

Cloud/Fog Computing for Smart Grid

Sabita Maharjan

Professor, University of Oslo, Norway

Feb. 29, 2024



Learning Objectives

Throughout this lecture, it is aimed for the students to be able to

- Learn the **basic concepts** and **architecture** of **cloud computing**
- Learn the **basic concepts** and **architecture** of **fog computing**
- Understand **why and how** cloud computing and fog computing can be **applied in smart grid**

Industry Invited Talk Today

Speakers

Espen Kåsin

Vice President

Services and Software, Embriq

Asle Jåsund

Senior Consultant

Embriq



Title

Digital Energy

Embriq

- Leading supplier of IT operations, software, industrial IoT technology and consultancy services
- Provides IT services to optimize the operations of energy companies

Outline



Cloud Computing




Fog Computing



Cloud/Fog Computing for Smart Grid

CLOUD COMPUTING

Why do we need cloud computing?



Comparison of electricity generation and cloud computing

UiO does not generate its own electricity but it is a user of electricity.

Similarly, IT services can be “generated” remotely by a factory-size bank of powerful computers (“servers”) and delivered over the Internet to consumers who can take as much, or as little as they need.

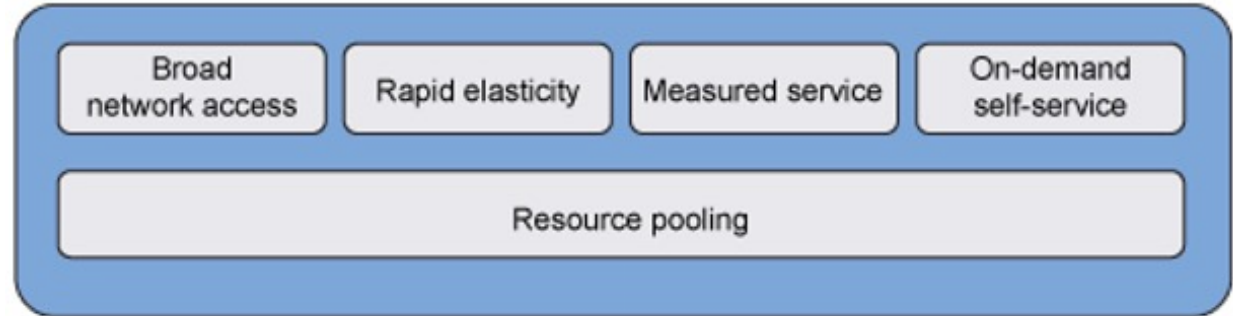
Cloud Computing definition

According to NIST (National Institute of Standard and Technology, USA)

Cloud computing is a model **for enabling 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.

NIST's 5-4-3 Principle for Cloud Computing


Five essential characteristics



Three service models




Four deployment models



The 5-4-3 principle is a simple, well-structured, and disciplined way of understanding cloud computing. 5 characteristics, 4 deployment models and 3 service models together explain the key aspects of cloud computing.

Five Essential Characteristics




On-demand self-service




Broad network access



Resource pooling

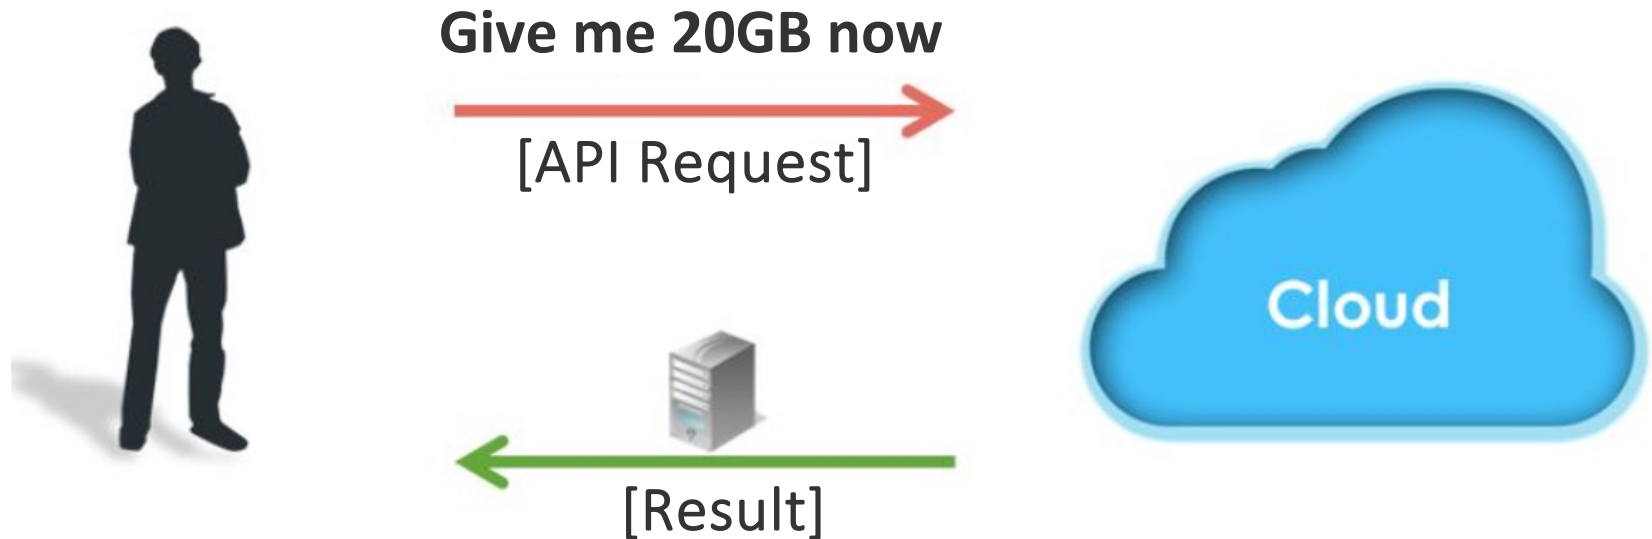


Rapid elasticity



Measured service


Five Essential Characteristics – On-demand self-service



On-demand self-service

A consumer can utilise **computing capabilities**, such as **server time** and **network storage**, as needed, without requiring human interactions with each service provider.

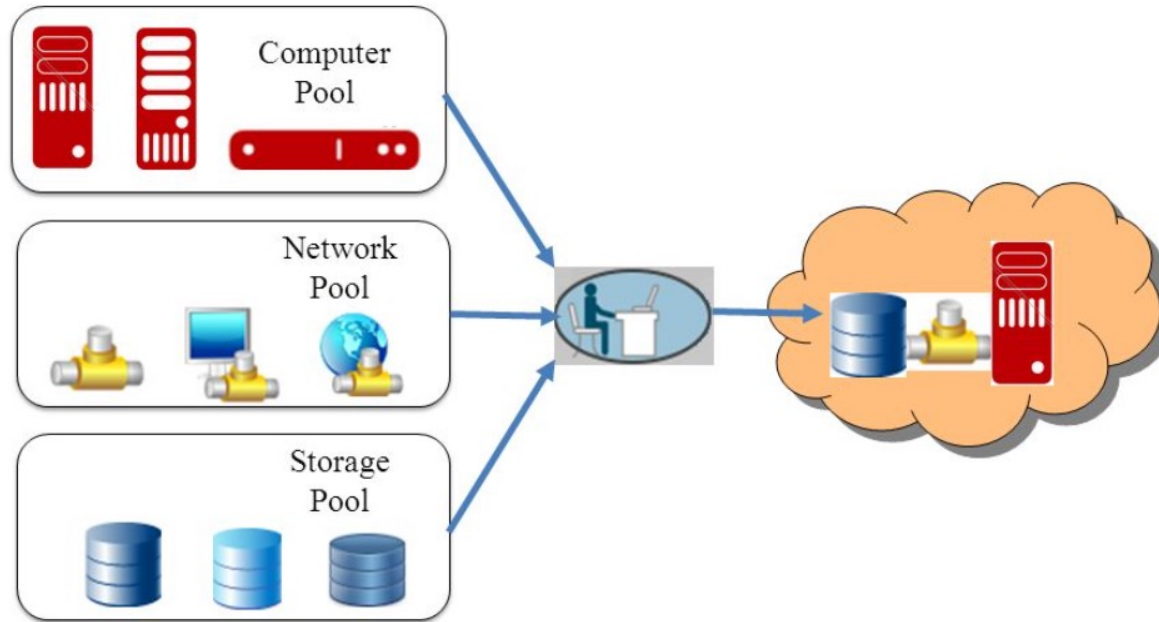
Five Essential Characteristics – Broadband networks access



Broadband network access

Capabilities are available over the network and accessed through standard mechanisms that are used by heterogeneous client platforms (mobile phones, tablets, laptops, and workstations) i.e., **numerous connectivity options**


Five Essential Characteristics – Resource pooling



Resource pooling

The provider's computing **resources are pooled** to serve multiple consumers with different **resources** and **dynamically (re)assigned** according to consumer demand. Examples of resources include storage, processing, memory, and network bandwidth.

Five Essential Characteristics – rapid elasticity



At time t_1 , there are many new visitors and more resources are provisioned

At time t_2 , there are low visitors and much less resources are provisioned

Rapid elasticity

Capabilities can be **elastically provisioned** and released to **scale rapidly** with **dynamic demand**.

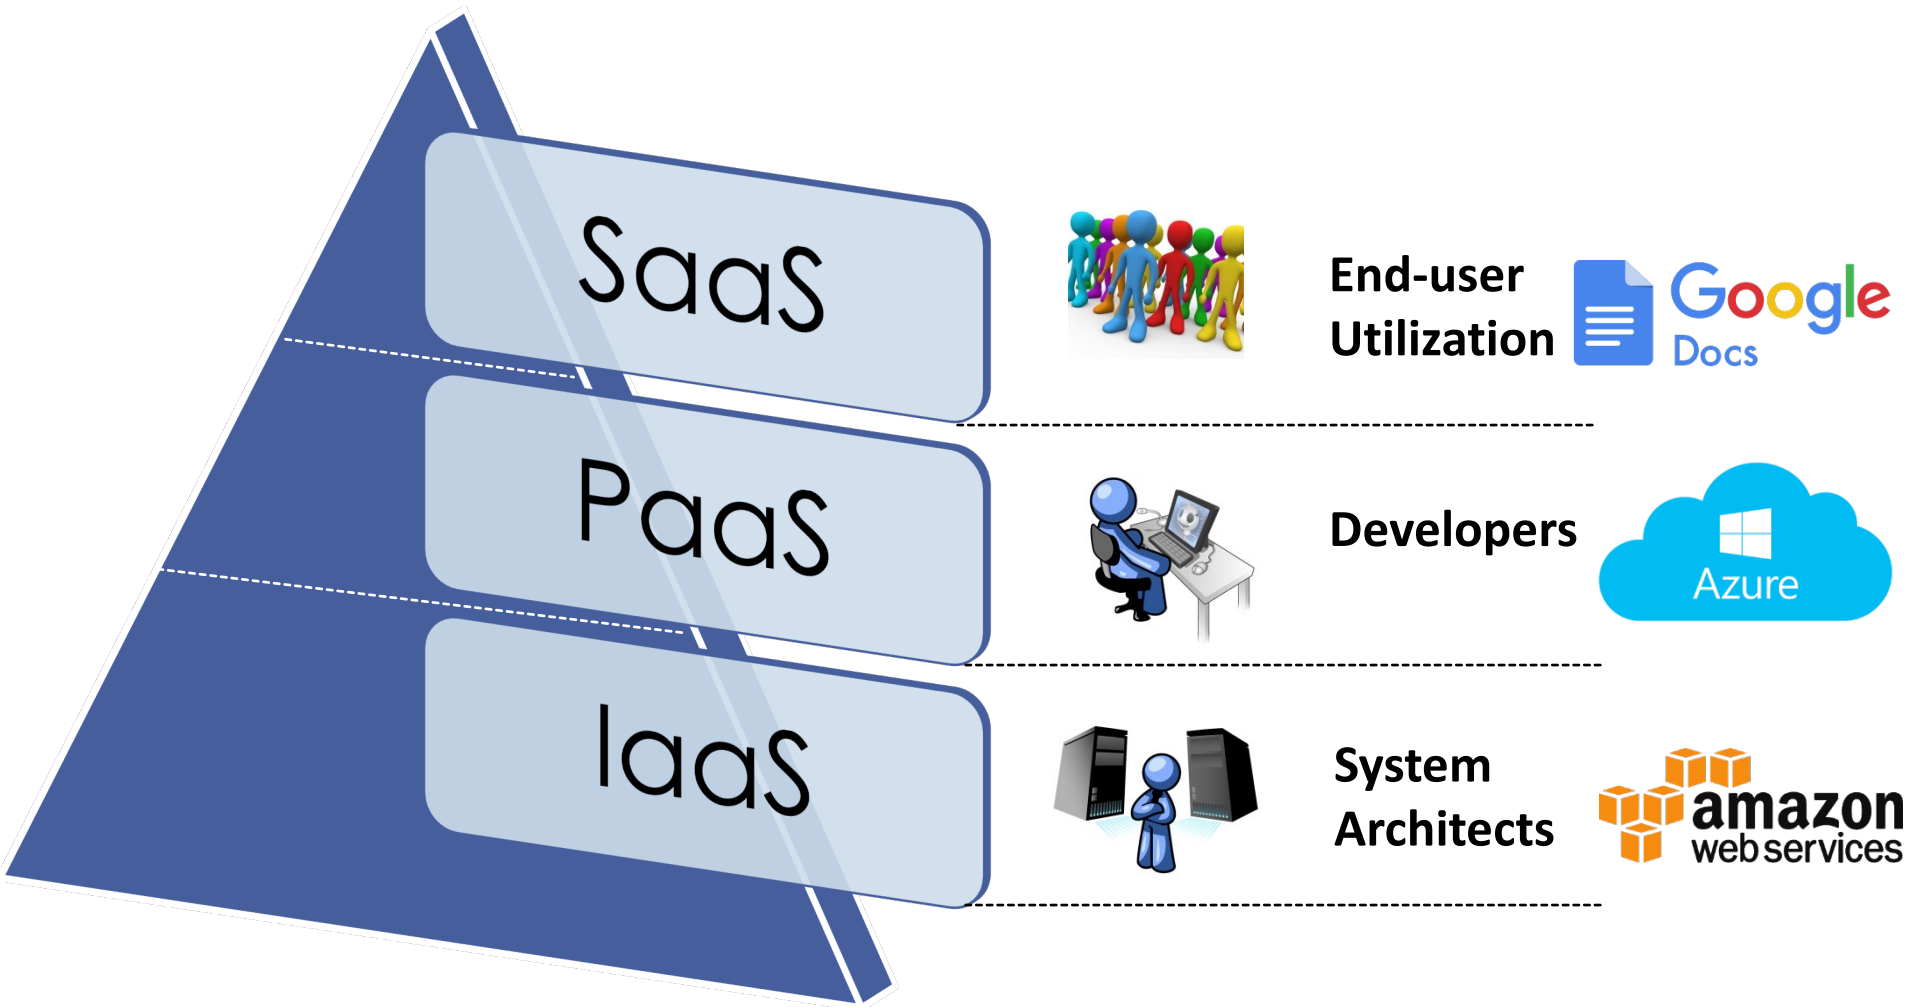
Five Essential Characteristics – measured service

Details			
Expand All Services Collapse All Services		Printer Friendly Version	
AWS Service Charges			\$9.28
<input type="checkbox"/> Amazon Elastic Compute Cloud			\$8.27
Download Usage Report »			
US East (Northern Virginia) Region			
Amazon EC2 running Linux/UNIX			
\$0.020 per Micro Instance (t1.micro) instance-hour (or partial hour)	275 Hrs		5.50
Amazon EC2 EBS			
\$0.10 per GB-month of provisioned storage	8,100 GB-Mo		0.81
\$0.10 per 1 million I/O requests	38,962 IOs		0.01
\$0.125 per GB-Month of snapshot data stored	0.069 GB-Mo		0.01
Elastic IP Addresses			
\$0.00 per Elastic IP address not attached to a running instance for the first hour	1 Hr		0.00

Measured service

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

Three Service Models

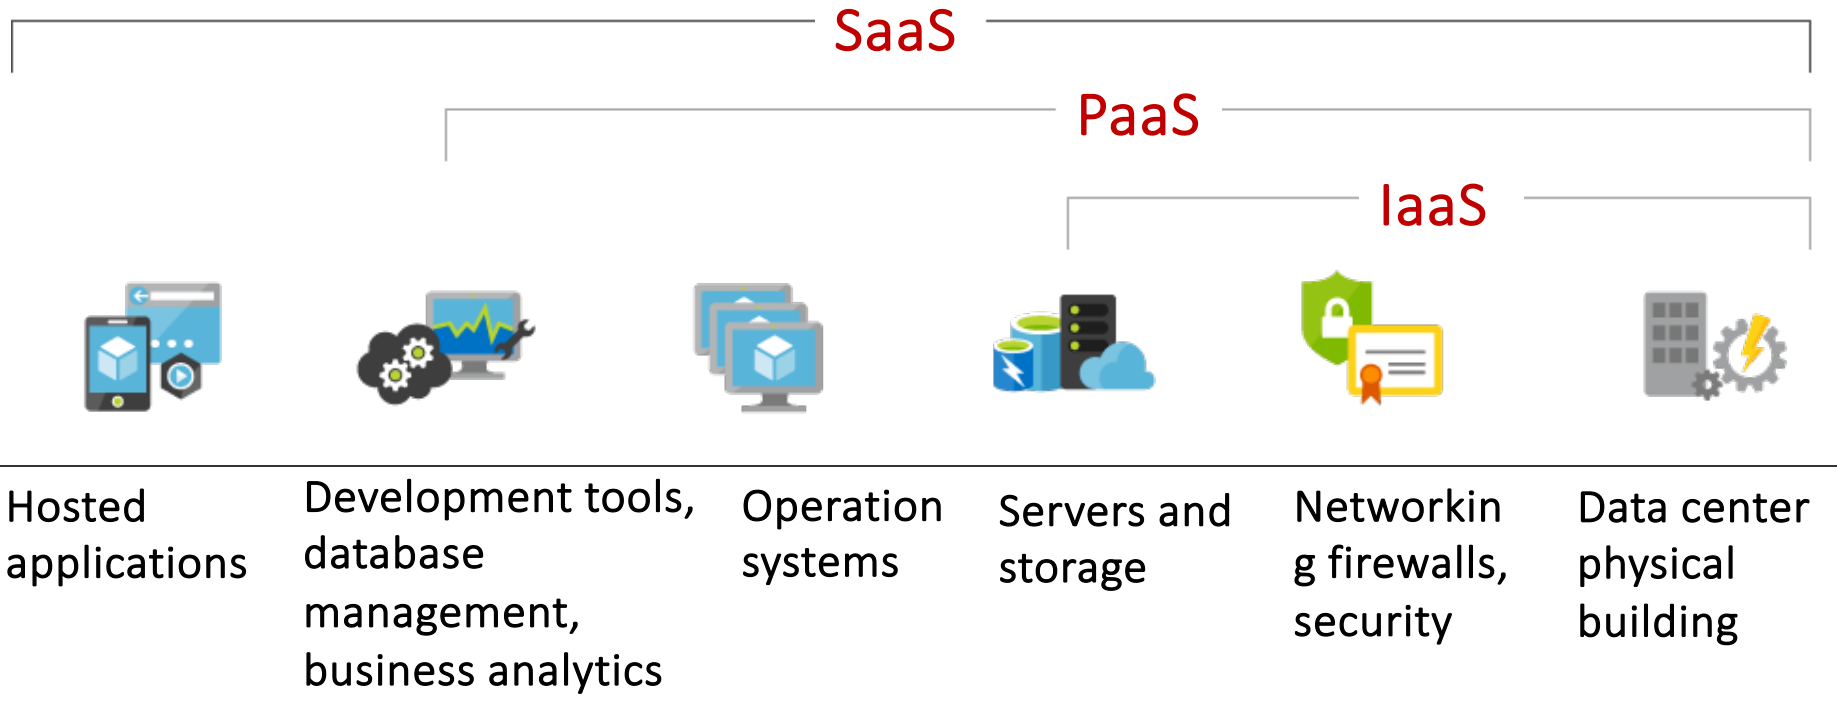


SaaS: Software as a Service

PaaS: Platform as a Service

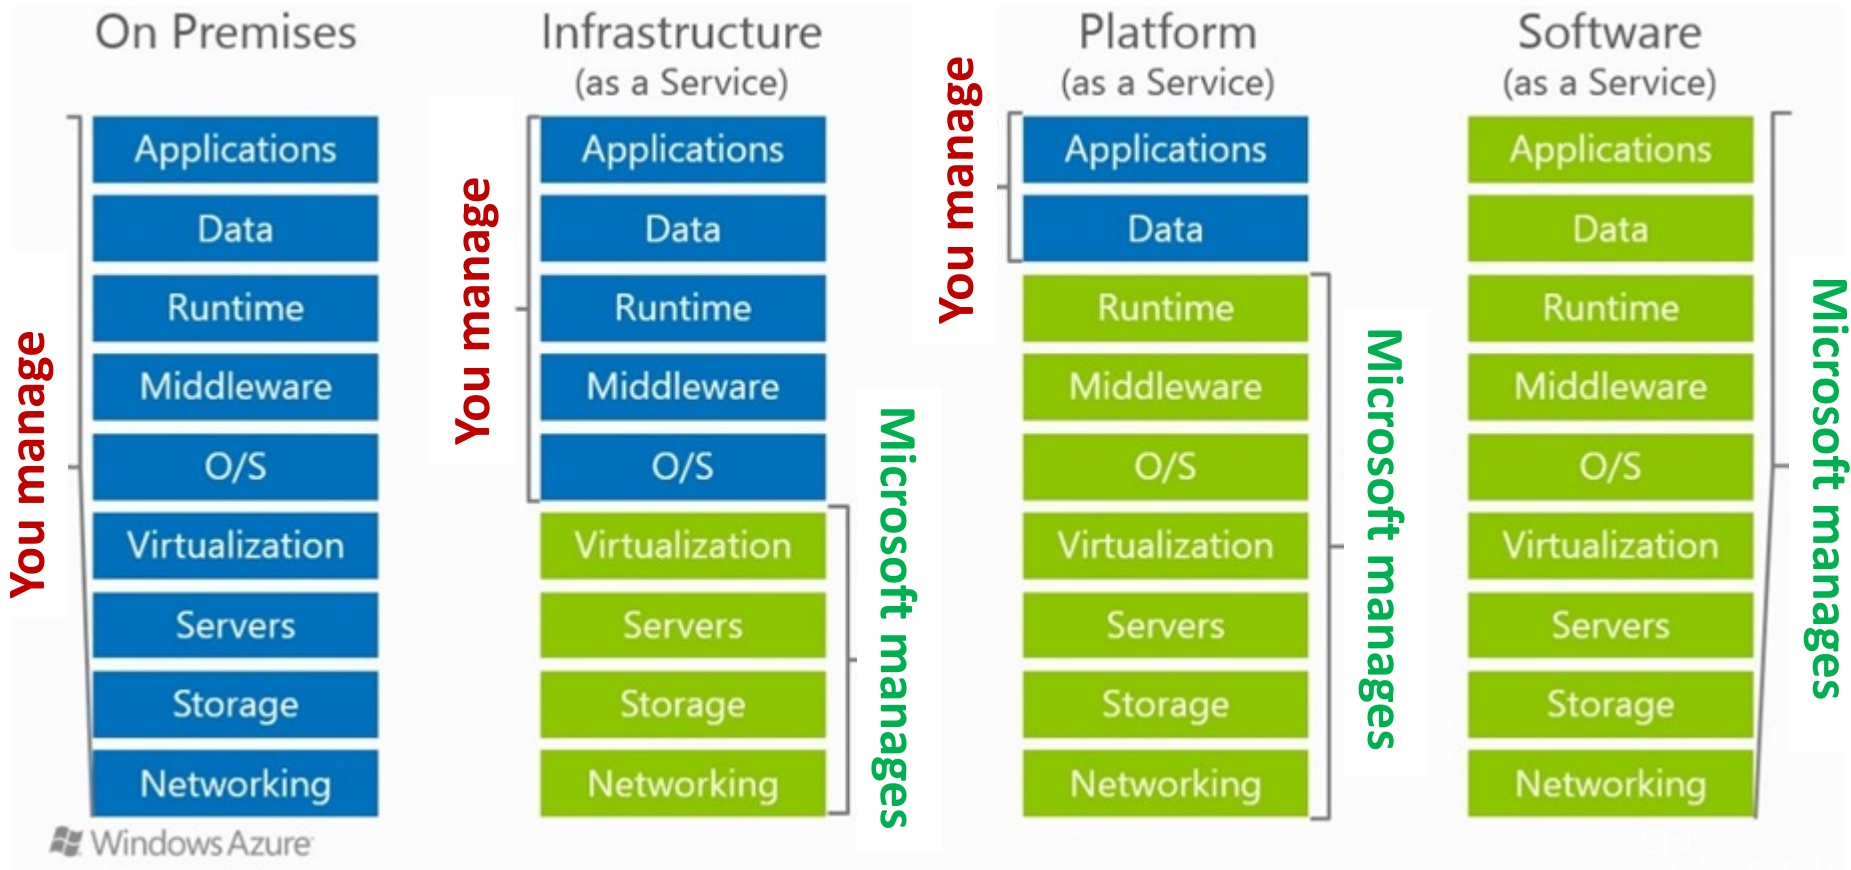
IaaS: Infrastructure as a Service

IaaS, PaaS, SaaS Provide Different Hardware/Software



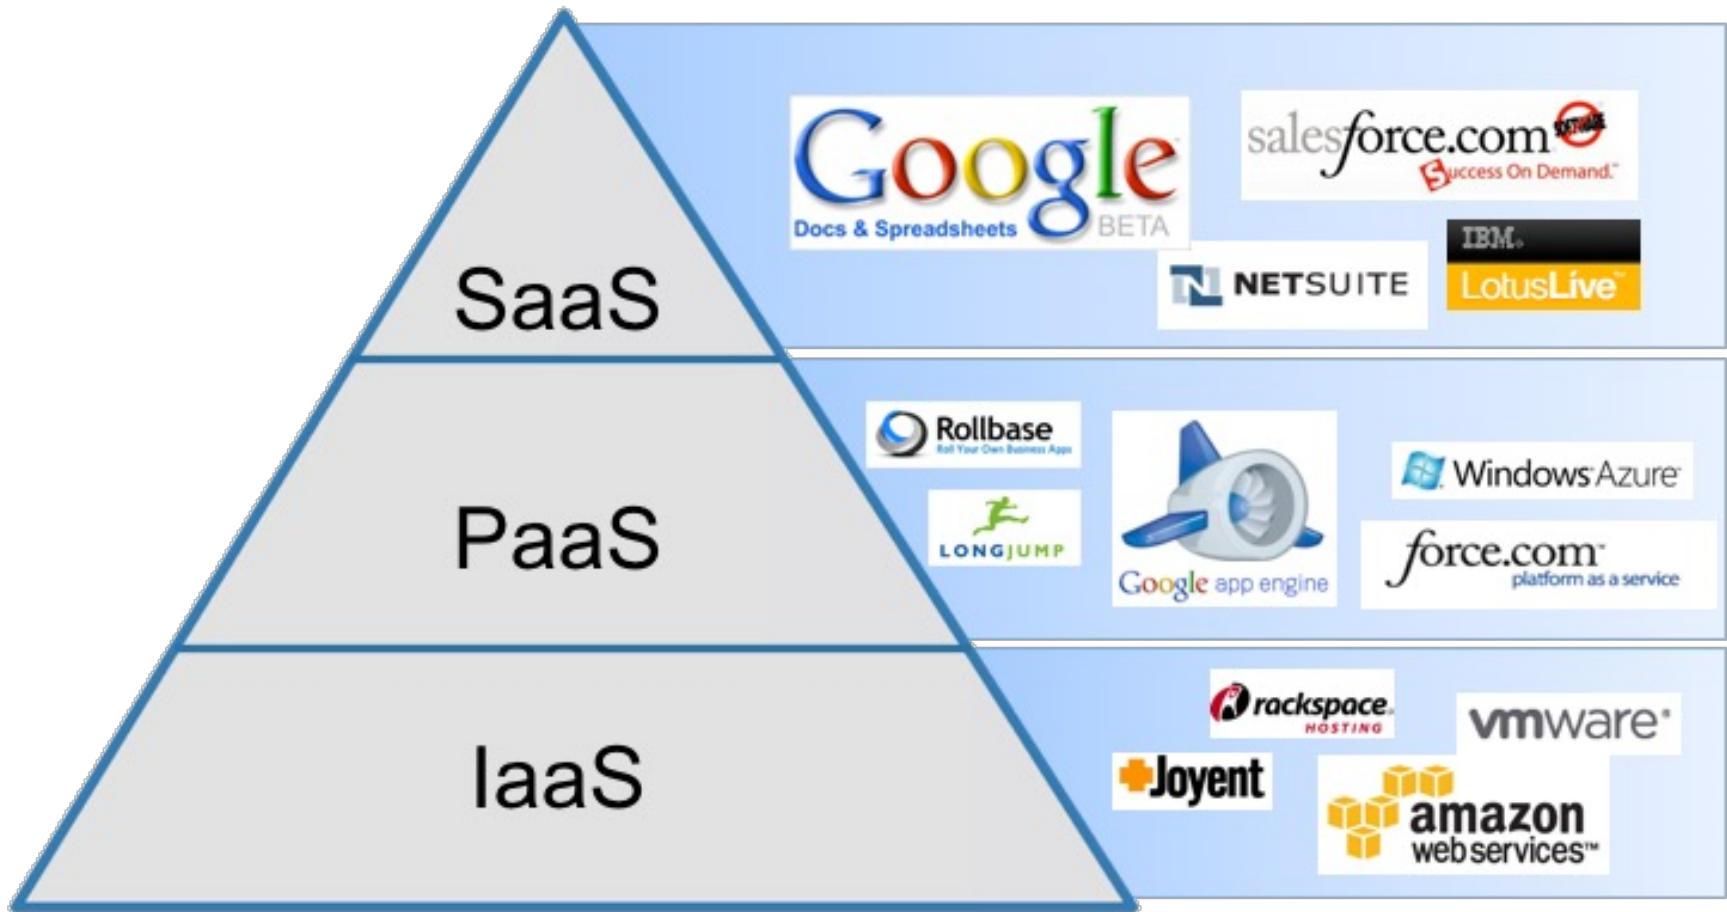
IaaS, PaaS, SaaS provide different level of hardware and software services

IaaS, PaaS, SaaS have Different Software Stack Control

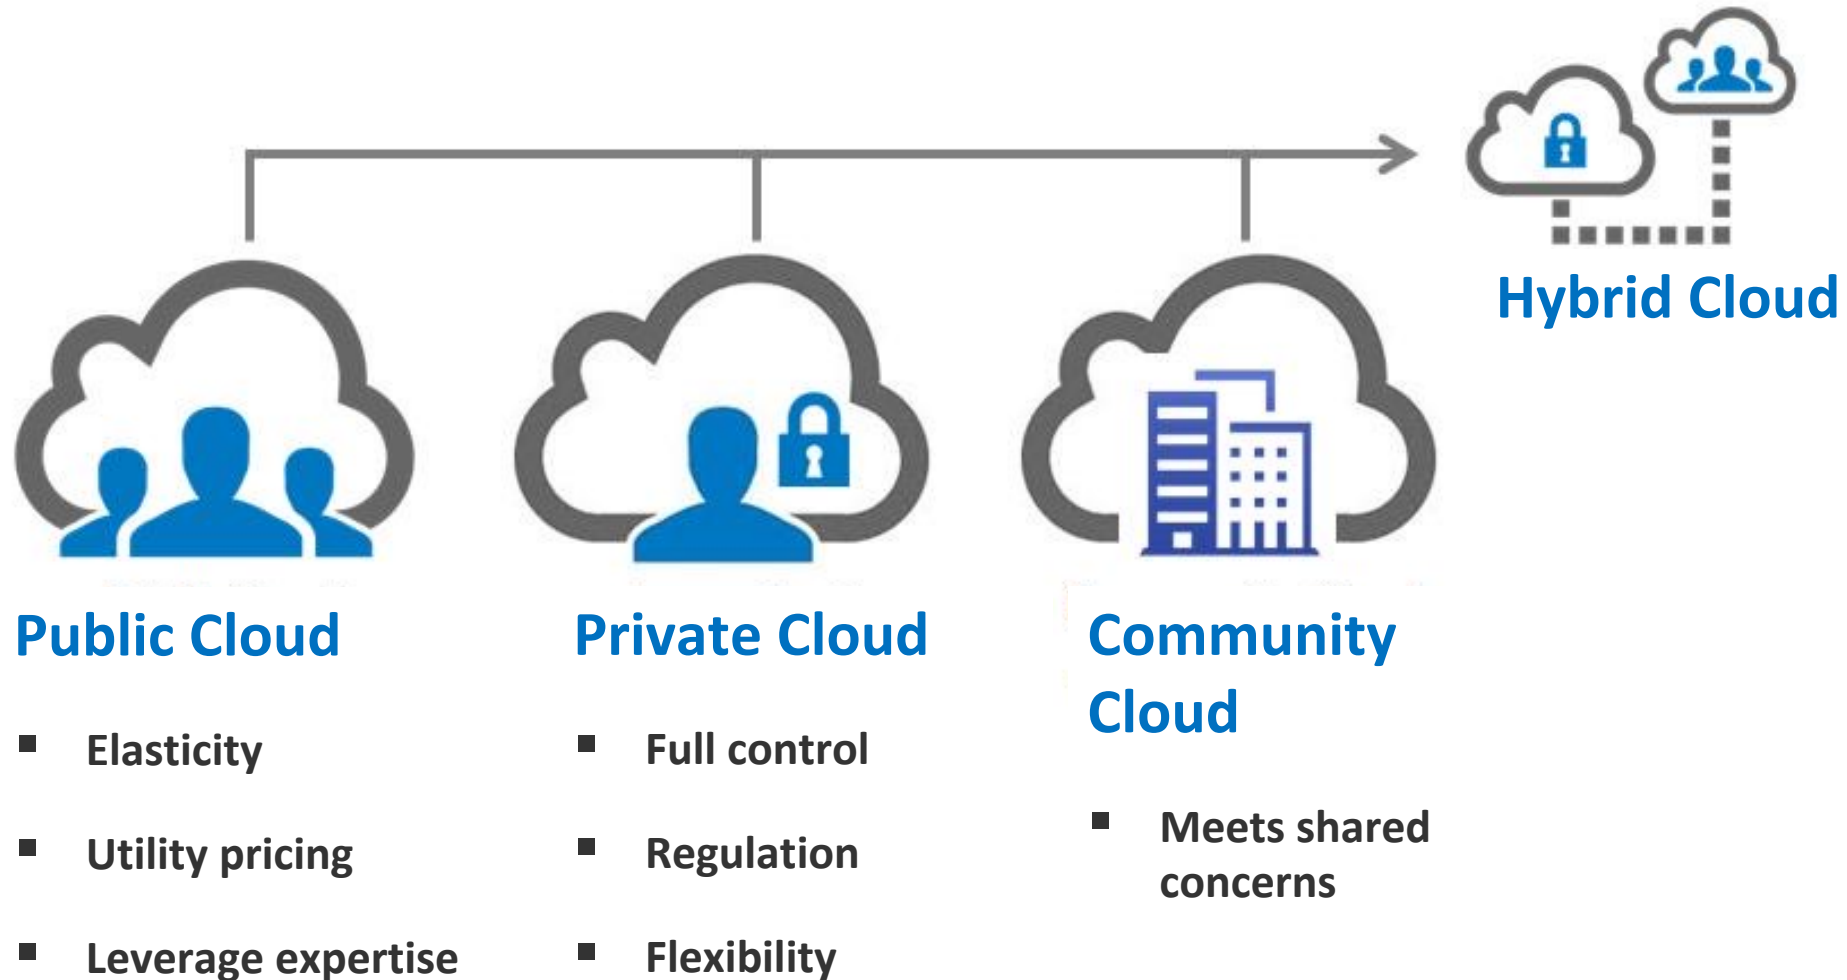


Microsoft Azure Cloud Services as an example to show IaaS, PaaS and SaaS software stack control, by Microsoft and by customers

Example Vendors



Four Deployment Models




Four Deployment Models – 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.

Example: public cloud service providers like Amazon Web Service (AWS)




Four Deployment Models – private cloud

The cloud infrastructure is provisioned for **exclusive use by a single organization** (multiple consumers), E.g., business units.

It may be owned, managed, and operated by an organization, a third party, or a combination of them.

Example: Amazon AWS GovCloud




Four Deployment Models – hybrid cloud

The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community or public).

They are bound together by standardized technology that enables data and application portability.

Example: An enterprise can deploy a **private cloud to host sensitive workloads**, but use a third-party public cloud provider, e.g., Google Compute Engine, to host less-critical resources



Four Deployment Models – 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 a combination of them.

U.S.-based dedicated IBM SoftLayer cloud for federal agencies



Cloud Computing for Transport in Norway




Background: Driving in slippery condition could be very dangerous. It will be easy to maneuver if you already know where to look out for on the road.

Volvo works with Norwegian Public Roads Administration to launch cloud-based system to share road friction information provided by individual cars

Approach: 50 test cars on the roads. The Volvo test car detects an icy road, the information is transmitted to Volvo's database via the network ICE and then transmitted to other vehicles that are approaching the slippery area

Cloud Computing for Healthcare - by Norwegian Health Network




Use case study:
Consuming cloud services from the Health Network

Nearly all health care providers in Norway are connected to the Norwegian Health Network, and approx. 150 third-party service providers.


Main focus: information security and control. Norwegian laws on public archives do not allow storage of archive material outside of Norway

Challenges for cloud computing...

New Observation: end-users are driving the big changes...



Data Consumer



Data Prosumer
(producer + consumer)

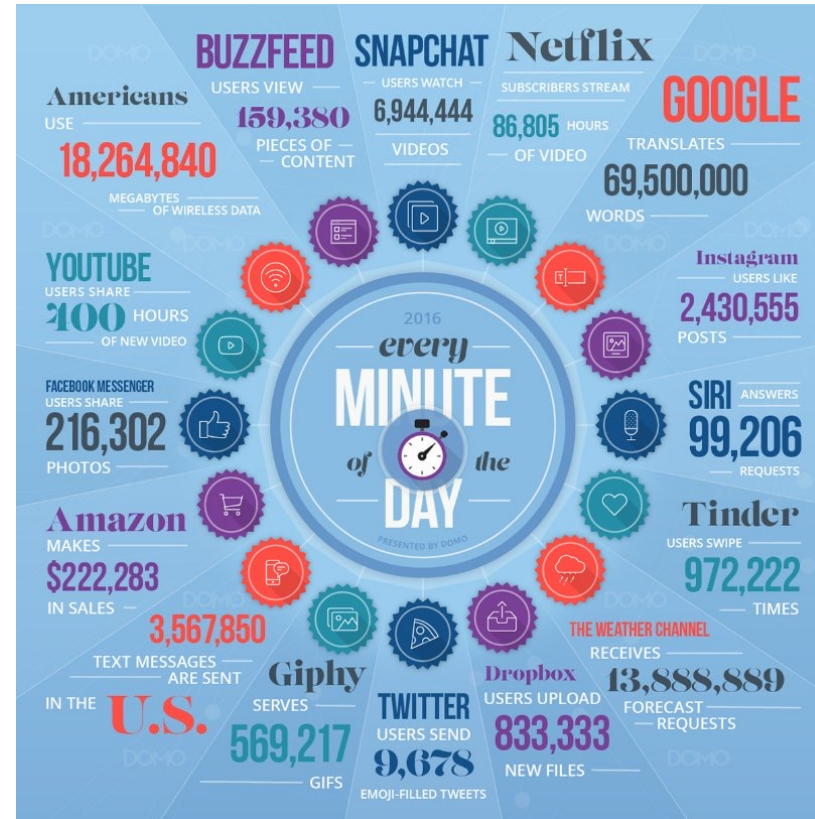
New Observation: end users are driving big changes

An end-user is changing from a pure data consumer to both data consumer and data producer. This change requires more new functions at the end-user side.

Data Consumer: an end-user watches a YouTube video on a smart phone.


Data Producer: people are producing data from their mobile devices.

What are new challenges for cloud computing when end-users' pattern are changing?



For example, **every single minute**, YouTube users upload **100 hours of new video** content; Instagram users post nearly **2430000 new photos**

1st New Challenge for Cloud Computing – Data Transmission




With the **growing data generation** at the end, **speed of data transmission is becoming the bottleneck** for the cloud-based computing paradigm

For example: **about 5 GB** data is generated by a Boeing 787 every second(*), but the bandwidth between the airplane and satellite or base station on the ground is not large enough for data transmission



* Boeing 787s to Create Half a Terabyte of Data Per Flight, Says Virgin Atlantic. Accessed on Dec. 7, 2016. [Online]. Available: <https://dataflog.com/read/self-driving-carscreate-2-petabytes-data-annually/172>

2nd New Challenge for Cloud Computing – Latency



Cloud may have high latency which does not fit certain applications

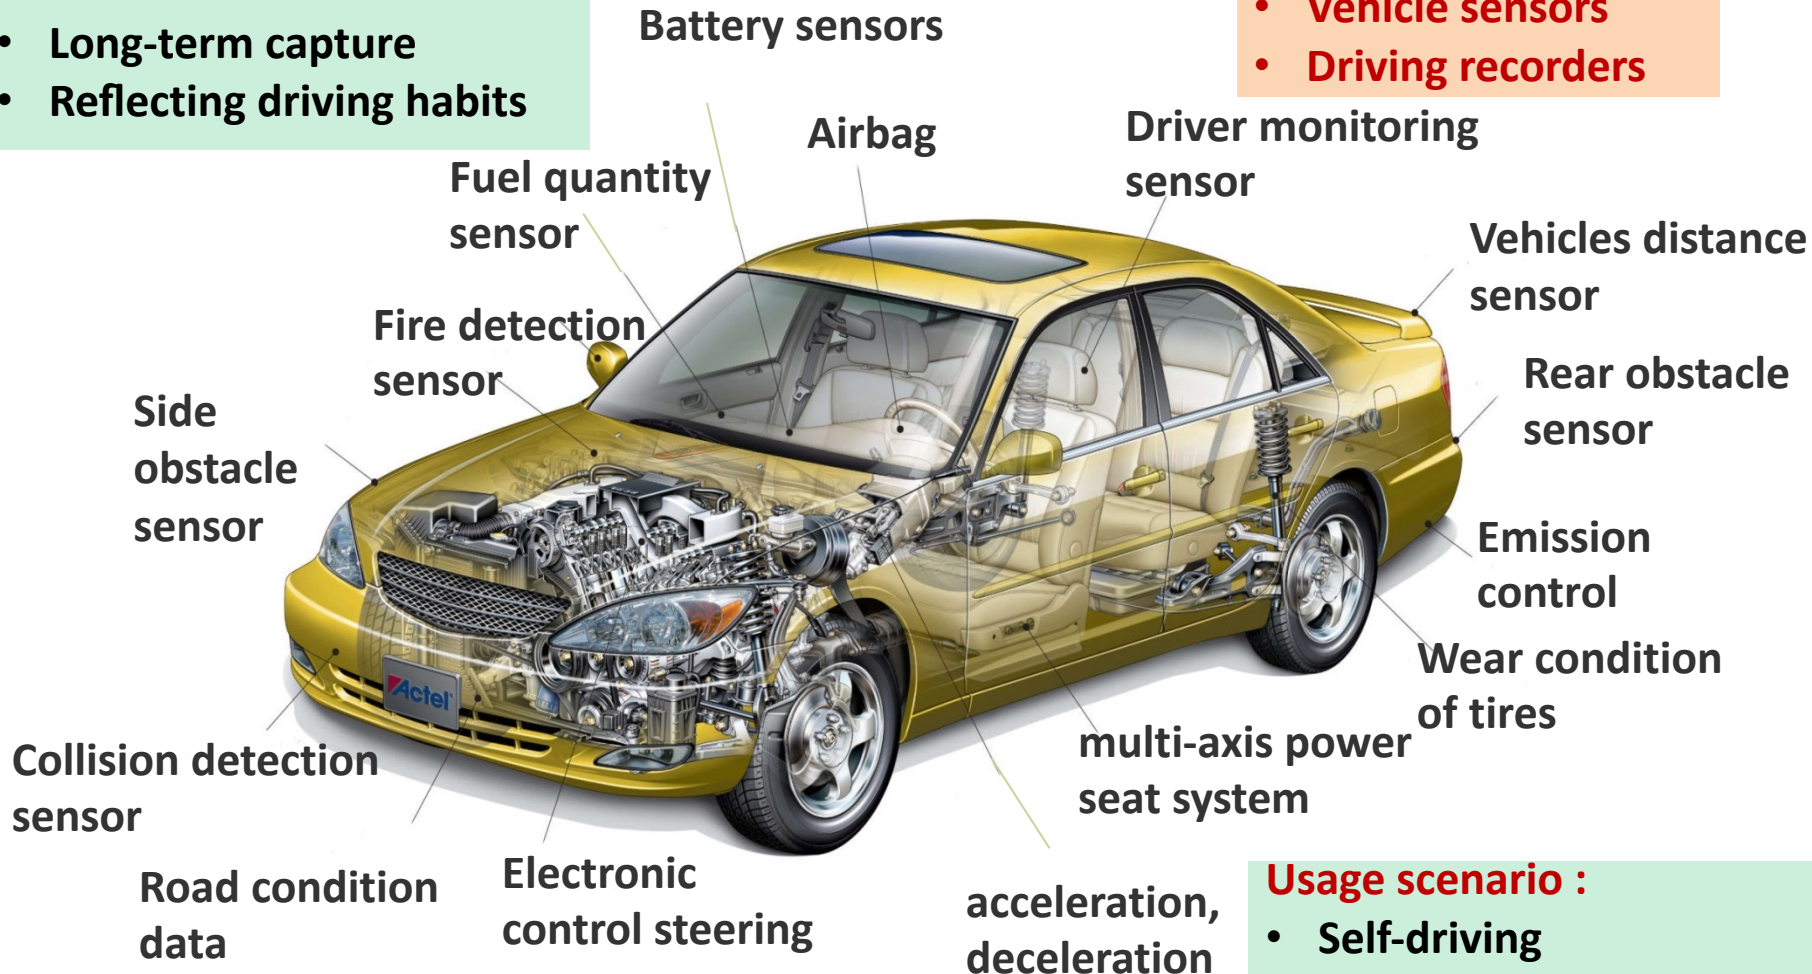
Vehicle sensors: vehicular big data

Feature:

- Periodical
- Long-term capture
- Reflecting driving habits

Source:

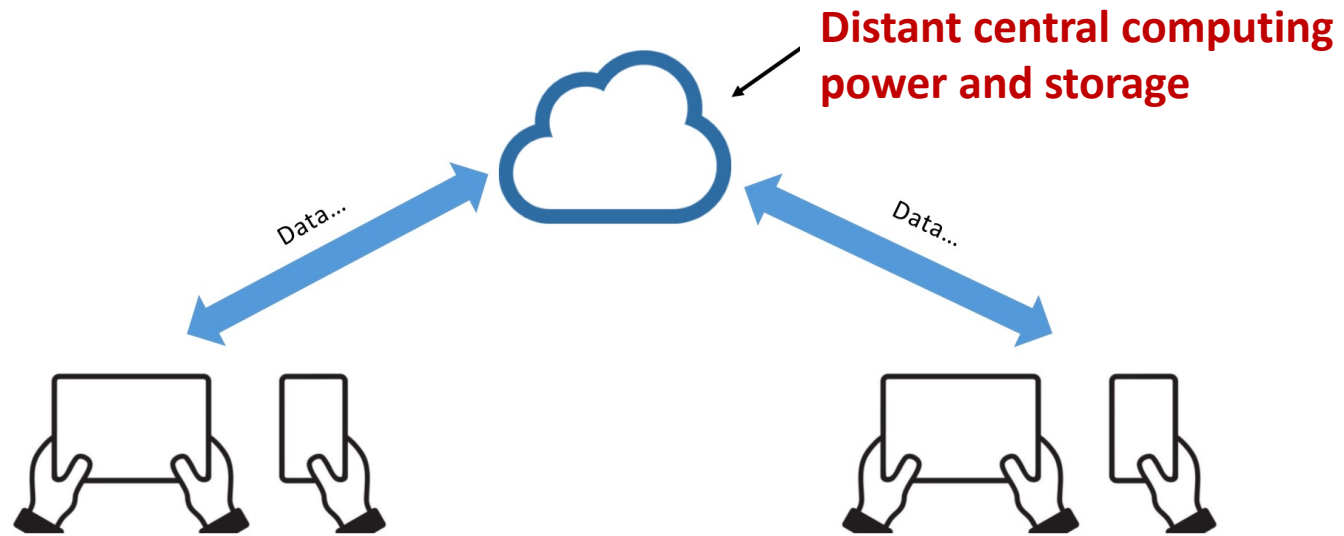
- Vehicle sensors
- Driving recorders



Usage scenario :

- Self-driving
- Collision avoidance

2nd New Challenge for Cloud Computing –Latency



Cloud may have high latency which does not fit certain applications

An autonomous vehicle generates **1 GB data every second** and it **requires real-time processing** to make correct decisions. If all the data needs to be sent to the cloud for processing, the response time would be too long.

8 Cameras + 12 Sensors
+ 1 Radar: 1GB/sec data



3rd New Challenge for Cloud Computing – Privacy

Storing data and important files on external service providers in cloud always opens up risks

For example: for wearable health devices, since the physical data collected by the things is usually private, uploading raw data to cloud has the risk of privacy issues

CHANGING FACE OF SECURITY

LinkedIn Lost 167 Million Account Credentials in Data Breach

Robert Hackett

May 18, 2016



Remember LinkedIn's [2012 data breach](#)?

A hacker [stole](#) 6.5 million encrypted passwords from the site and posted them to a Russian crime forum. Now it appears that data theft was just the tip of the



These New Challenges Motivate Fog Computing Paradigm



Data transmission



Low latency



Privacy & security

Fog Computing

New challenges motivate to put computation, storages, **services close to the end-users** to significantly reduce latency and protect private data.

Expectation: 45% of data will be stored, processed, analyzed, and acted upon close to, or **at the edge of the network**

FOG COMPUTING

Fog Computing definition

According to CISCO

Fog Computing extends the cloud computing paradigm that provides computation, storage, and networking services between end devices and traditional cloud servers. Fog computing nodes are typically located at the edge of network located away from the main cloud data centers.

Edge computing and fog computing



We define “edge” as any computing and network resources between data sources and cloud data centers.

Fog Computing Architecture

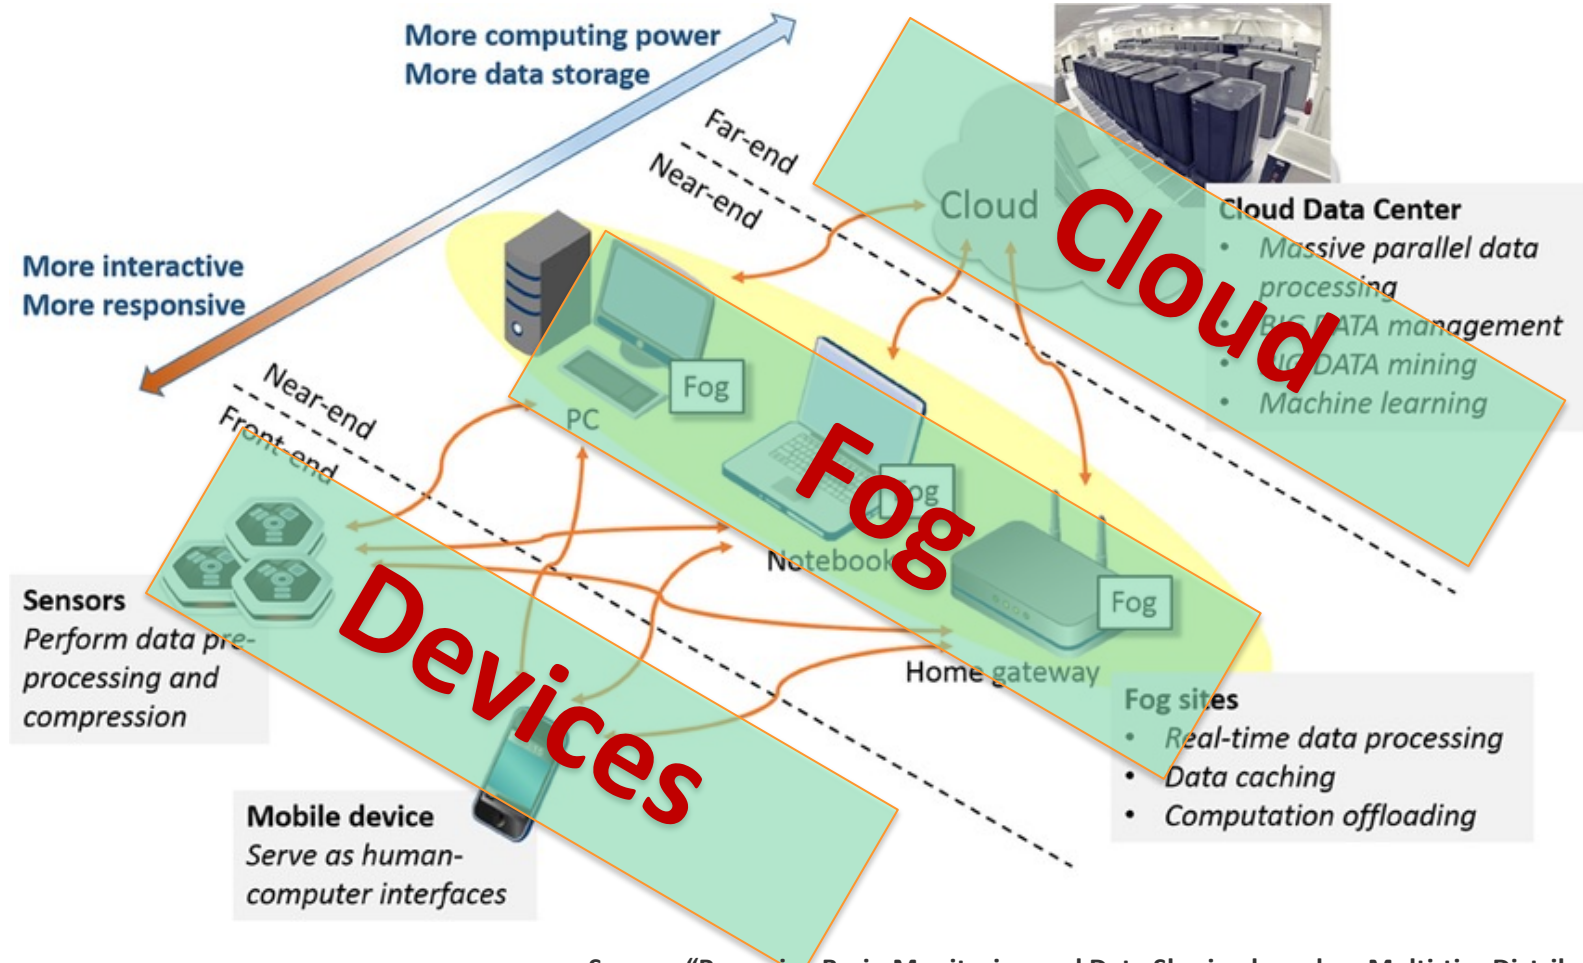
Fog computing refers to a three-tier architecture

Centralized cloud servers **coexist** with fog nodes **but are not essential** for the execution of fog services

Fog nodes can range from resource-poor devices (e.g. end devices) to **more powerful cloud servers** (e.g. Internet routers, 5G base stations).



Where is Fog or Edge? – an example




Source: "Pervasive Brain Monitoring and Data Sharing based on Multi-tier Distributed Computing and Linked Data Technology", Frontiers in Human Neuroscience, June 2014

Three-layer architecture: End-devices → Fog layer → Cloud layer.
We also call this Fog-to-Cloud (F2C) architecture

Main Difference between Cloud Computing and Fog Computing

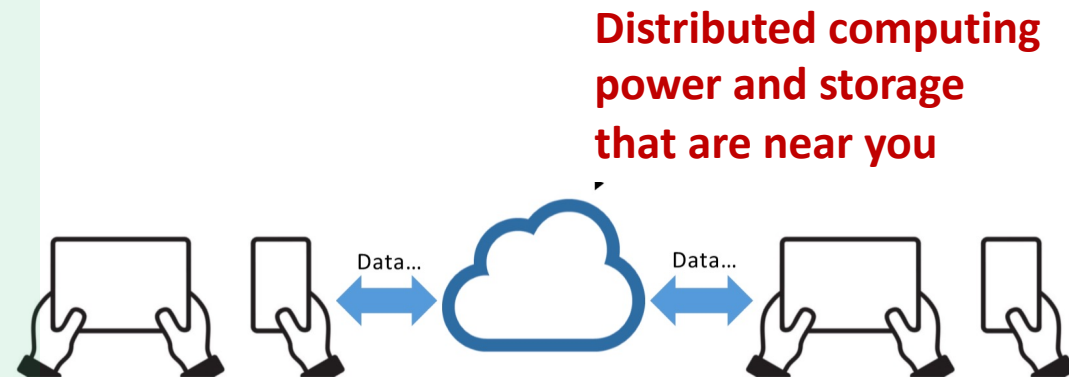
Cloud computing

- Servers, computing power, storage are **centralized**.
- The **specific location** of resources are **not known to end users**.
- They have **no or limited control of the cloud infrastructure**.

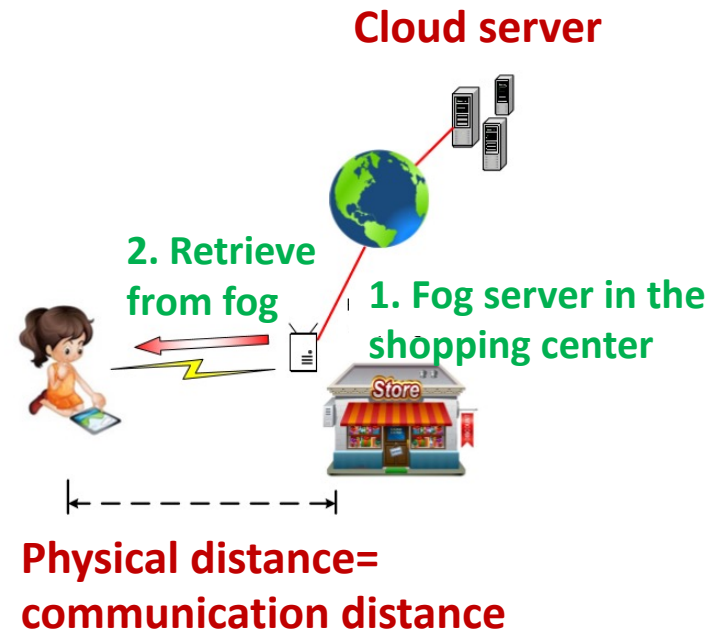
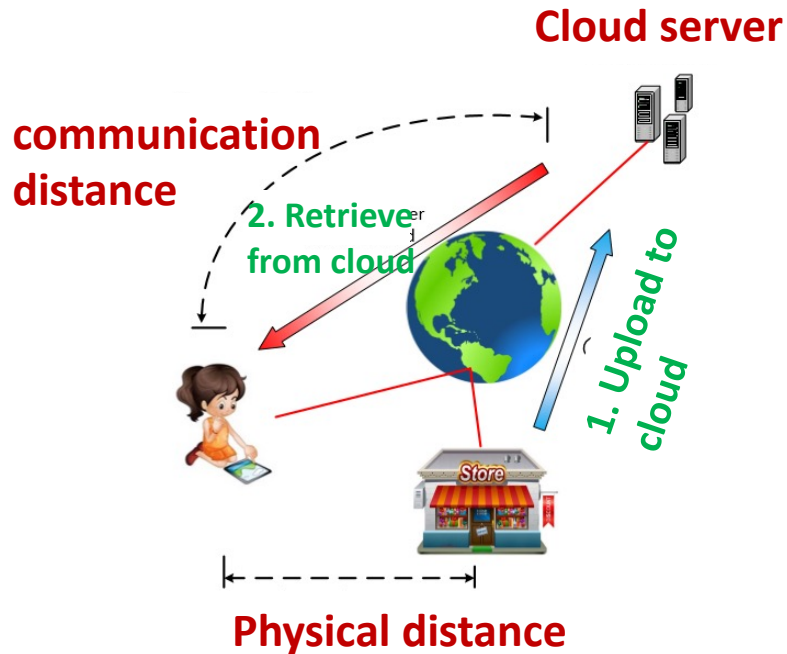


Fog computing

- Servers, computing power and storage are close to end users.
- The **specific location** of the resources **may be known**
- End users may be able to **manage partially**.



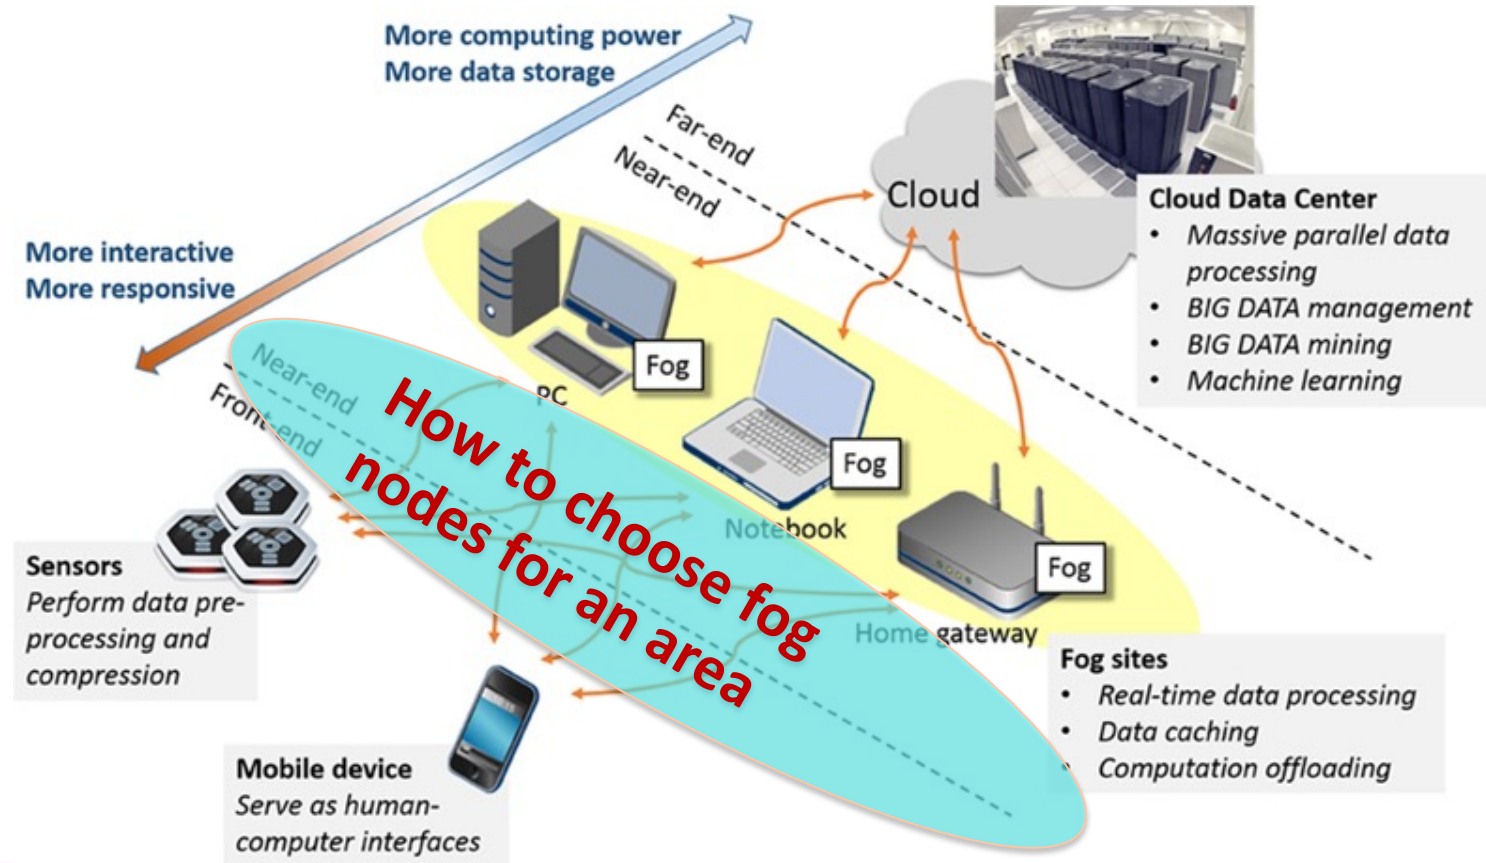
Fog Computing Advantage – An example



Face recognition (*): A user uses an app in a smart phone to capture a face photo of herself or any other people, and the app will transmit the face photo to a server, either in a fog or in Amazon EC2 cloud. The server will recognize the face by matching it in the local face photo database.

Response time: reduced **from 900 to 169 ms** by moving computation from cloud to the edge. The time duration from when the smartphone begins to upload the face photo to when the smartphone receives the result.

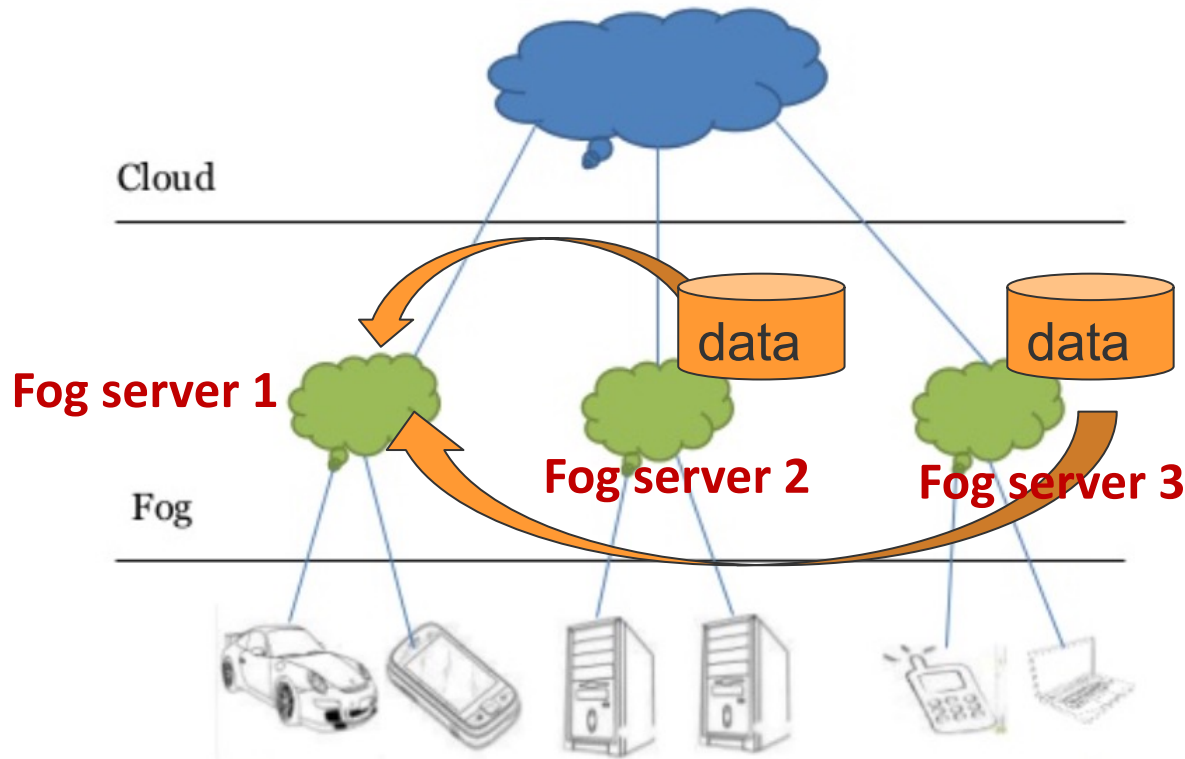
Fog Computing Challenges – Quality-of-Service (I)



Connectivity

- Fog nodes provide new opportunities to reduce cost and expanding connectivity.
- **The selection of fog node from end users will heavily impact the performance.**
- We can dynamically select a subset of fog nodes to increase the availability of fog services for a certain area.

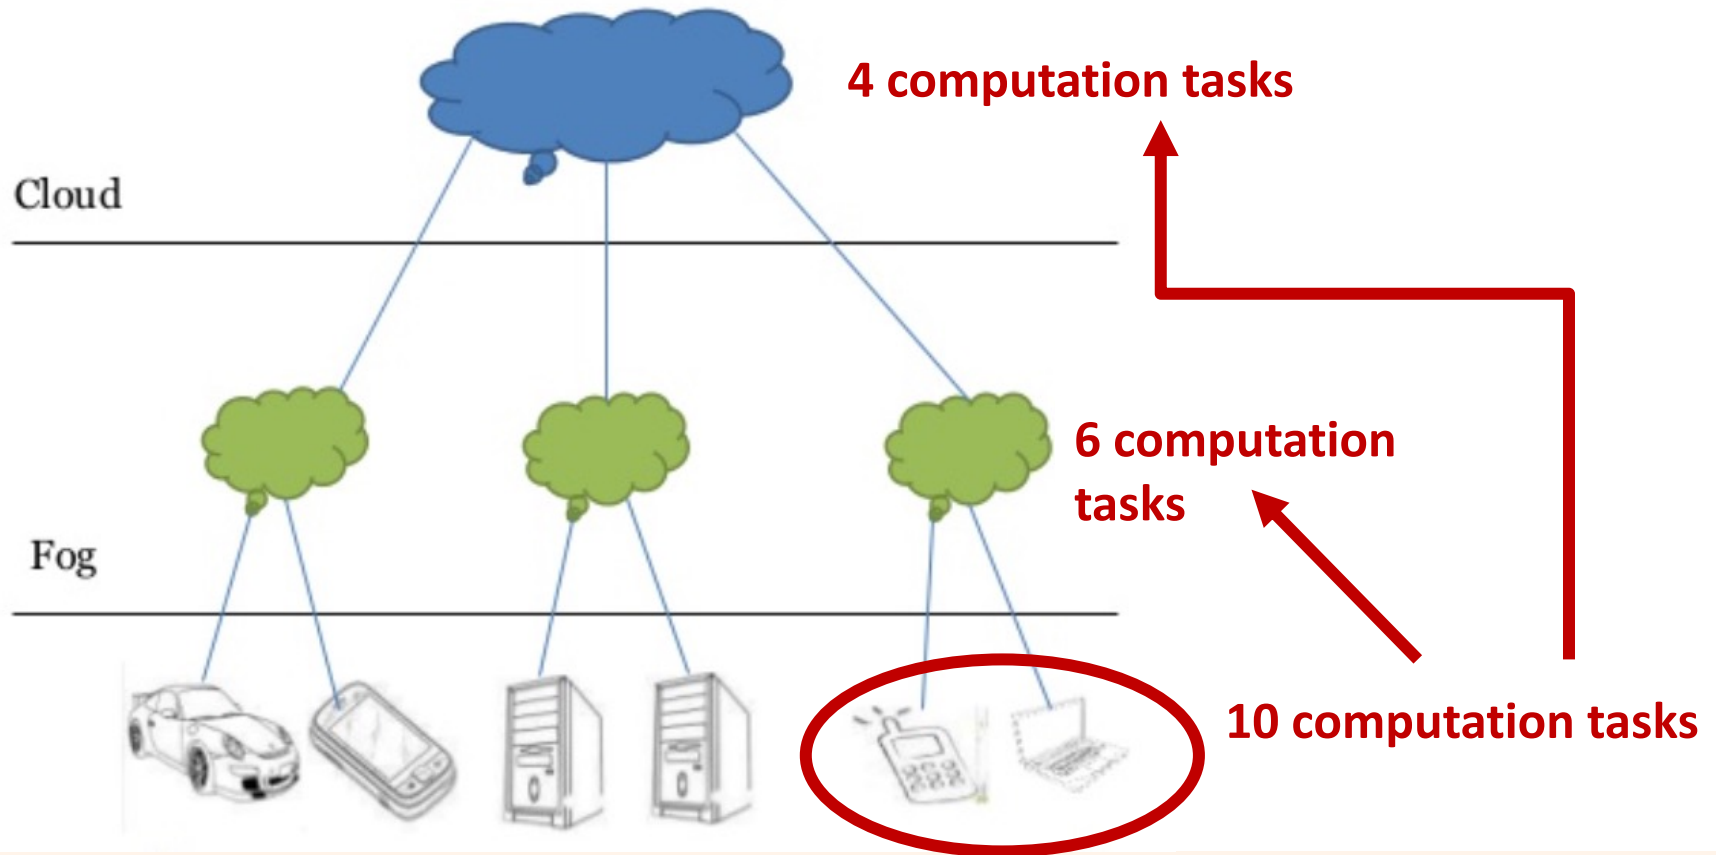
Fog Computing Challenges – Quality-of-Service (II)



Capacity: It is important to investigate how data is placed in fog network since data locality for computation is very important. Need to consider two aspects: 1) **network bandwidth**, 2) **storage capacity**.

For example, a fog node may need to compute on data that is distributed in several nearby nodes. The computation cannot start before the completion of data aggregation, which adds delay to services.

Computation Offloading



Computation offloading: Heavy computation tasks can be executed in fog or cloud instead of mobile devices. This saves storage and battery lifetime.

- The main challenge is how to deal with the **dynamics**.
- In three-layered architecture: device-fog-cloud, a typical question: **how to partition tasks to offload on fog and cloud**.

Fog Computing Applications (I)

Fog/Edge computing is able to deal with new applications that suffer from the limitations and poor scalability of the centralized cloud paradigm.


(Typical) Applications susceptible to latency

- **manufacturing**
- **urban transport**
- **smart energy systems**

where actuation has to be driven in real-time.



Geo-distributed applications: Typical examples include **environmental monitoring**, which are based on the collection and processing of streams from thousands or even millions of sensors.



Fog Computing Applications (II)

Mobile applications: typical applications involving **fast mobility**, e.g., autonomous driving. They require moving objects to access local resources (computing, storage) residing at their vicinity.



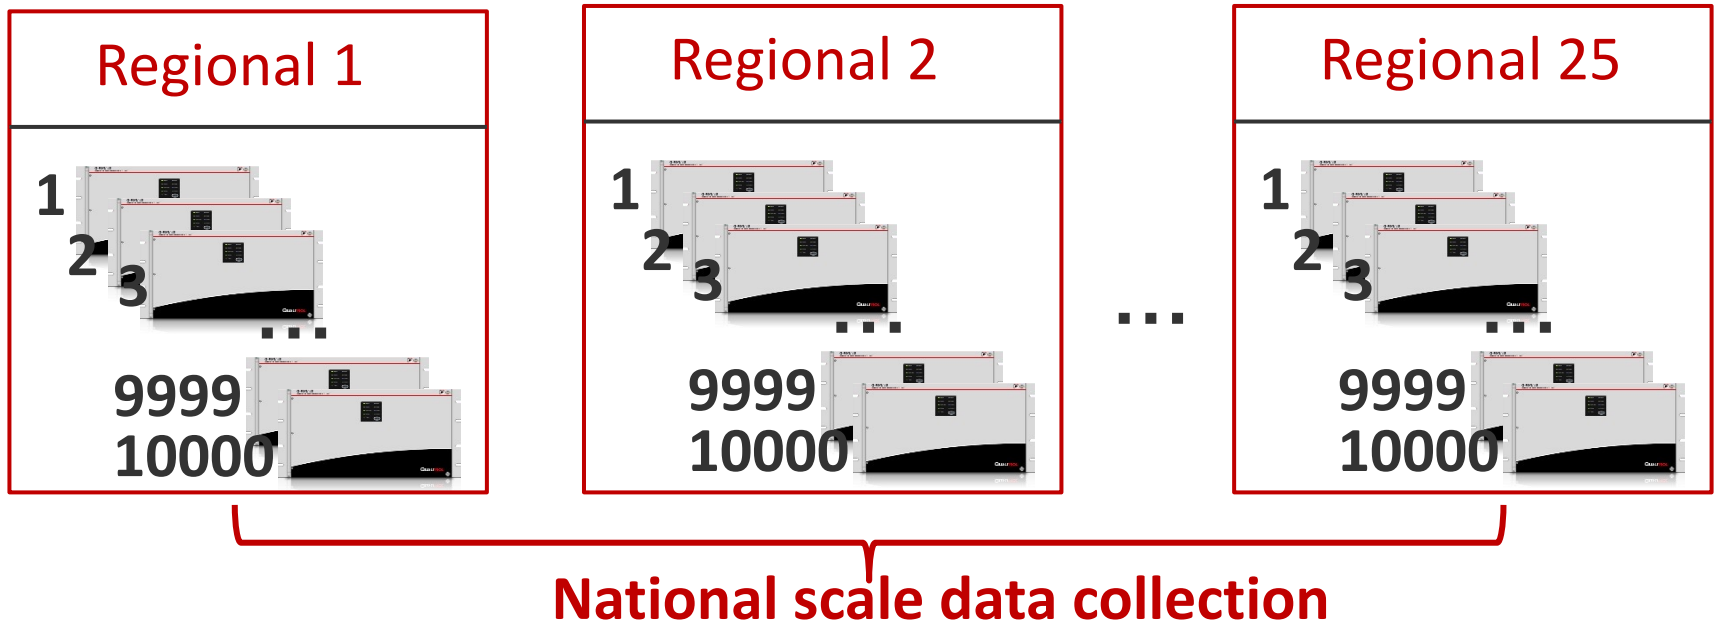
- **Distributed multi-user applications with privacy implications and need for fine-grained privacy control.**
- These applications can benefit from a **decentralization of the storage and management of private data** to various edge servers, thus alleviating the risk of transferring, aggregating and processing all private datasets at the centralized cloud.



CLOUD/FOG COMPUTING FOR SMART GRID

National Scale Phasor Data Collection – an example

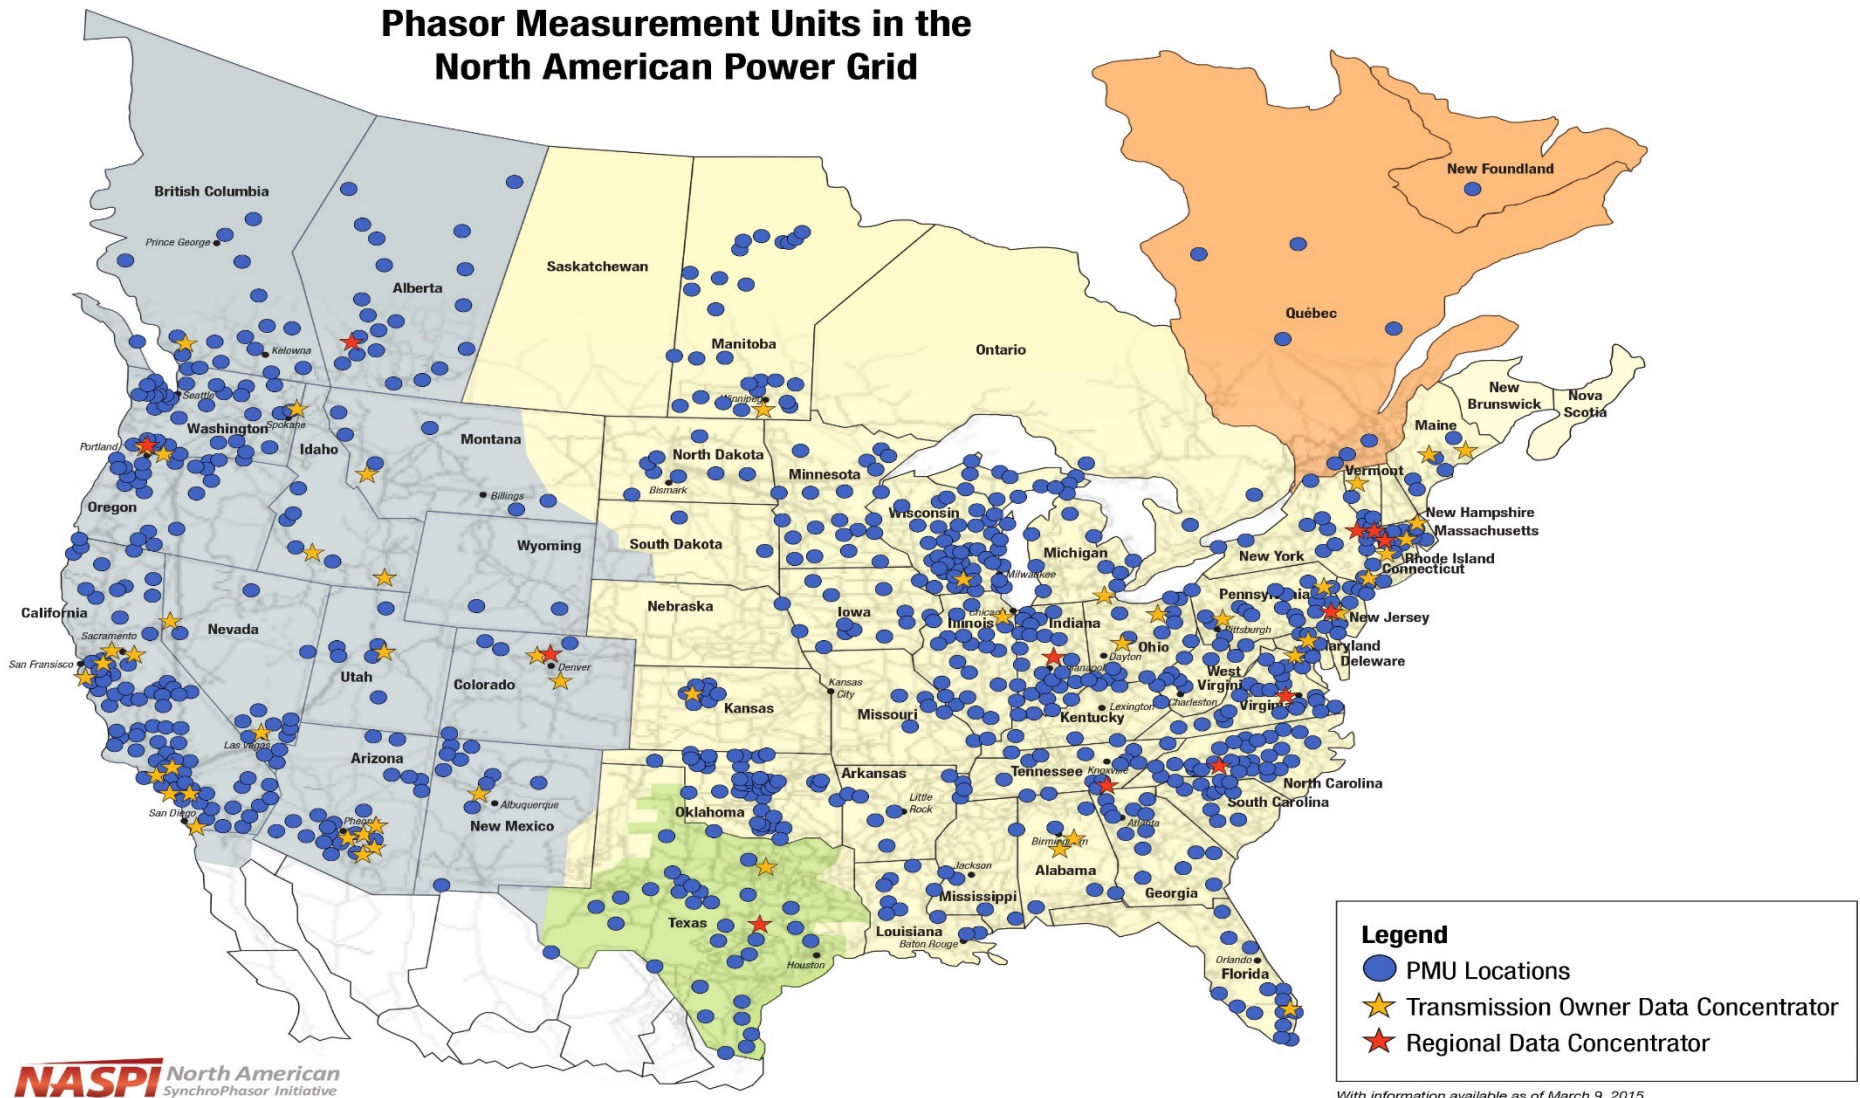
Phasor Measurement Unit (PMU) is a device which measures the electrical waves on an electricity grid. The PMU data is analyzed for detection of power system anomalies, hence is crucial for power grid stability.



A country has **25 regions**. Each region has **10000 PMU**. Each PMU has **30 measurements/sec** and each measurement has **256 bytes**.

Generated data: $256 * 30 * 10000 * 25 = 15 \text{Gbit/sec!}$

NASPI PMU Map (as of Mar 9, 2015)



Reliable synchrophasor data communication and analysis is essential for the effective monitoring and control of the power system

A blackout in 2003 caused an economic loss of **10 billion U.S dollars**


Subsequent investigation of the blackout revealed that the catastrophe **could have been prevented if there was an early warning system**

Tremendous Smart Meters Data in Smart Grid

The amount of data in smart grid may be tremendous since: i) there are a huge number of smart meters/sensors; ii) each meter shall transmit data periodically to the control center.

22 million smart meters in Zhejiang Province, China, as required by China State Grid.

A smart meter should be **96 measurements/day**, which leads to **2.1 billion records per day**.

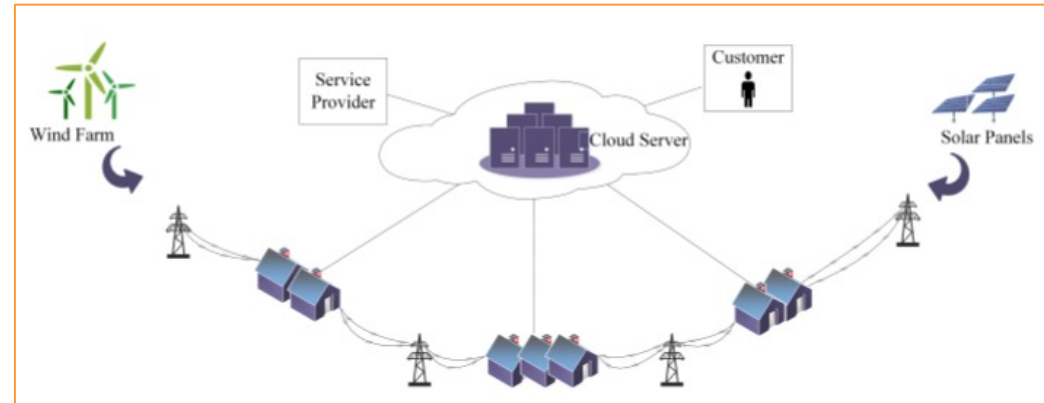


Zhejiang province

Cloud/Fog Architecture for Smart Grid

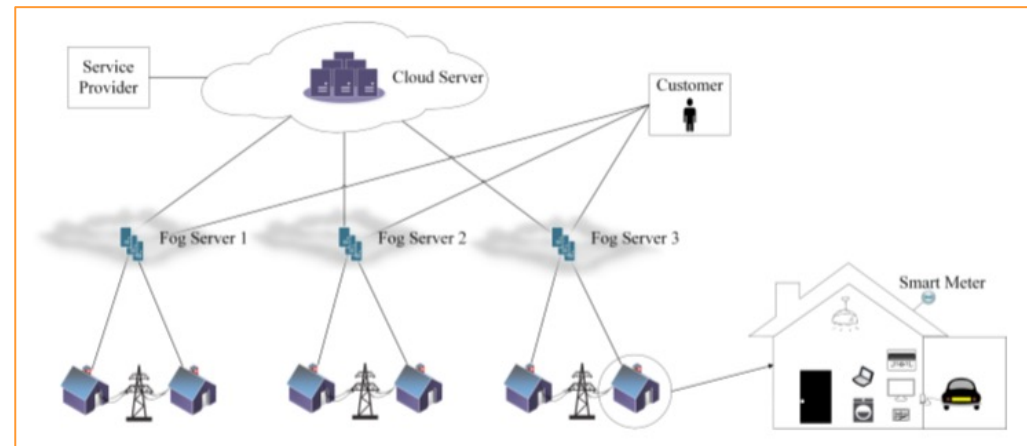
Cloud computing

All smart meter data is stored and processed in the cloud.



Cloud/Fog Architecture

Can support low latency requirement for communication and computation and improved privacy, which can be addressed by fog computing.




Central Cloud Computing for Data Management

Meter Data Management Systems (MDMS) uses cloud computing to perform **long-term data storage and management for the vast quantities of usage data and events**

MDMS enables interactions with operation and management systems that

- manages the billing and customer information,
- provides power quality report and load forecasting based on meter data

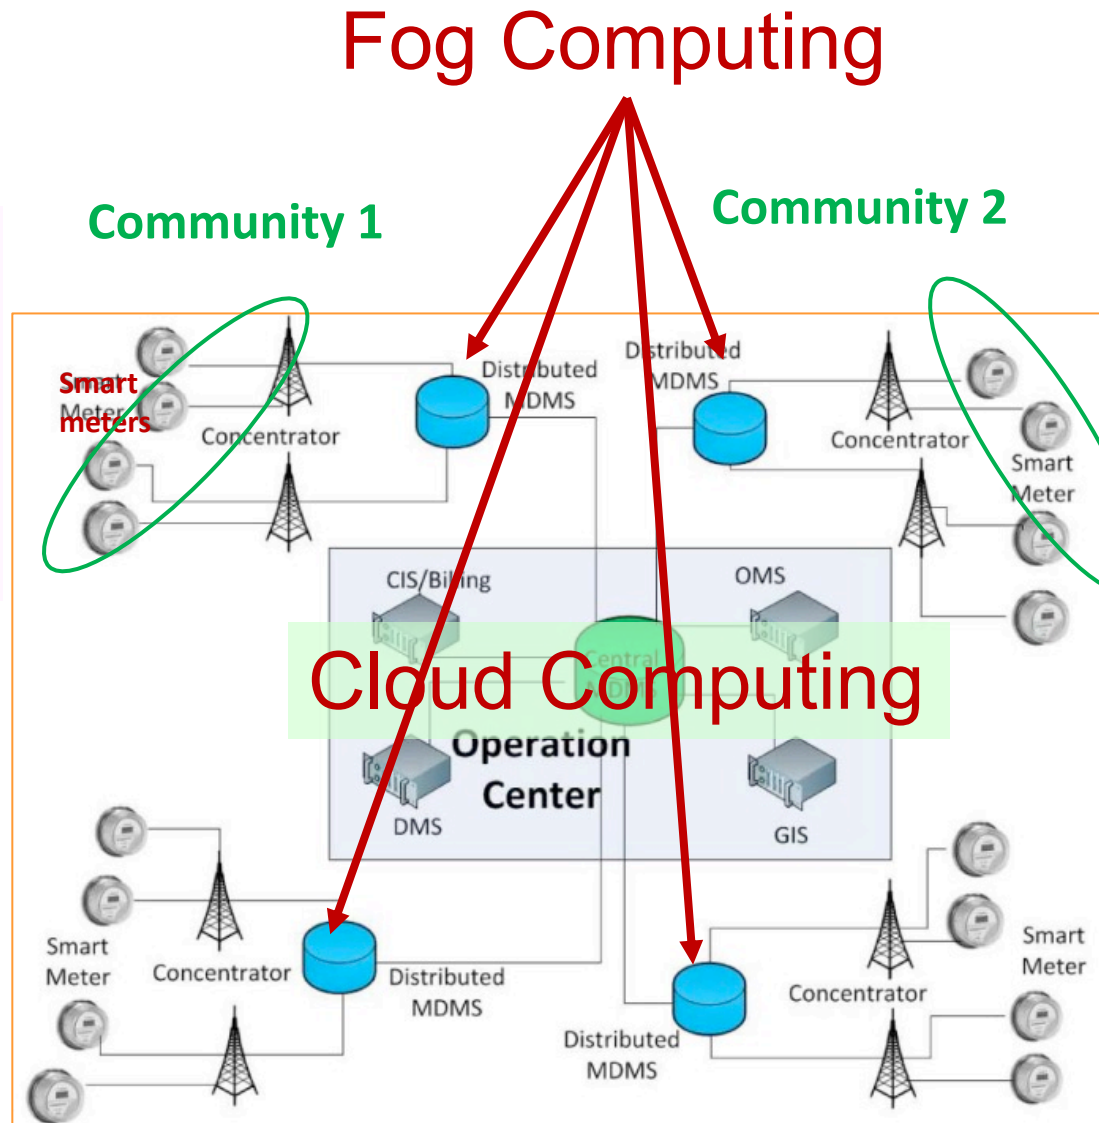


All data from the smart meters need to go through the centralized cloud, it makes the system non-scalable. Why?

One approach: Cloud/fog computing for data management

- Several distributed MDMSs using fog computing are close to smart meters.
- Each MDMS is responsible for the specific area.

Communication distance for data collection is largely reduced, and the corresponding resources needed are also less.




Fog Computing for Smart Home

In smart house, the home of the future might be equipped with a variety of **power use meters** and **monitoring devices**, adapting behavior to match cost of power, load on the grid, and activities of the residents.

A fog node/server at home can collect, process, analyze, aggregate and transmit a large amount of data.

The fog server can also manage the smart house for **decision making**, e.g., demand response, temperature control.



M. A. Al Faruque and K. Vatanparvar. *Energy Management-as-a-Service Over Fog Computing Platform*. IEEE Internet of Things Journal, 3(2):161–169, April 2016

Fog Computing for Wind Farm

Fog computing at turbines



- A large wind farm may consist of hundreds of wind turbines and cover an area of hundreds of square miles.
 - Modern wind turbines **are very large control structures aimed at improving wind power capture and power quality.**
-
- Wide geographical deployment of a large system consists of a number of autonomous yet coordinated turbines.
 - A turbine is supposed to quickly respond to external weather and environment, which gives rise to the need of a fog computing platform.

Fog Computing for Wind Farm - example



ParStream software is used by Envision, a power generation company that manages a 13 GW fleet of wind and solar panels.

The real-time output of wind turbines increased by 15 percent by analyzing data at the edge of a network with response times of less than a second

MORE CONSIDERATIONS...

Comparison of Fog Computing and Cloud Computing

	Cloud Computing	Fog Computing
Target Users	General Internet users.	Mobile users
Service Type	Global information collected from worldwide	Limited localized information services related to specific deployment locations
Hardware	Ample and scalable storage space and compute power	Limited storage, compute power and wireless interface
Distance to Users	Faraway from users and communicate through IP networks	In the physical proximity and communicate through single-hop wireless connection
Working Environment	Warehouse-size building with air conditioning systems	Outdoor (streets, parklands, etc.) or indoor (restaurants, shopping malls, etc.)
Deployment	Centralized and maintained by Amazon, Google, etc.	Centralized or distributed in regional areas by local business (local telecommunication vendor, shopping mall retailer, etc.)

Fog Computing in Industry (I)



The OpenFog Reference Architecture, to be released in 1Q 2017, is based on eight pillars.

IT industry leaders (e.g., ARM, Cisco, Dell, Intel, Microsoft, and the Princeton University Edge Laboratory) formed a coalition and have joined forces to create the OpenFog Consortium

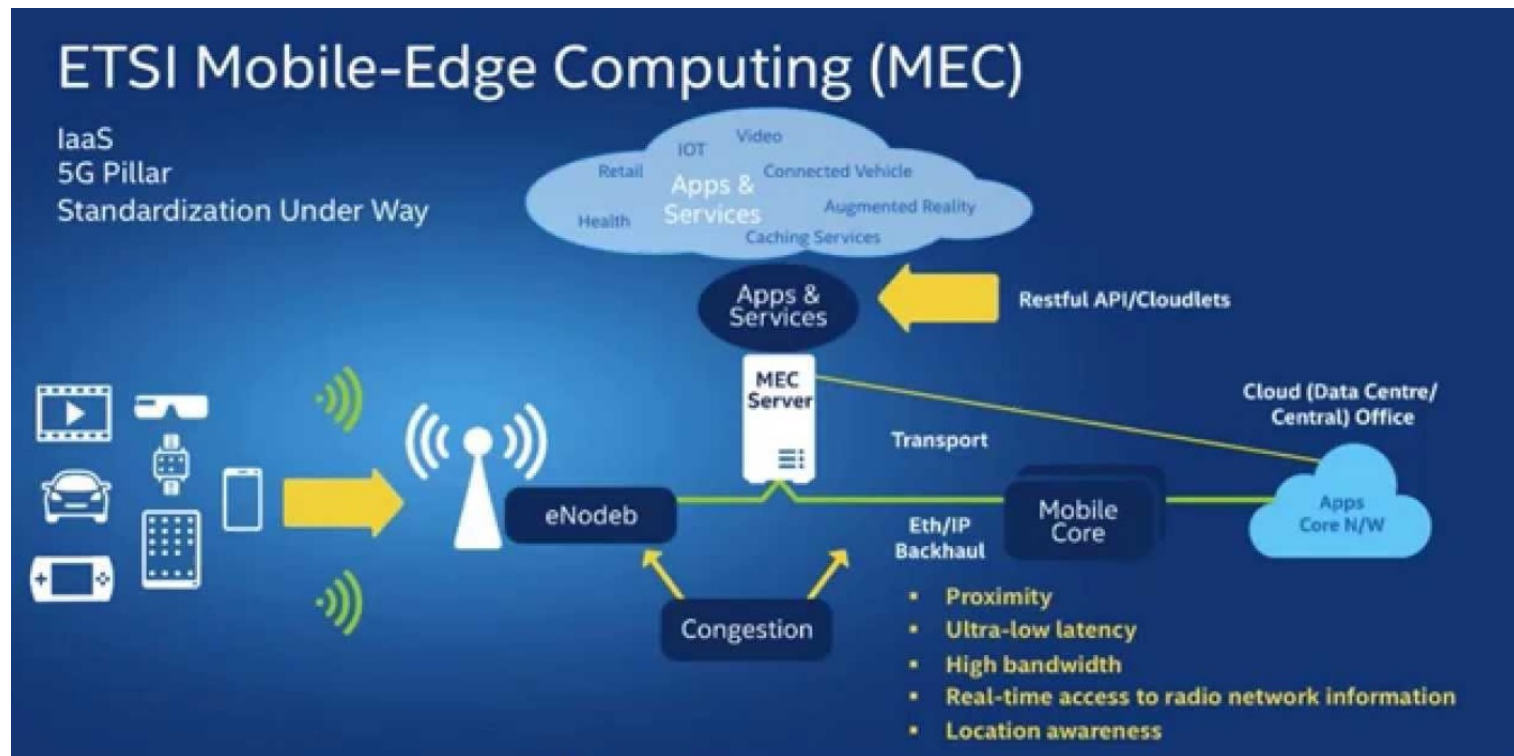
Goal: acceleration of Fog computing technology deployment for the good of the industry <https://www.openfogconsortium.org/>

Fog Computing in Industry (II)

Mobile Edge Computing (MEC) is developed by ETSI Specification. MEC pushes the cloud computing capabilities close to the Radio Access Networks in 4G. MEC server can be a mini-data center

located in either the base station


storage close to the radio access networks



Self-driving Car: use cloud computing or fog computing?

A self-driving car is a complex robotic system equipped with sensors, actuators and ICT capabilities.

Requirement: approximately 1 GB of data will need to be processed each second. This data will need to be analyzed quickly enough so that the vehicle can react to changes in its surroundings in less than a second.



Shall we use cloud computing or fog computing for self-driving car?

The requirements mean a self-driving vehicle would require a lot of computing power to minimize application latencies, as well as very low network latencies.

Cloud Computing Future: Anything as a Service (XaaS)

Recent advances: Drones as a Service, Robot as a Service, Blockchain as a Service

Drones as a Service: For example: Amazon delivers packages to your home with their drone delivery system



Drones are used as data mule to deliver data from wind farms to edge server or control center



Too far to send massive data directly to control center



References

The NIST Definition of Cloud Computing.

<http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf>

Cisco's understanding of Fog Computing:

<https://www.youtube.com/watch?v=OVLJvtJCT1M&index=35&list=WL>

W. Shi et al., "Edge Computing: Vision and Challenges", IEEE Internet of Things Journal, vol.3, no.5, Oct. 2016.