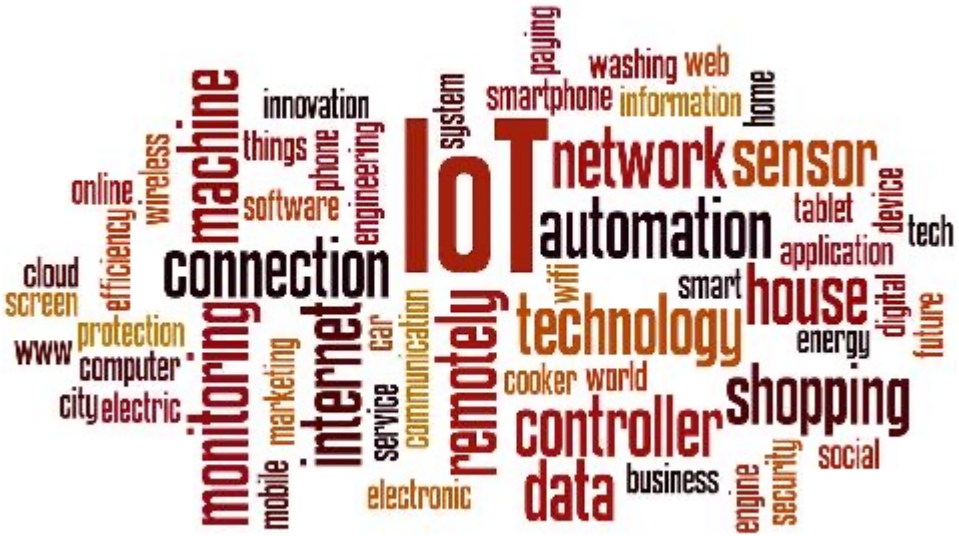

Why IoT ?

Smart Systems and Internet of Things are driven by a combination of:

1) Sensors
2) Connectivity
3) People & Processes

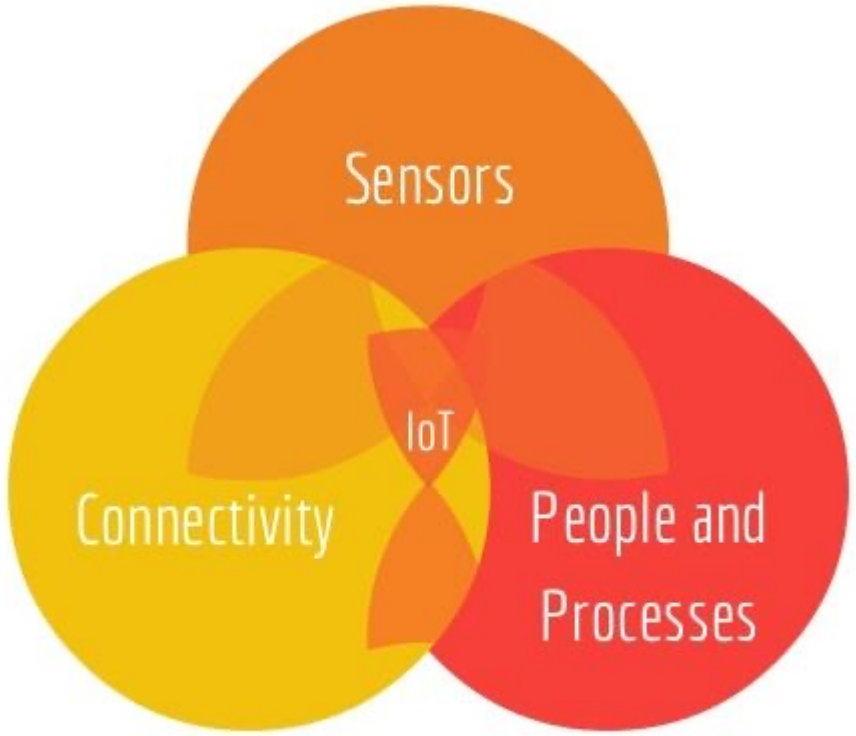

IoT applications

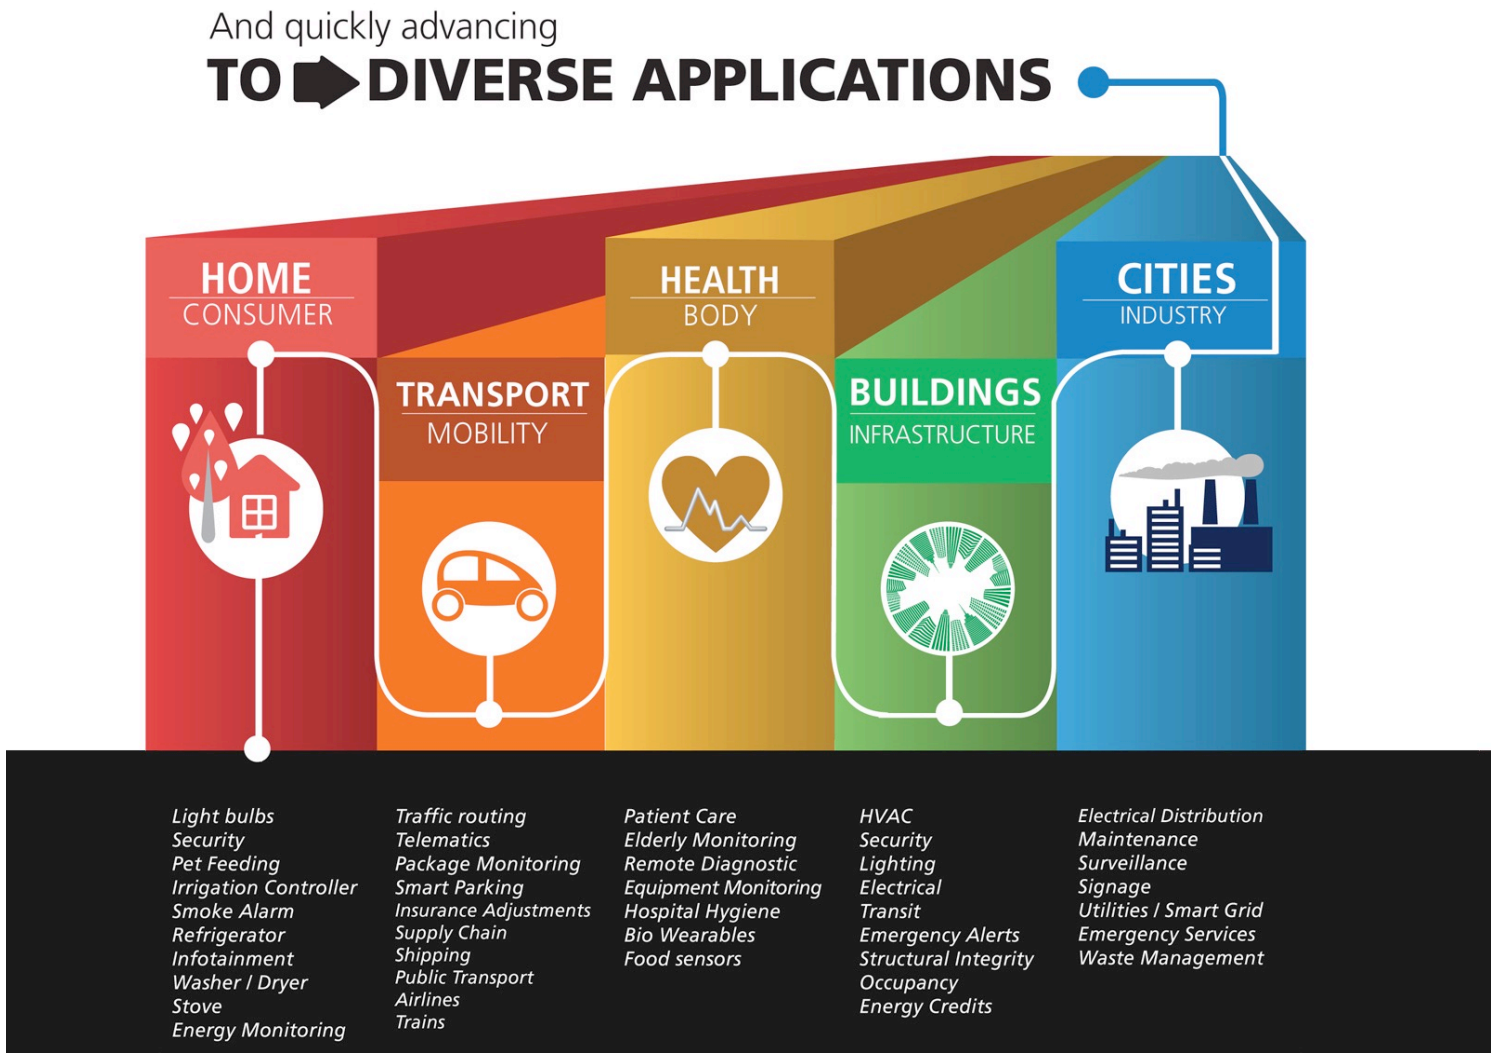

IoT issues


WHAT WOULD CONCERN YOU ABOUT A WORLD OF CONNECTED DEVICES?
PRIVACY 62%
SECURITY 54%
PHYSICAL SAFETY 27%
UNABLE TO REPAIR 24%
MACHINES TAKING OVER THE EARTH 21%
NOT KNOWING HOW TO USE THEM 17%
NO TANGIBLE BENEFITS 11%

Cloud Computing introduction

Cloud Computing: an historical view

Cloud Computing concept is the result of the evolution of the computing concept driven by the technology improvements and by users requirements.

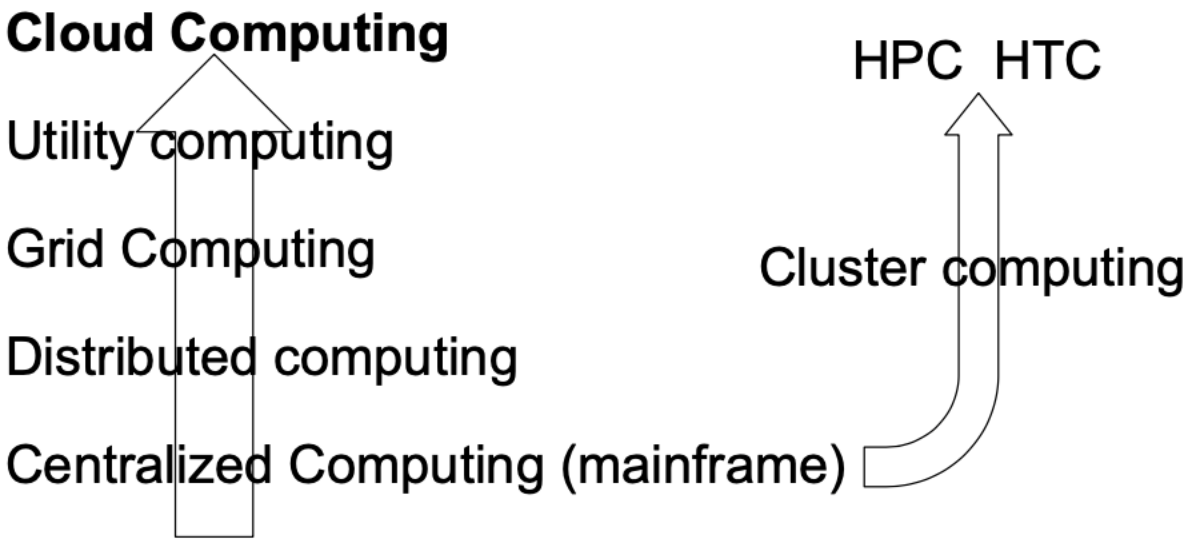

The concept of computing

«Computing is the process of using computer technology to complete a given goal-oriented task. [...] Computing may encompass the design and development of software and hardware systems for a broad range of purposes»

(Association of Computing Machinery, 2005)

Computing as a “numerical laboratory”

- Each scientific instrument is critically dependent on computing for sensor control, data processing, international collaboration, and access
- Computational modeling and data analytics are applicable to all areas of science and engineering
- Capture and analyze the torrent of experimental data being produced by a new generation of scientific instruments

What is Cloud?

Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (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 five essential characteristics, three service models, and four deployment models. (2011 - NIST)

5 Essential characteristics

1 - On-demand self-service. A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.

computing capabilities are: server time, network storage, number of servers etc.

2 - Broad network access. Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, tablets, laptops, and workstations).

5 Essential characteristics

3 - Resource pooling. The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, and network bandwidth.

5 Essential characteristics

4 - Rapid elasticity. Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time.

5 - Measured service. Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

3 Service Models

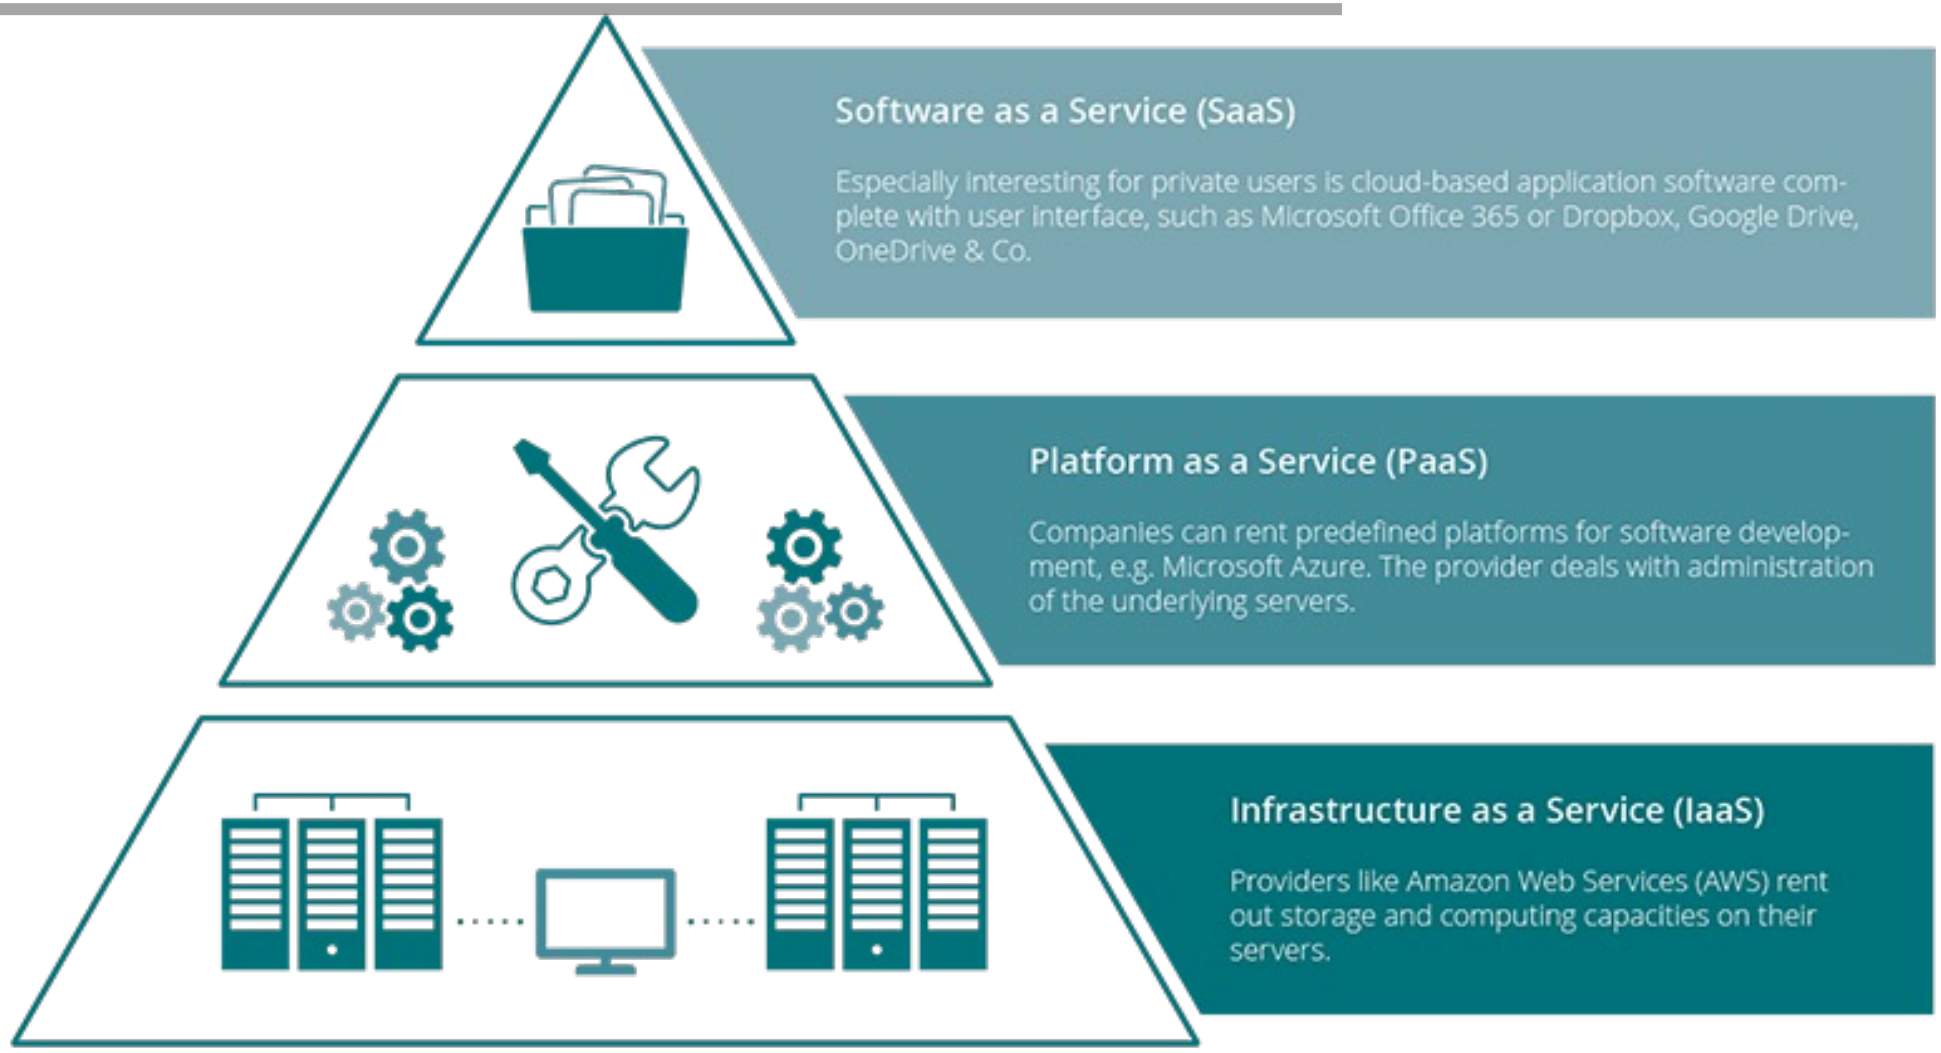

Software as a Service (SaaS)

Especially interesting for private users is cloud-based application software complete with user interface, such as Microsoft Office 365 or Dropbox, Google Drive, OneDrive & Co.

Platform as a Service (PaaS)

Companies can rent predefined platforms for software development, e.g. Microsoft Azure. The provider deals with administration of the underlying servers.

Infrastructure as a Service (IaaS)

Providers like Amazon Web Services (AWS) rent out storage and computing capacities on their servers.

3 Service Models

1 - Software as a Service (SaaS). The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure . The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited userspecific application configuration settings.

3 Service Models

2 - Platform as a Service (PaaS). The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.

Google App Engine

Microsoft Azure

3 Service Models

3 - Infrastructure as a Service (IaaS). The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls).

4 Deployment Strategies

1 - Private cloud. The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.

2 - Community cloud. The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.

4 Deployment Strategies

3 - Public cloud. The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.

4 - Hybrid cloud. The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).

4 Deployment Strategies

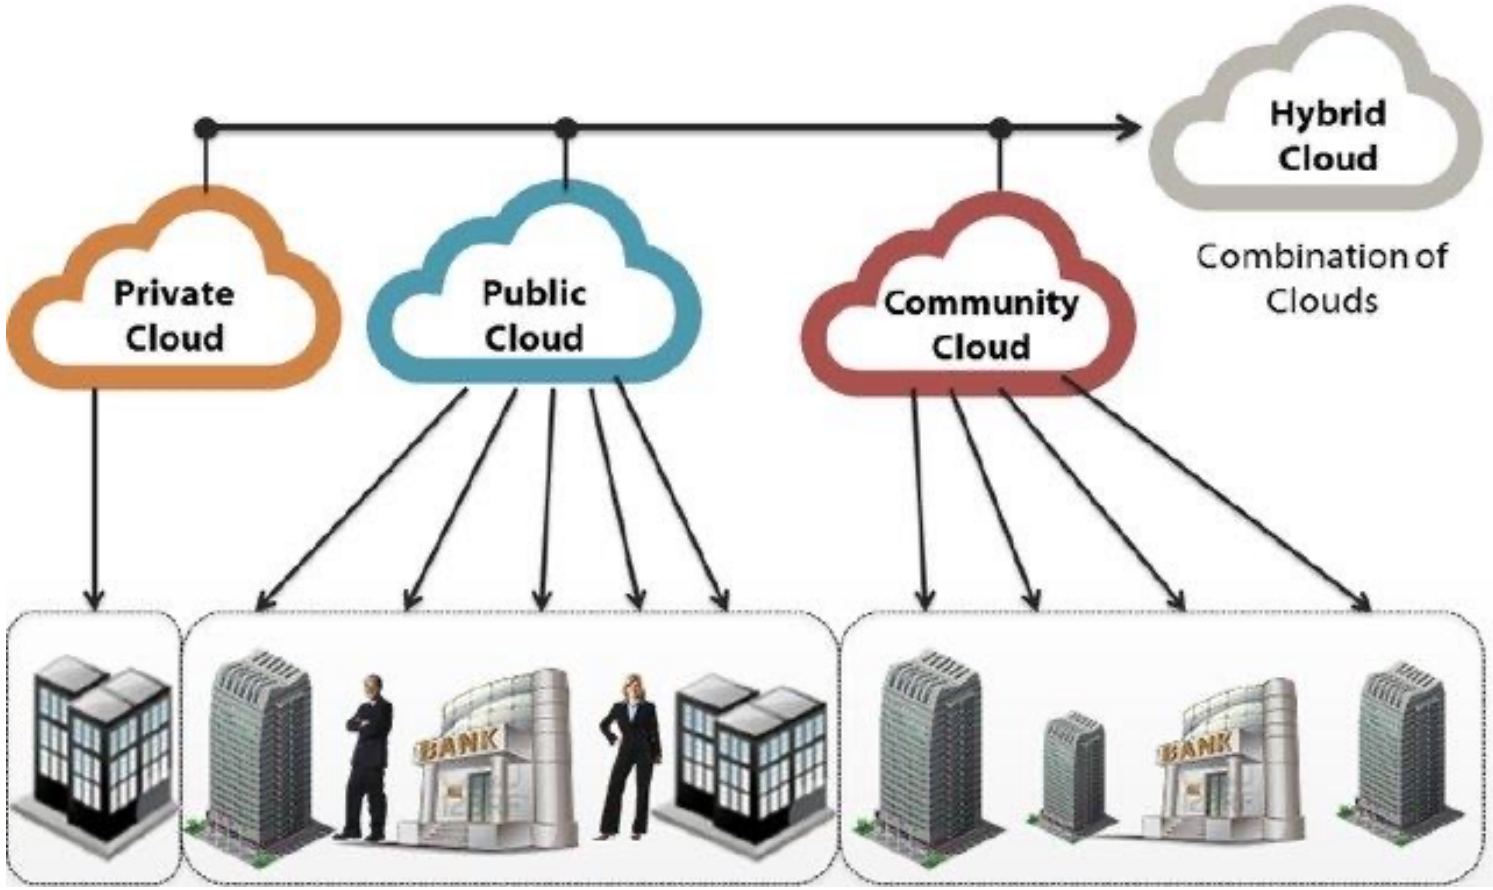

Advantages

- IT infrastructures is no longer a barrier
- Easier to test different solutions
- No need to wait for provisioning
- Shorter development cycles
- No initial investment needed
- No termination fees
- No commitments
- Clear pricing model

Advantages

- Pay for servers “by-the-hour” (on-demand)
- Pay for storage “per Gigabyte” per month
- Pay for data transfer “per Gigabyte”
- Easy to turn resources on/off (running costs)

Advantages

- Automation, less repetitive tasks (“70/30” rule)
- Better management tools
- Focus on your business
- No need to rebuild from scratch, reuse
- No need to work “Undifferentiated heavy lifting”

Traditional IT provisioning

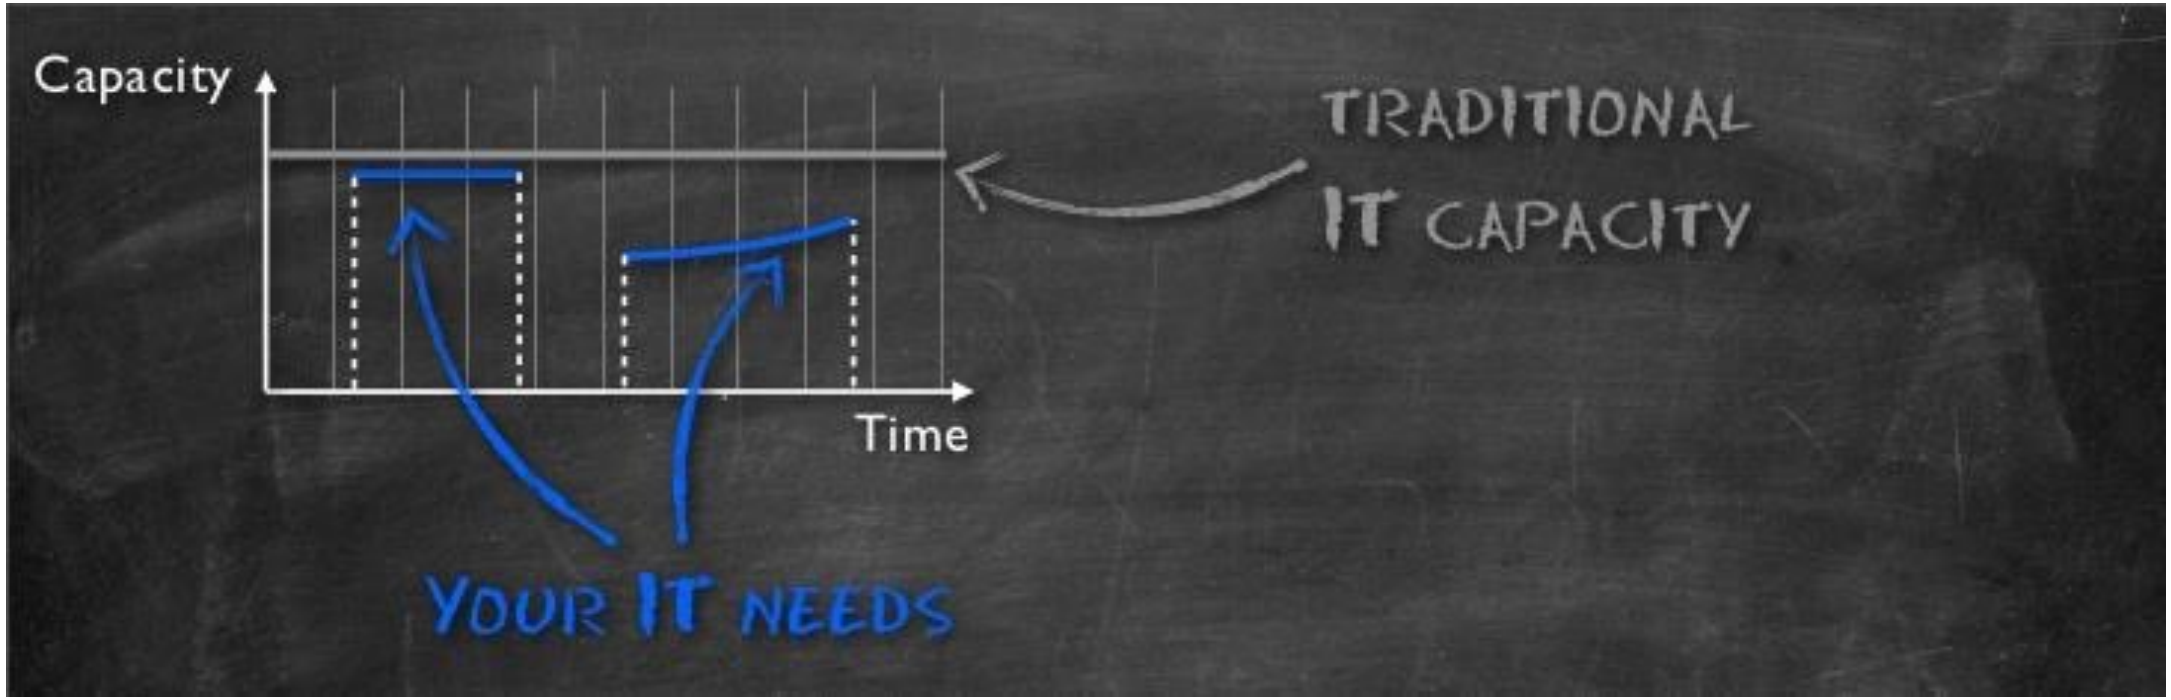

Traditional IT provisioning

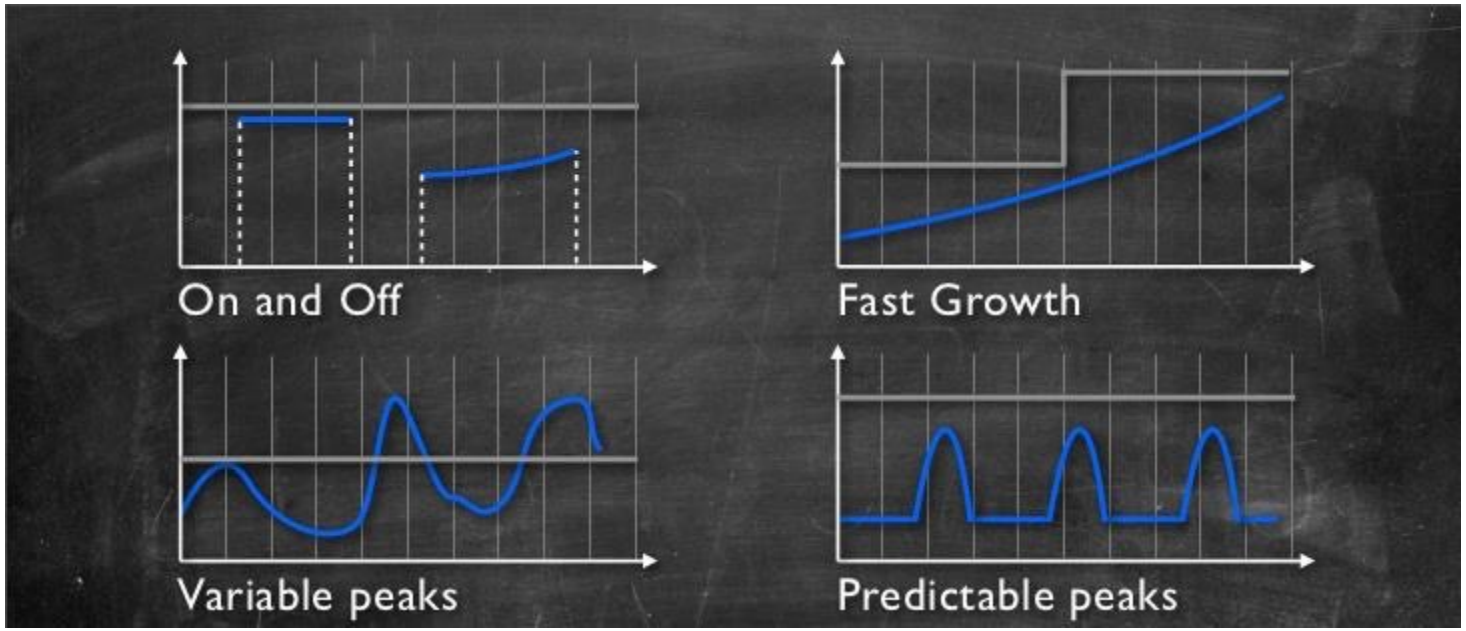

We may have different resource demands. Think about resources in a broad sense, everything can be a resource: storage, computational power, memory…

Traditional IT provisioning

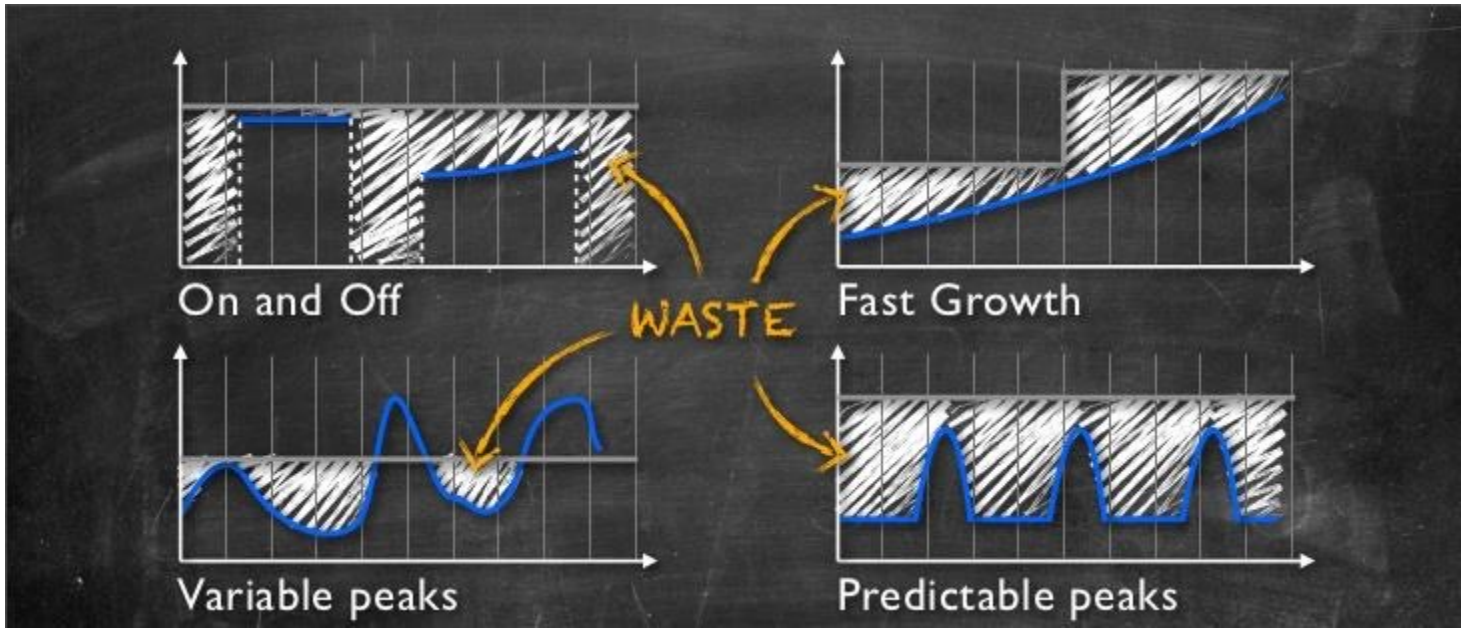

Traditional IT provisioning

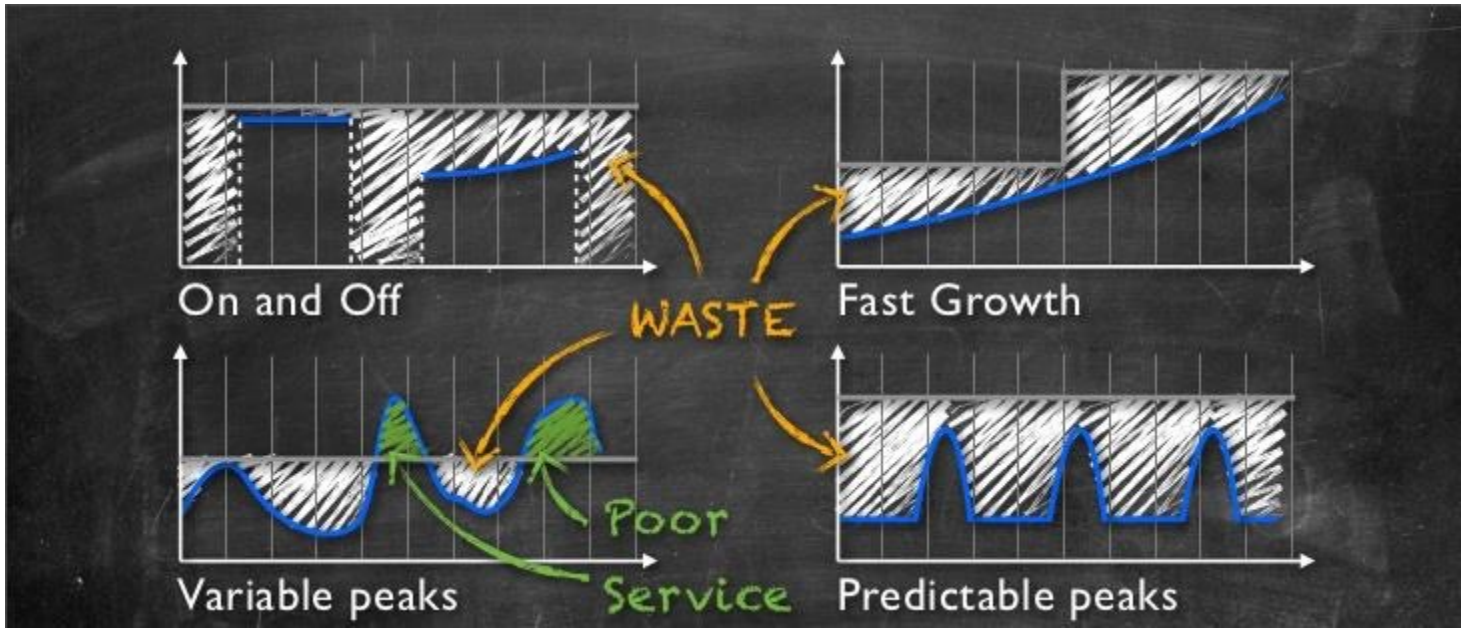

In the traditional IT increasing the capacity (the grey line) means buy more hardware, but it's not just the hardware cost. Think about a company, probably it has a server room.. Air conditioning, smoke detection system, alarm system, access control policy, security, IT management employee.. Costs!

Cloud provisioning

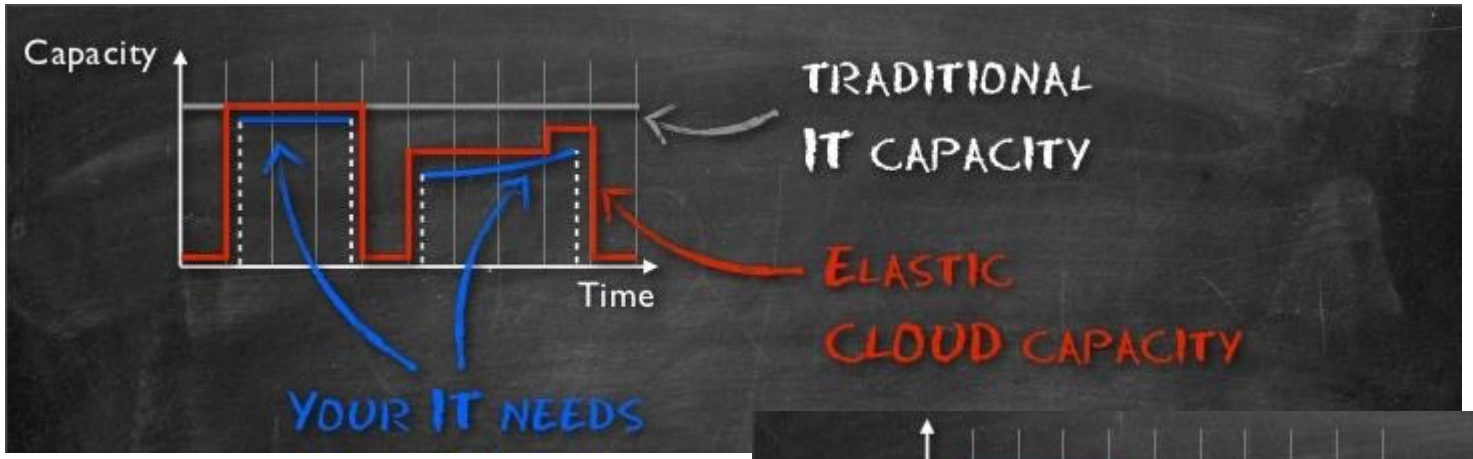

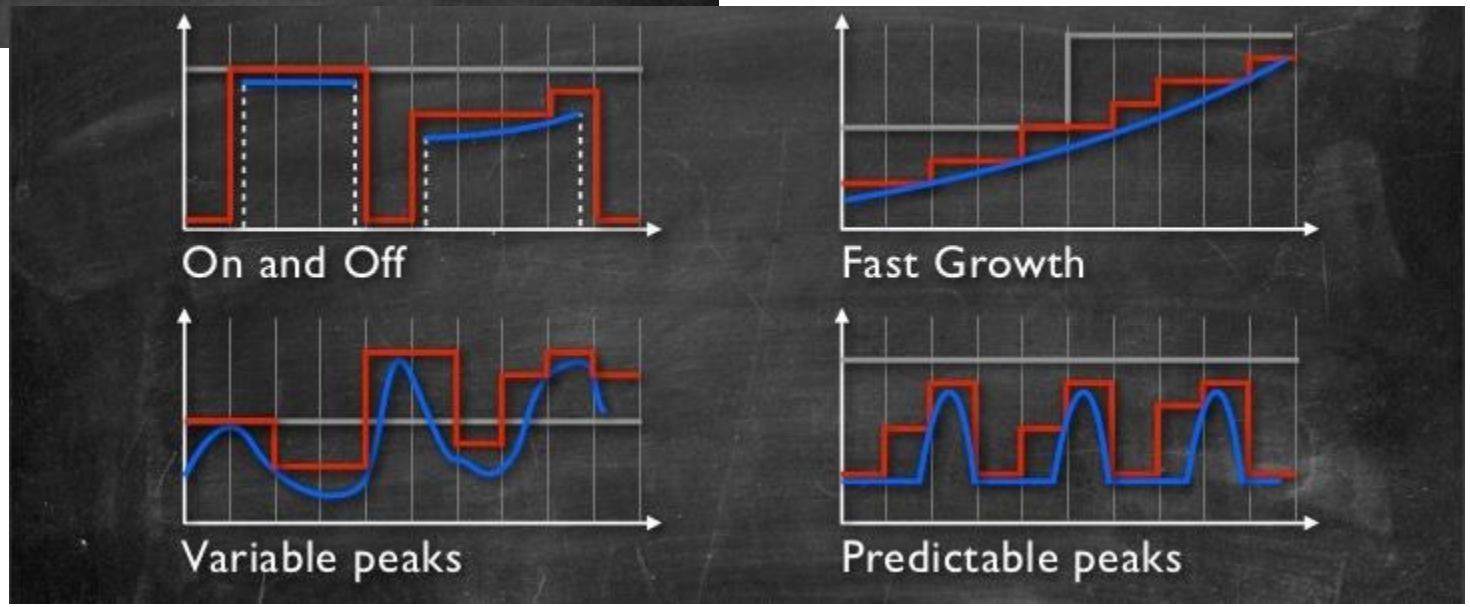

IoT - Cloud?

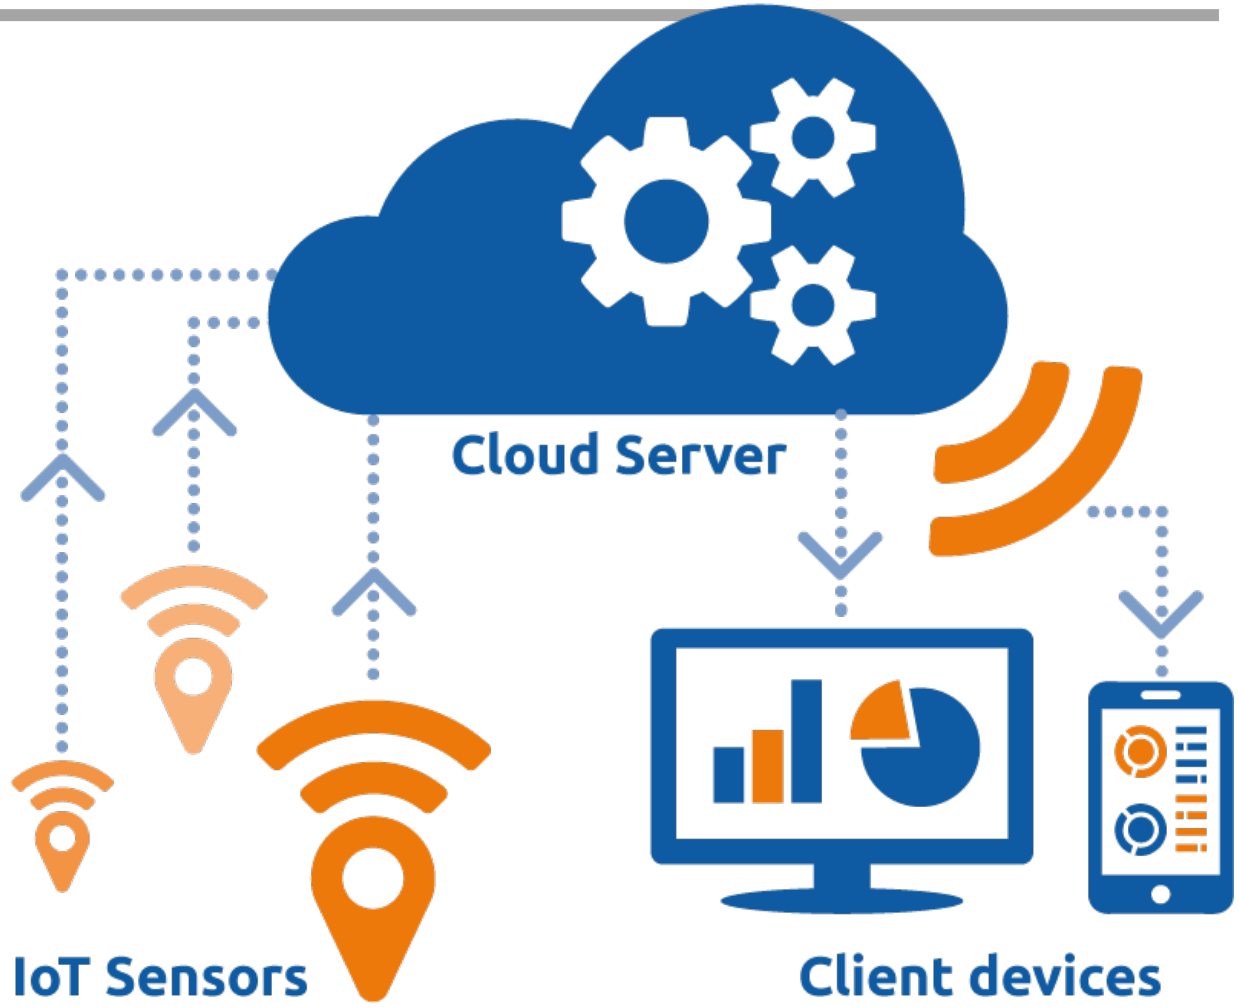

Look at the schema..Connectivity is really important!

IoT devices are the **Edge** of the network.

Edge vs Cloud Computing

We cannot process everything in the cloud!

As Internet of Things (IoT) devices become more common and incorporate more processing power, vast amounts of data is being generated on the outer “edge” of computing networks. Traditionally, the data produced by IoT devices is relayed back to a central network server, usually housed in a data center. Once that data is processed, further instructions are sent back to the devices out on the edge of the network.

There are two problems with this setup. First, it takes time for data to travel from the edge device back to the center for processing. This delay might only be a matter of milliseconds, but it can be critical. Secondly, all that data traveling back and forth between the edge and the center of the network puts tremendous strain on bandwidth. This combination of distance and high volume traffic can slow the network down to a crawl.

Edge vs Cloud Computing

Network latency can have serious consequences for IoT devices. Take, for example, self-driving cars. Autonomous vehicles gather a tremendous amount of data from their surroundings and from other devices nearby. If the vehicle's reaction time is dependent upon instructions from the computing resources at the core of the network, the slightest delay could literally be a matter of life and death. Think again about connectivity..

So.. Cloud computing is good, but it depends on the application!

Edge and Cloud usually works together, they are not mutually exclusive.

How? Virtualization..

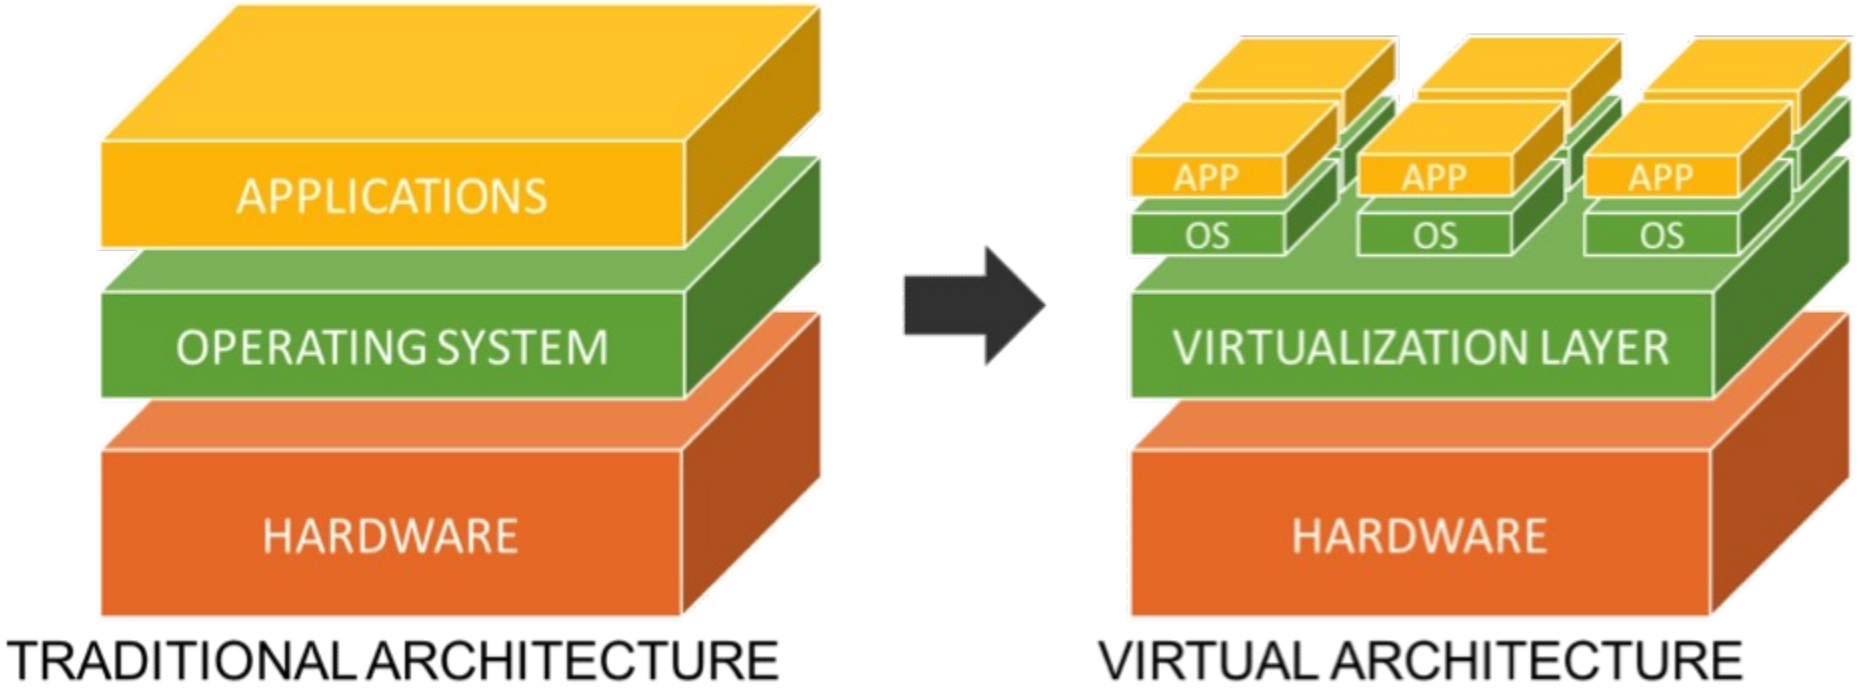

How? Virtualization..

Virtualization refers to the act of creating a virtual (rather than actual) version of something, including virtual computer hardware platforms, storage devices, and computer network resources.

In cloud computing this means that physical (hardware) components can be virtualized using appropriate software (“the hypervisor”).

The final user will have the same experience as he would use real hardware. Cloud service provider uses big hardware resources, sharing them between users using virtualization.

How? Virtualization..

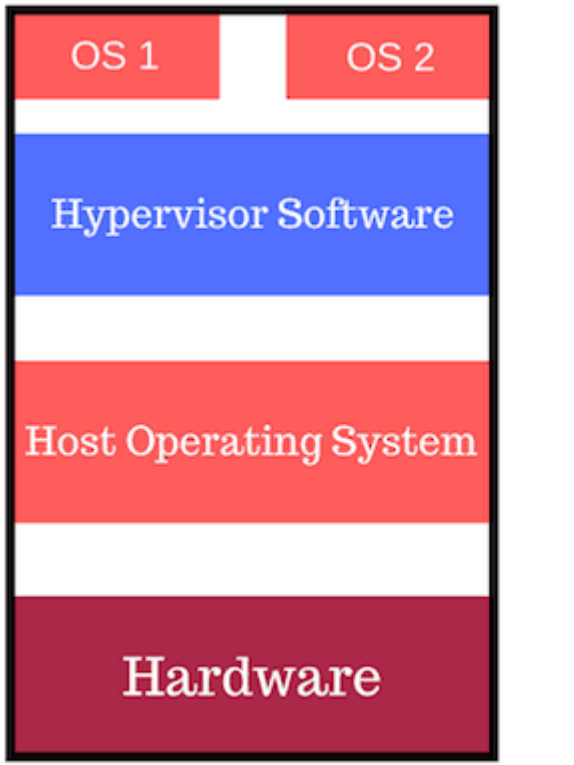

Standard Hosting Architecture

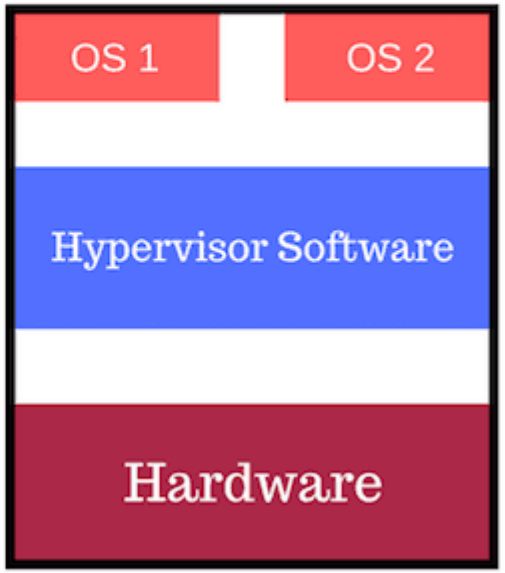

Bare Metal With Hypervisor

There are basically two possibility for virtualization. The standard way is a software that you can install on any Operating system. An example is VirtualBox that you can use on your laptop to create Virtual Machines.
The other possibility is the Bare Metal, it does not require admins to install a server operating system first. Bare-metal virtualization means the hypervisor has direct access to hardware resources, which results in better performance, scalability and stability. This is used by the cloud providers.

Virtualization: VirtualBox

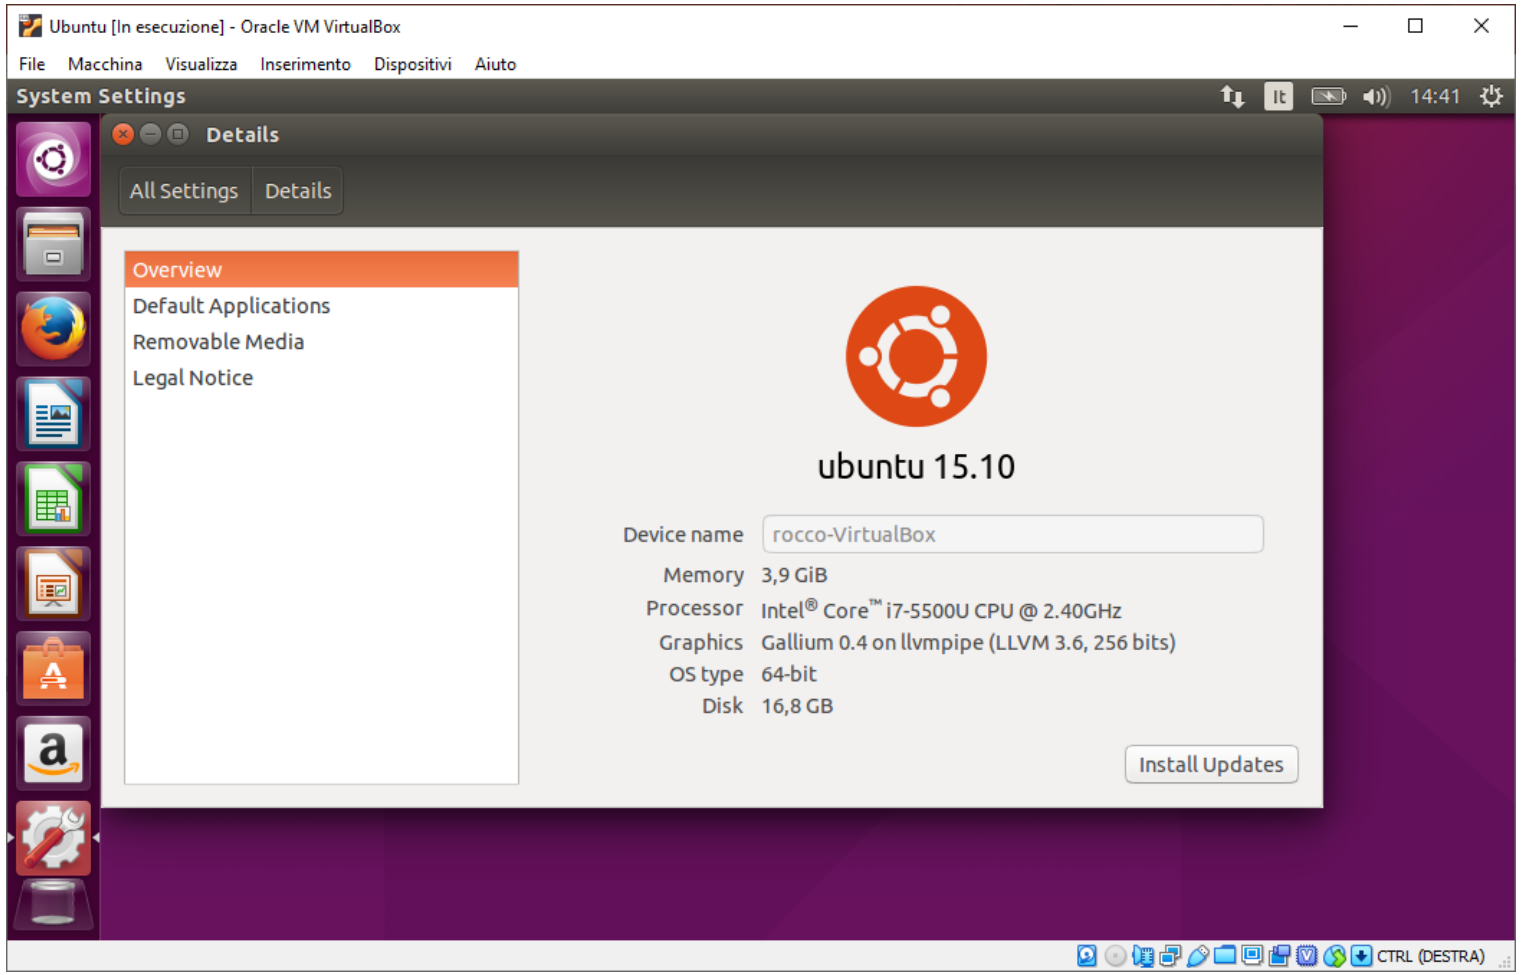

Everything is virtual inside the virtual Machine, we can see 4Gb of Memory.. But in Total my laptop has 8GB.
The same for Disk.. For the virtual machine is like having a 20Gb of Hard Disk, but the hard disk of my laptop Is 500GB..

Virtualization: VirtualBox

Main advantages of Virtual Machines are:

- Easy to create, destroy and distribute (it's just a file..)
- Can run on different hardware
- Useful for development in different architectures (windows, linux, MacOS, android..)
- No worries about the underlying hardware

AWS

https://youtu.be/jOhbTAU4OPI

AWS

amazon.com amazon.com amazon web services

Consumer Business	Seller Business	IT Infrastructure Business
Tens of millions of active customer accounts	Sell on Amazon websites	Web-scale cloud computing infrastructure for developing, deploying & operating applications
13 countries: US, UK, Germany, Japan, France, Canada, China, Italy, Brazil, Mexico, India, Spain, Australia	Use Amazon technology for your own retail website	Over 1 million registered customers in over 190 countries
	Leverage Amazon's massive fulfilment centre network	

AWS

The cloud gave Amazon 11.5 percent of its revenue in the second quarter.

AWS has become crucial to Amazon's profitability. In the second quarter of 2018 AWS' operating income totaled $1.64 billion, which was **55 percent** of all of Amazon's operating income.

AWS

Some AWS customers..

https://aws.amazon.com/it/solutions/case-studies/enterprise/

AWS has 70+ services.. Just some examples..

AWS has 70+ services.. Just some examples..

EC2: Amazon Elastic Compute Cloud. Allow users to rent virtual computers on which to run their own computer applications. EC2 encourages scalable deployment of applications by providing a web service through which a user can boot an Amazon Machine Image (AMI) to configure a virtual machine, which Amazon calls an "instance", containing any software desired. A user can create, launch, and terminate server-instances as needed, paying by the second for active servers – hence the term "elastic". EC2 provides users with control over the geographical location of instances that allows for latency optimization and high levels of redundancy.

RDS: Amazon Relational Database Service. is a distributed relational database service by Amazon Web Services (AWS). It is a web service running "in the cloud" designed to simplify the setup, operation, and scaling of a relational database for use in applications.

AWS has 70+ services.. Just some examples..

S3: Amazon Simple Storage Service is a "simple storage service" offered by Amazon Web Services (AWS) that provides object storage through a web service interface. Amazon S3 uses the same scalable storage infrastructure that Amazon.com uses to run its global e-commerce network.

Redshift: Amazon Redshift is an Internet hosting service and data warehouse product which forms part of the larger cloud-computing platform Amazon Web Services. It is built to handle large scale data sets and database migrations. Redshift differs from Amazon's other hosted database offering, Amazon RDS, in its ability to handle analytic workloads on big data.