CLOUD COMPUTING SYSTEMS

Lecture 12

Nuno Preguiça

(nuno.preguica_at_fct.unl.pt)

PROBLEMS OF CLOUD COMPUTING PLATFORMS

Reliability and security depends on the reliability and security of the cloud platform

- If the cloud platform fails, applications running in the cloud will also fail
- If there is a security breach in the cloud platform, data may be compromised

Vendor lock-in

- Reliance on cloud-specific services leads to vendor lock-in problems

OPENSTACK

OpenStack is an open-source software to manage a cloud platform.

Used to control large pools of compute, storage, and networking resources in a single or multiple datacenters.

Aims at providing all the necessary services to run a cloud platform, providing open standard interfaces.

OpenStack services

Nova is the service to provision compute instances (similar to EC2).

Supports for creating:

- Virtual machines;
- System containers;
- Baremetal servers (through the use of ironic).

OPENSTACK SERVICES (2)

Zun is the OpenStack Container service. It provides an API for running application containers without the need to manage servers or clusters.

The service runs two types of nodes:

- Controller nodes run all management code - e.g. image management, etc.
- Compute nodes run the containers.

OpenStack services (3)

Cinder is the OpenStack Block Storage service.

Manages volumes used by the Nova virtual machines, Zun containers, etc.

Provides **Highly available and Fault-Tolerant**, by relying on replication.

OpenStack services (4)

Swift is a highly available, distributed, eventually consistent object/blob store.

- DHT-based organization.
- Asynchronous replication.

MULTI CLOUD COMPUTING

Multi cloud computing is the use of multiple cloud computing and storage services in a single heterogeneous architecture.

The idea is to build an application that runs in multiple cloud platforms.

Why?

- Avoid vendor lock-in.
- Provide better reliability and security by combining resources in multiple platforms.
- Build on the best services on different platforms.

OpenStack MultiCloud

OpenStack supports multi-cloud solutions, by relying on its open standard APIs.

Allow to deploy and run services in multiple site – there is an Orchestration service responsible to manage the multi-cloud deployment.

DevOps specify the configuration of the multi-cloud solution in an Orchestration template, which specify the resource types to be used.

Other solution for Multi-Cloud

There are a number of third-party applications that can be used to build multi-cloud deployments.

Some examples include:

- Google Anthos
- VMware vRealize Cloud Management
- …

SERVICE-LEVEL MULTI-CLOUD SOLUTIONS

There are also service-level multi-cloud solutions.

Example: blob storage.

For data storage, several cloud platform provide blob storage service with S3 interface.

Several systems integrate S3 from multiple cloud providers for providing additional reliability and security, lower latency, etc.

OUTLINE

Moving to the edge

- Multi cloud computing
- **Hybrid cloud**
- Edge computing

HYBRID CLOUD

Hybrid cloud is a cloud combining two different cloud infrastructures. Hybrid cloud is most commonly used to refer a system that combines 1 private cloud and at least 1 public cloud.

Hybrid cloud systems incorporate some degree of workload portability, orchestration, and management across 2 or more environments.

BENEFITS

The hybrid cloud model tries to get the best of private and public cloud.

Sensitive, highly regulated, and mission-critical applications and workloads or workloads with constant performance and capacity requirements can run on private cloud infrastructure.

Less-sensitive, more-dynamic, or even temporary workloads (such as development and test environments for a new application) can run on the public cloud.

It is possible to leverage additional public cloud capacity to accommodate a spike in demand for a private cloud application (this is known as “cloud bursting”).

TYPES: HYBRID MONOCLOUD

Hybrid monocloud is hybrid cloud with one cloud provider — essentially it is an extension of a single public cloud provider's software and hardware stack to the customer's on-premises environment so that the exact same stack runs in both locations. The two environments are tethered together to form a single hybrid environment, managed from the public cloud with the same tools used to manage the public cloud provider's infrastructure.

TYPES: HYBRID MULTICLOUD

Hybrid multicloud is an open standards-based stack that can be deployed on any public cloud infrastructure. That means across multiple providers as well as on premises. As with hybrid monocloud, the environments are tethered together to form a single hybrid environment, but management can be done on- or off-premises and across multiple providers, using a common set of management tools chosen by the customer.

Hybrid multicloud architecture gives an organization the flexibility to move workloads from vendor to vendor and environment to environment as needed and to swap out cloud services and vendors for any reason.

CLOUD OPEN STANDARDS

Open standards are open to the public for use by anyone. Typically, the purpose of open standards is to allow for consistency and repeatability in approach.

In the case of hybrid cloud, open standards can help support interoperability, integration, and management. Some examples of open standards that support hybrid cloud include Kubernetes, OpenStack, and Cloud Foundry.

HYBRID CLOUD INTEGRATION

Integration across applications and data – in the cloud and on- and off-premises – is key to ensuring the components of the hybrid ecosystem work together quickly and reliably.

HYBRID CLOUD MANAGEMENT

Management includes, but is not limited to, provisioning, scaling, and monitoring across environments.

In a hybrid monocloud environment, management is relatively straightforward because with a single vendor, you can use the same tools to manage or provision across the infrastructure.

In a hybrid multicloud environment encompassing multiple cloud vendors, it is more of a challenge to manage consistently.

Kubernetes can help with management tasks like scaling containerized apps, rolling out new versions of apps, and providing monitoring, logging, debugging, etc.

HYBRID CLOUD STORAGE

Cloud storage allows to save data and files to an off-site storage provider. The provider hosts, secures, manages, and maintains the servers and associated infrastructure and ensures access to the data whenever needed.

A hybrid cloud storage model gives the choice of which data to store in which cloud. For instance, highly regulated data subject to strict archiving and replication requirements is usually more suited to a private cloud environment, whereas less-sensitive data (such as email that doesn't contain business secrets) can be stored in the public cloud. Some organizations use hybrid clouds to supplement their internal storage networks with public cloud storage.

OUTLINE

Moving to the edge

- Multi cloud computing
- Hybrid cloud
- **Edge computing**

EDGE COMPUTING

Nascent technologies and applications for mobile computing and the Internet of Things (IoT) are driving computing toward dispersion.

Edge computing is a new paradigm in which substantial computing and storage resources – known as cloudlets, micro datacenters, or fog nodes – are placed at the Internet's edge in close proximity to mobile devices or sensors.

From CDNs to Edge computing

Edge computing generalizes and extends the CDN concept by leveraging cloud computing infrastructure. As with CDNs, the proximity of cloudlets to end users is crucial. However, instead of being limited to caching web content, a cloudlet can run arbitrary code just as in cloud computing.

This code is typically encapsulated in a virtual machine (VM) or a lighter-weight container for isolation, safety, resource management, and metering.

BENEFITS OF EDGE COMPUTING

Highly responsive cloud services. A cloudlet's physical proximity to a mobile device makes it easier to achieve low end-to-end latency, high bandwidth, and low jitter to services located on the cloudlet. This is valuable for applications such as augmented reality that offload computation to the cloudlet.

Scalability via edge analytics/pre-processing. The cumulative ingress bandwidth demand into the cloud from a large collection of high-band- width IoT sensors, such as video cameras, is considerably lower if the raw data is analyzed on cloudlets. Only the (much smaller) extracted information and metadata must be transmit- ted to the cloud.

BENEFITS OF EDGE COMPUTING (2)

Privacy-policy enforcement. By serving as the first point of contact in the infrastructure for IoT sensor data, a cloudlet can enforce the privacy policies of its owner prior to release of the data to the cloud.

Masking cloud outages. If a cloud service becomes unavailable due to network failure, cloud failure, or a denial-of-service attack, a fallback service on a nearby cloudlet can temporarily mask the failure.

HIGHLY RESPONSIVE CLOUD SERVICES

Some computations can be slow (or have high energy consumption) when executed in mobile devices.

Offloading computations to the fixed infrastructure can help, but:

- Latency to data centers is high.
- Executing in a nearby (edge) location is much more efficient.

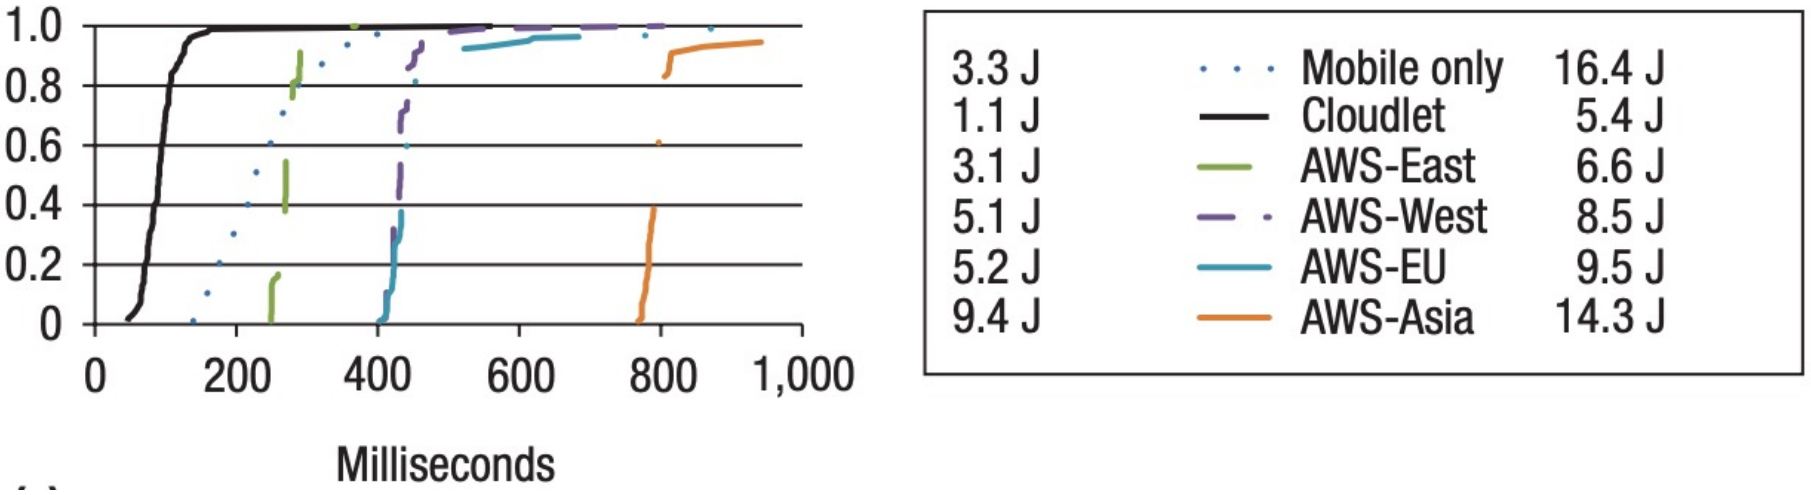

CHALLENGES

What to run in the edge (cloudlets)? What are the services and programming abstractions and API?

How to integrate with the core of the cloud?

- AWS Lambda @ Edge is proposing to address this challenge by supporting Lambda functions.
- Microsoft Azure has solutions for IoT.

EDGE ANALYTICS

Large IoT systems will generate too much information to be sent to the cloud.

A way to address this issue is to filter/preprocess information at the edge.

Note: the learning phase does not need to be done at the edge… only the classification.

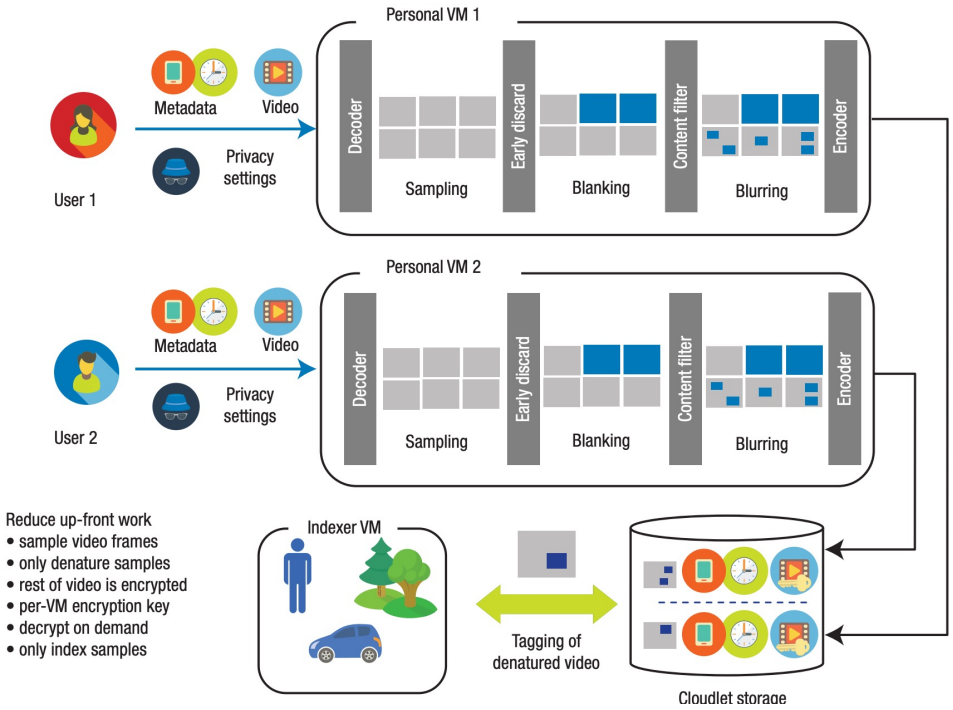

EDGE ANALYTICS (2)

All major cloud providers have solutions for running analytics at the edge for IoT.

Challenges:

- How to improve models based on new data?
- How to decide what can be discarded?

PRIVACY POLICY ENFORCEMENT

Executing computations at the edge may help addressing privacy (and legal) concerns.

Users and organizations desire finer-grain control over data released at IoT hub. For example, they wan to be able to delete some data, decrease the quality/only send aggregates of other data, etc.

Cloudlets can be used to execute the necessary transformations.

PRIVACY POLICY ENFORCEMENT: EXAMPLE

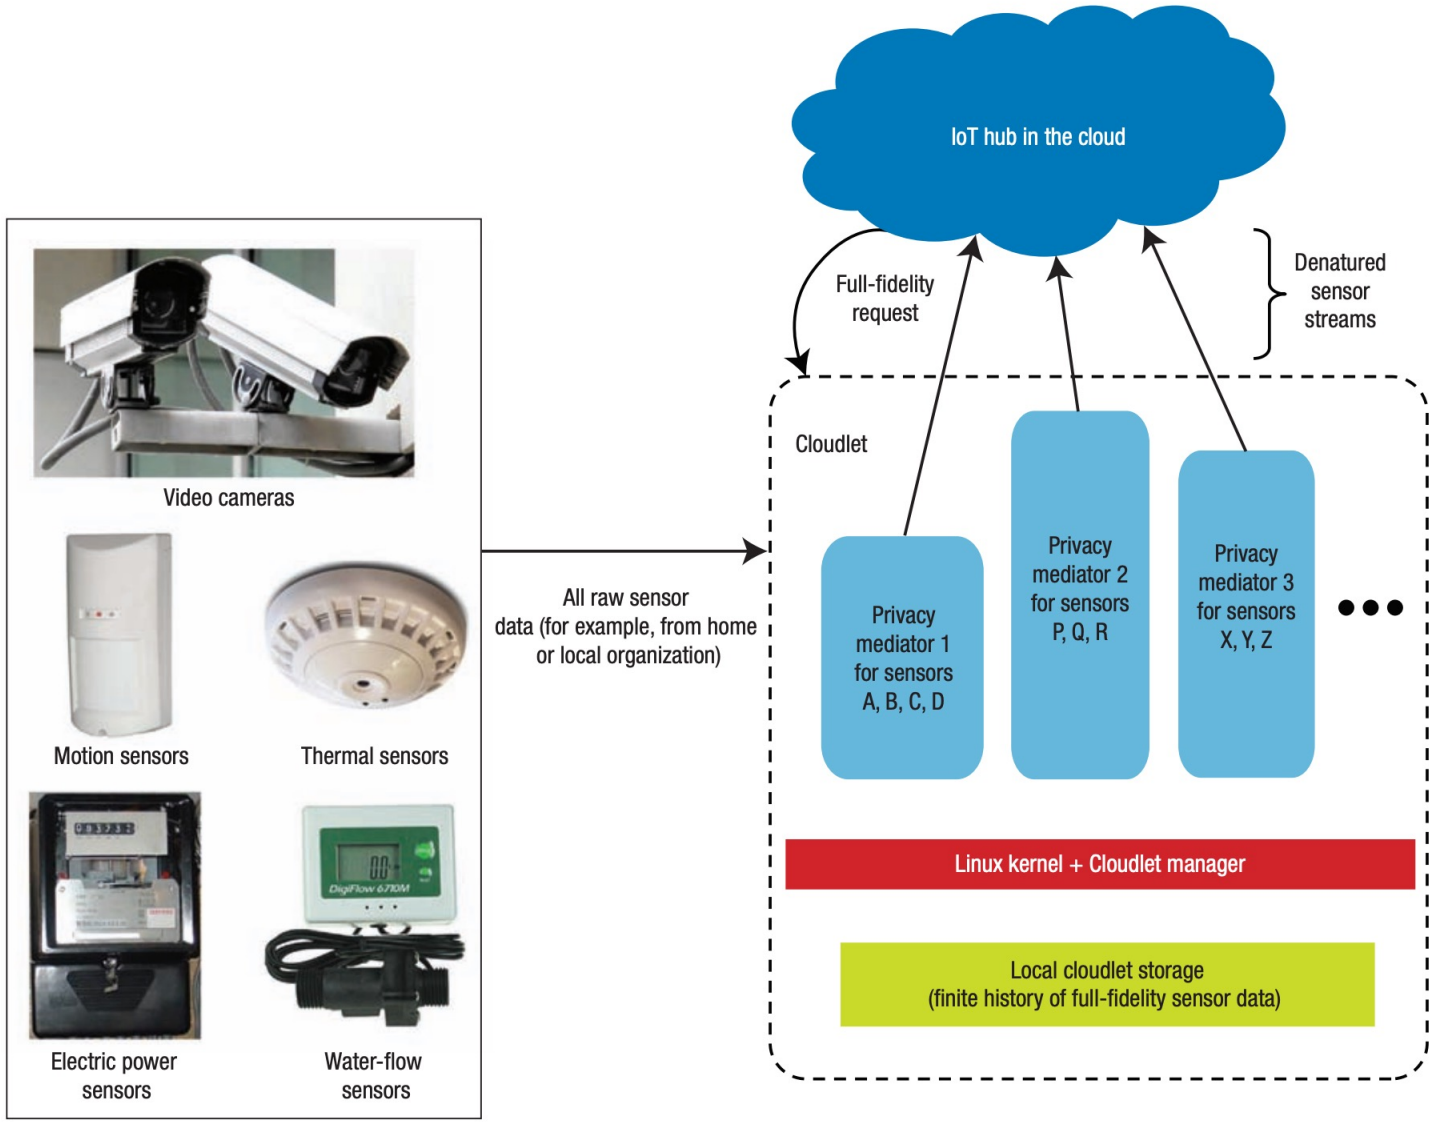

CHALLENGES

How to define what can data can be propagated and to where?

How to enforce the defined policies?

How to be able to execute useful computations with partial/degraded data?

How to guarantee that computation will not compromise privacy?

MASKING CLOUD OUTAGES

As the dependence on the cloud grows, so does the vulnerability to cloud outages. Implicit in the convergence of mobile and cloud computing is the assumption that the cloud is easily accessible at all times – in other words, there is good end-to-end network quality and few network or cloud failures.

There are scenarios where this might not be true:

- Military scenarios;
- Recovery from natural disasters;
- Weak networking at some places.

MASKING CLOUD OUTAGES (2)

Moving functionality to the edge can help addressing these scenarios.

Challenges:

- Guarantee the needed data is available – requires pre-fetching/hoarding;
- Be able to run while disconnected.

To know more

https://www.ibm.com/cloud/learn/hybrid-cloud

The Emergence of Edge Computing. Mahadev Satyanarayanan. 2016. http://elijah.cs.cmu.edu/DOCS/satya-edge2016.pdf

ACKNOWLEDGMENTS

Some text and images from the referenced papers.

Some slides based on a previous version by Paulo Lopes and Vitor Duarte.