Future generation of computational infrastructures and the role of cloud computing

Professor Sasu Tarkoma, Head of the Department
Department of Computer Science
University of Helsinki

Contents

University of Helsinki

- The largest and the oldest university in Finland
- Key data for 2014
  - 35 000 students
  - 8 200 employees
  - 300 subjects
  - 5 850 degrees/year
  - 480 PhDs/year
- Founded in Turku 1640
- Moved to Helsinki 1828

Faculty of Science / Kumpula Campus

Departments

- Chemistry
- Computer science
- Geosciences and Geography
- Mathematics and Statistics
- Physics

Department of Computer Science

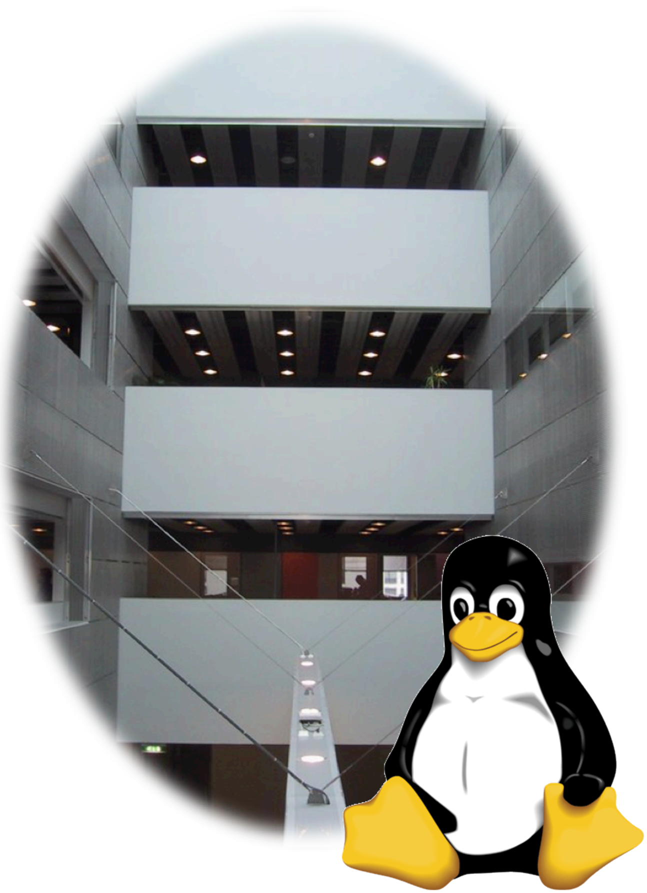

- Leading institution in Computer Science in Finland
- Students and employees 2014
  - 1 727 students (53 PhD students; nearly 30 000 credits)
  - 266 employees (168.7 FTE, 31.9% foreign, 20.3 % female)
- Part of the Faculty of Science
- Located in Exactum, Kumpula Campus
- Renowned for high quality of research and teaching
- The Linux kernel was originally developed at the Department by Linus Torvalds

Overview of Cloud Technology

Computing Environment

- Scaling up
  - Having more powerful servers
- Scaling out
  - Having more servers
- Clusters provide computing resources
  - Space requirements, power, cooling
  - Most power converted into heat
- Datacenters
  - Massive computing units
  - **Warehouse-sized computer** with hundreds or thousands of racks
- Networks of datacenters

Cluster Computing Environment

- Big Data compute and storage nodes are stored on racks based on common off the shelf components
- Typically many racks in a cluster or datacenter
- Intra-rack and inter-rack communication have differing latencies
- Nodes can experience failures
  - Computation in tasks or jobs
  - Replication for fault tolerance
- **Placement of tasks and data is important**
- **Software ensures fault-tolerance and availability**

Cloud Computing

Definition by NIST:

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.

Typically common-of-the-shelf servers

Compute nodes, storage nodes, …

Virtualized resources running on a cloud platform

IaaS, PaaS, SaaS, XaaS

Workloads are Evolving
Server vs Cloud

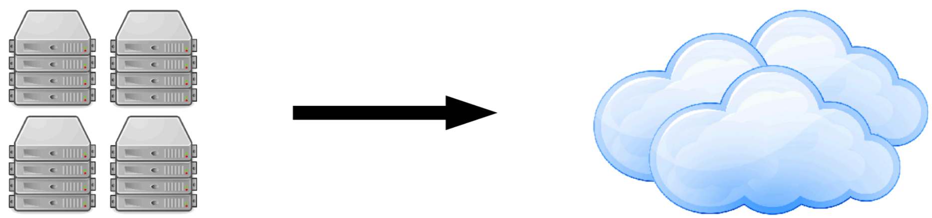

TRADITIONAL WORKLOADS	CLOUD WORKLOADS
Physical servers	Virtualized elastic resources
Needs expensive high availability tools	Tolerates VM failure – if one fails, another replaces it
Application scales up rather than out	Fault tolerance often built into the workload
	Application scales out rather than up

OpenStack Architecture

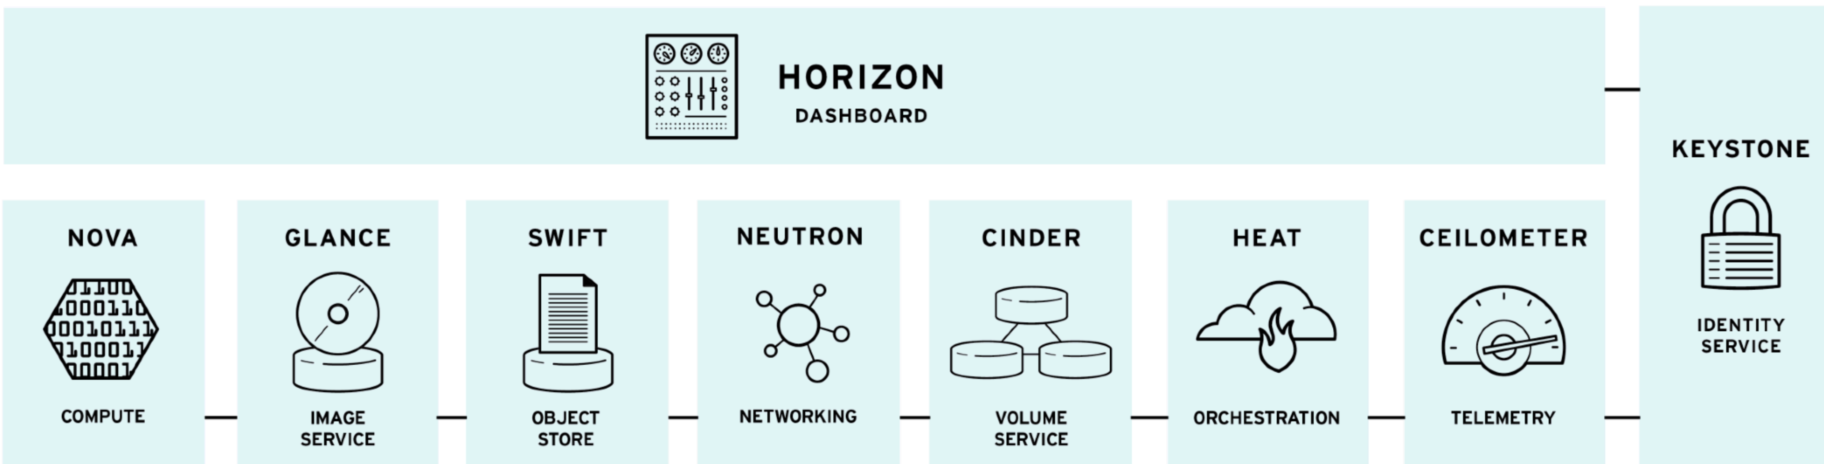

Modular architecture made of individual autonomous components

Components expose RESTFul API for communications

Components have stateless worker nodes and rely on messaging

Framework that is designed for scaling out

Based on a set of core services (largely Python and based on Linux)

Big Data (Hadoop/Spark) support with Sahara

Analytics cloud for smartphones

Carat Team

carat.cs.helsinki.fi

Motivation

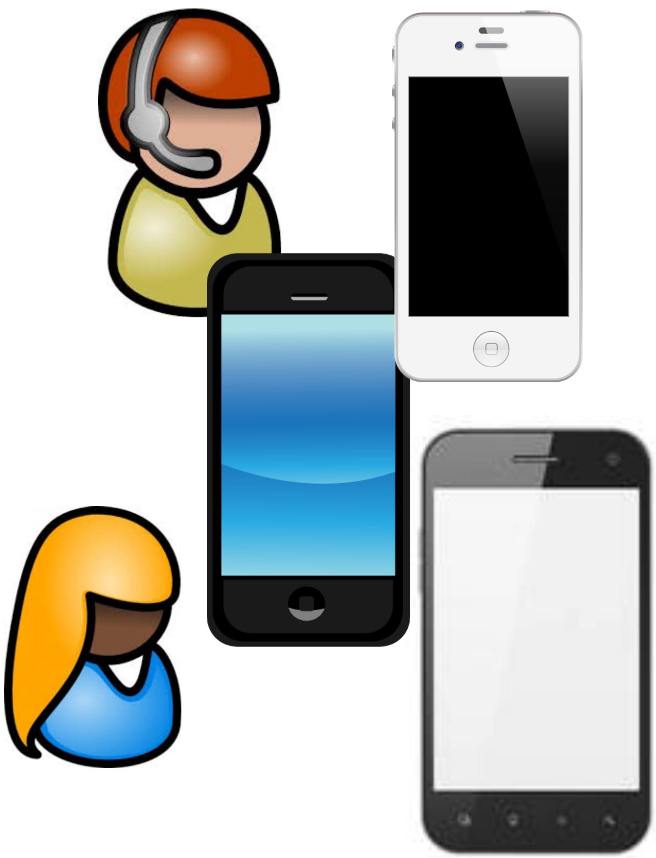

Battery lifetime?

Risk level?

A lot of heterogeneous, active devices and lot of users with different intents. – What kind of behavior is **normal** or **typical**?

Introducing Carat

Carat is the **first system** to use the mobile device community to detect and correct energy problems

Our method for **diagnosing** energy anomalies uses the community to infer a specification (expected energy use), and we call deviation from that inferred specification an anomaly

Carat

- Originated in UC Berkeley, in collaboration with University of Helsinki
- Mobile app for Android and iOS
- Currently over 848 000 users
- >2TB of data, > 100 million measurements
- Research project with many directions
- http://carat.cs.helsinki.fi

The Carat project: System

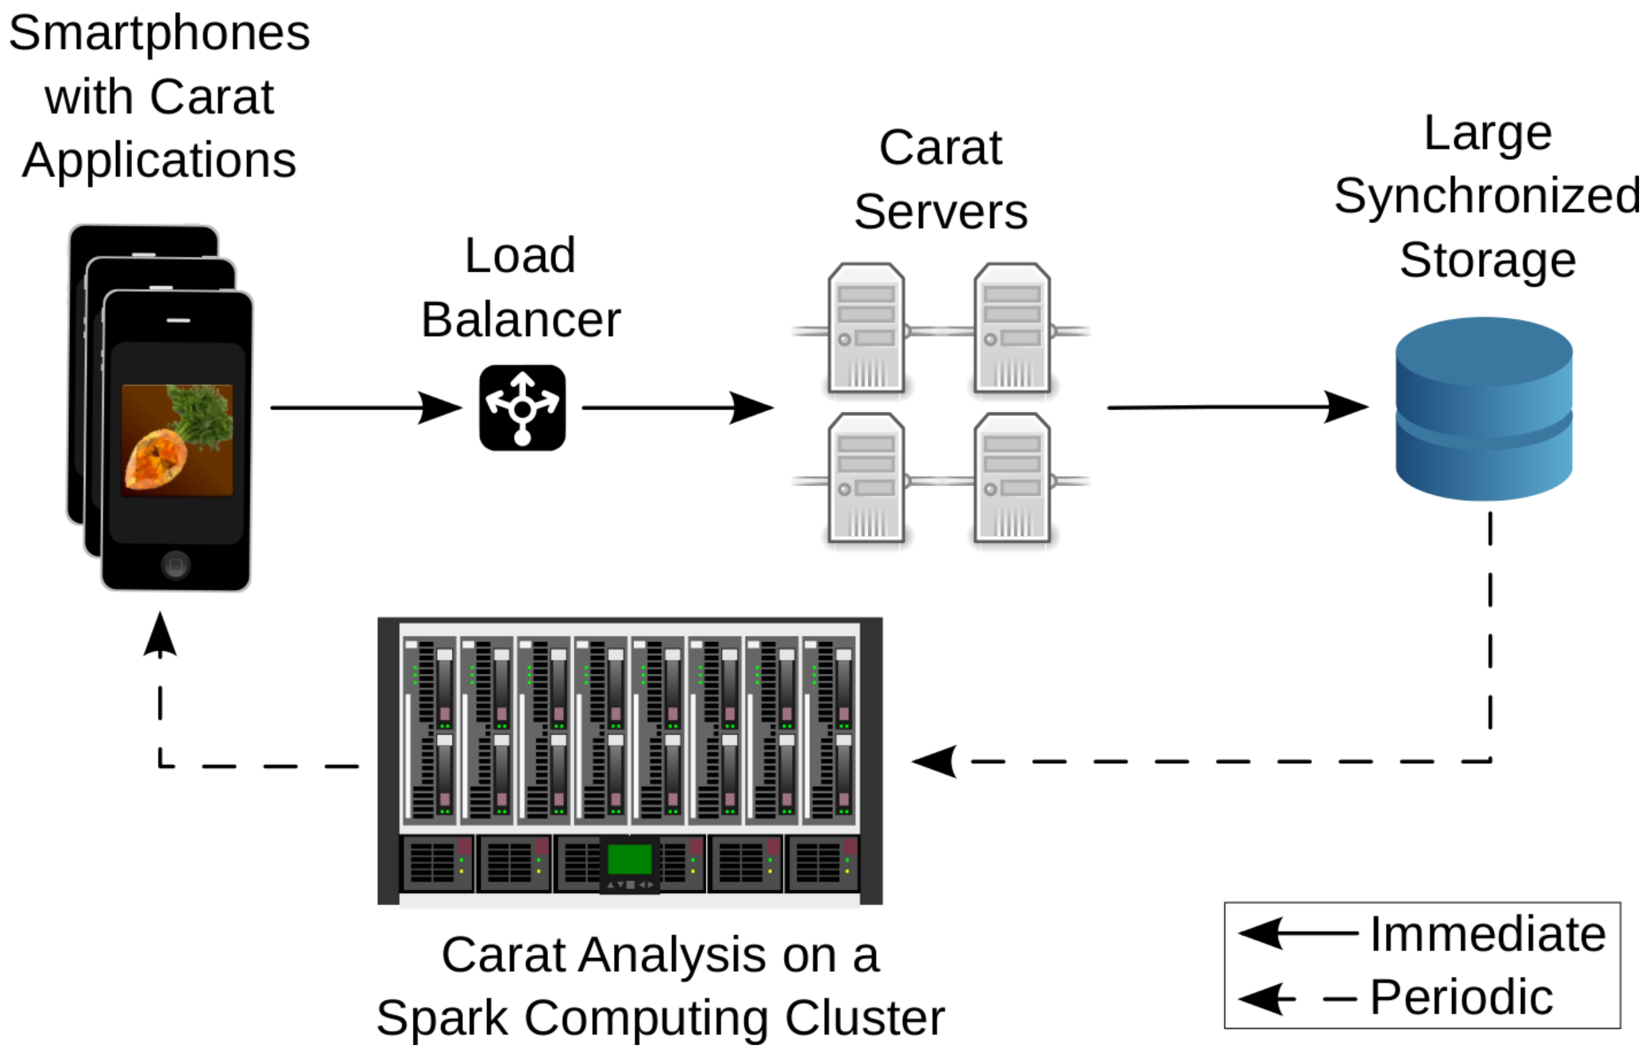

What is Carat?

- Users see Hogs, high energy use apps
- And Bugs that use energy faster on THEIR device than on others
- Users with these issues quickly see battery life benefits once they are taken care of

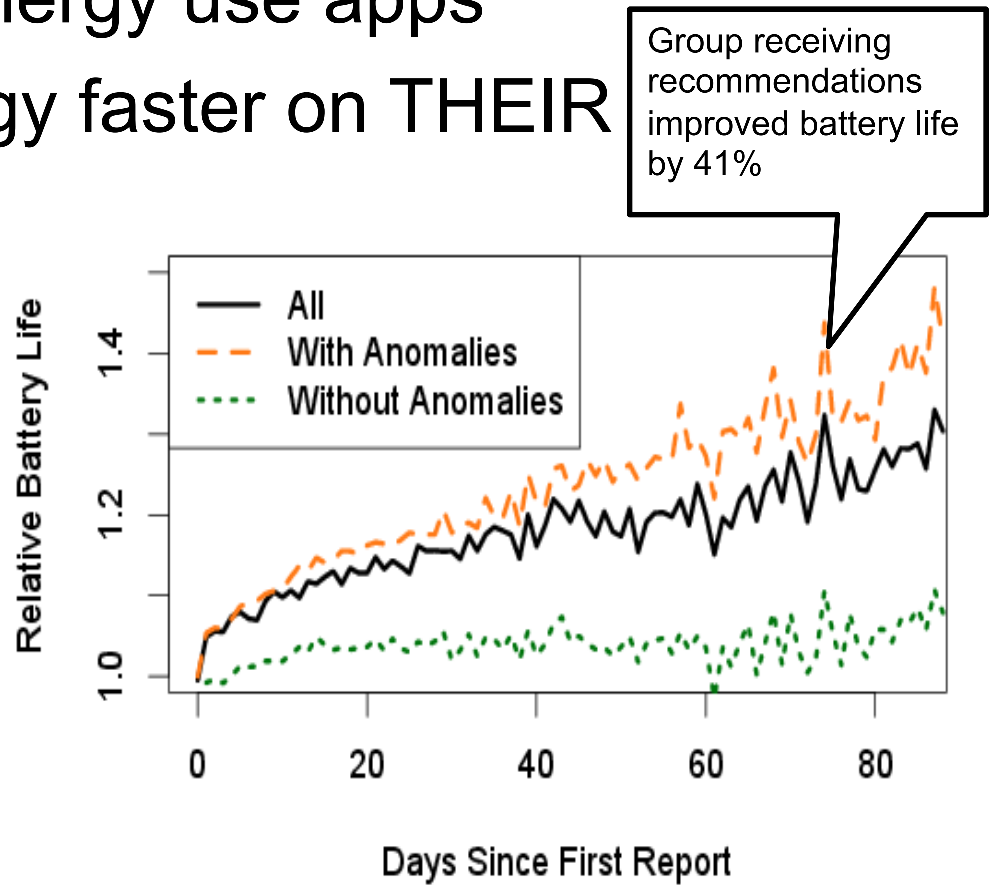

Collaborative Data Gathering

Each device collects

- Battery life, timestamp, running apps, system settings

The data is combined and results for your apps and your device sent back to you

Collaborative aspect: We know trends in the community, as well as how your device is different

This can be used for phones, sensors, houses, base stations, servers, laptops, … anything that generates measurements

An Early Warning System for Malware

A lightweight technique for identifying devices at risk

By looking at applications that occur with malware, it is possible to predict infection 5x better than choosing devices at random

- Useful for administrators, organisations (**Bring Your Own Device** scenario)

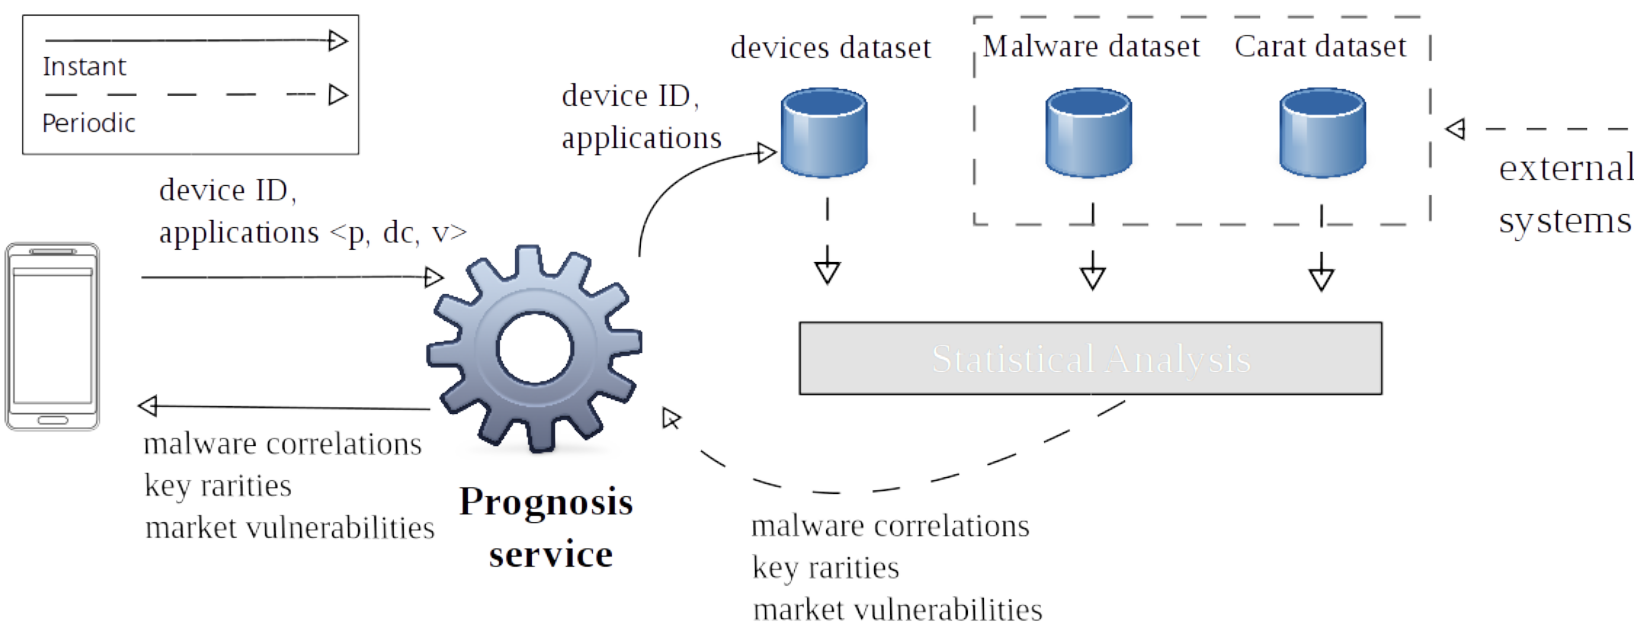

Related Publications

- A. J. Oliner, A. P. Iyer, I. Stoica, E. Lagerspetz, S. Tarkoma. Carat: Collaborative Energy Diagnosis for Mobile Devices. In ACM SenSys 2013.
- A. J. Oliner, A. Iyer, E. Lagerspetz, S. Tarkoma, I. Stoica. Carat: Collaborative energy debugging for mobile devices. In HotDep 2012.
- A. J. Oliner, A. P. Iyer, E. Lagerspetz, I. Stoica, and S. Tarkoma. Carat: Collaborative Energy Bug Detection. Poster and demo at the proceedings of the 9th USENIX Symposium on Networked Systems Design and Implementation (NSDI '12), San Jose, California.
- K. Athukorala, E. Lagerspetz, M von Kügelgen, A. Jylhä, A. J. Oliner, S. Tarkoma, G. Jacucci. How Carat Affects User Behavior: Implications for Mobile Battery Awareness Applications. ACM CHI 2014.
- H.T. T. Truong, E. Lagerspetz, P. Nurmi, A. J. Oliner, S. Tarkoma, N. Asokan, S. Bhattacharya, The Company You Keep: Measuring Mobile Malware Infection Rates and Identifying Inexpensive Predictors of Susceptibility to Infection, Proceedings of WWW 2014.
- E. Lagerspetz, H. Truong, S. Tarkoma, N. Asokan. Mdoctor - A Mobile Malware Prognosis Application. DASec workshop in conjunction with ICDCS 2014.
- E. Peltonen, E. Lagerspetz, P. Nurmi, and S. Tarkoma. Energy Modeling of System Settings: A Crowdsourced Approach, IEEE PerCom '15. **(Best paper award).**
- S. Tarkoma, M. Siekkinen, E. Lagerspetz, Y. Xiao. “Smartphone Energy Consumption: Modelling and Optimization”, August 2014, Cambridge University Press.
- E. Lagerspetz. Collaborative Mobile Energy Awareness. PhD thesis. University of Helsinki. November 2014. (**UH Dissertation Award 2014**).

Secure Scientific Cloud: Datacenter Indirection Infrastructure for Secure HEP Data Analysis

Collaboration between

Computer Science Department, University of Helsinki

Sasu Tarkoma, Lirim Osmani

Helsinki Institute of Physics (HIP)

Paula Eerola, Tomas Lindén, John White, Salman Toor

Funded by Academy of Finland 2012 - 2014

Cloud based setup

We have a production CMS site based on private cloud setup

OpenStack

Gluster Filesystem

Advanced Resource Connector (ARC) middleware for providing grid interfaces

CERN VM File System (CVMFS)

OpenStack deployed on Ubuntu 12.04 LTS

VMs based on Scientic Linux CERN 6.4

System Architecture

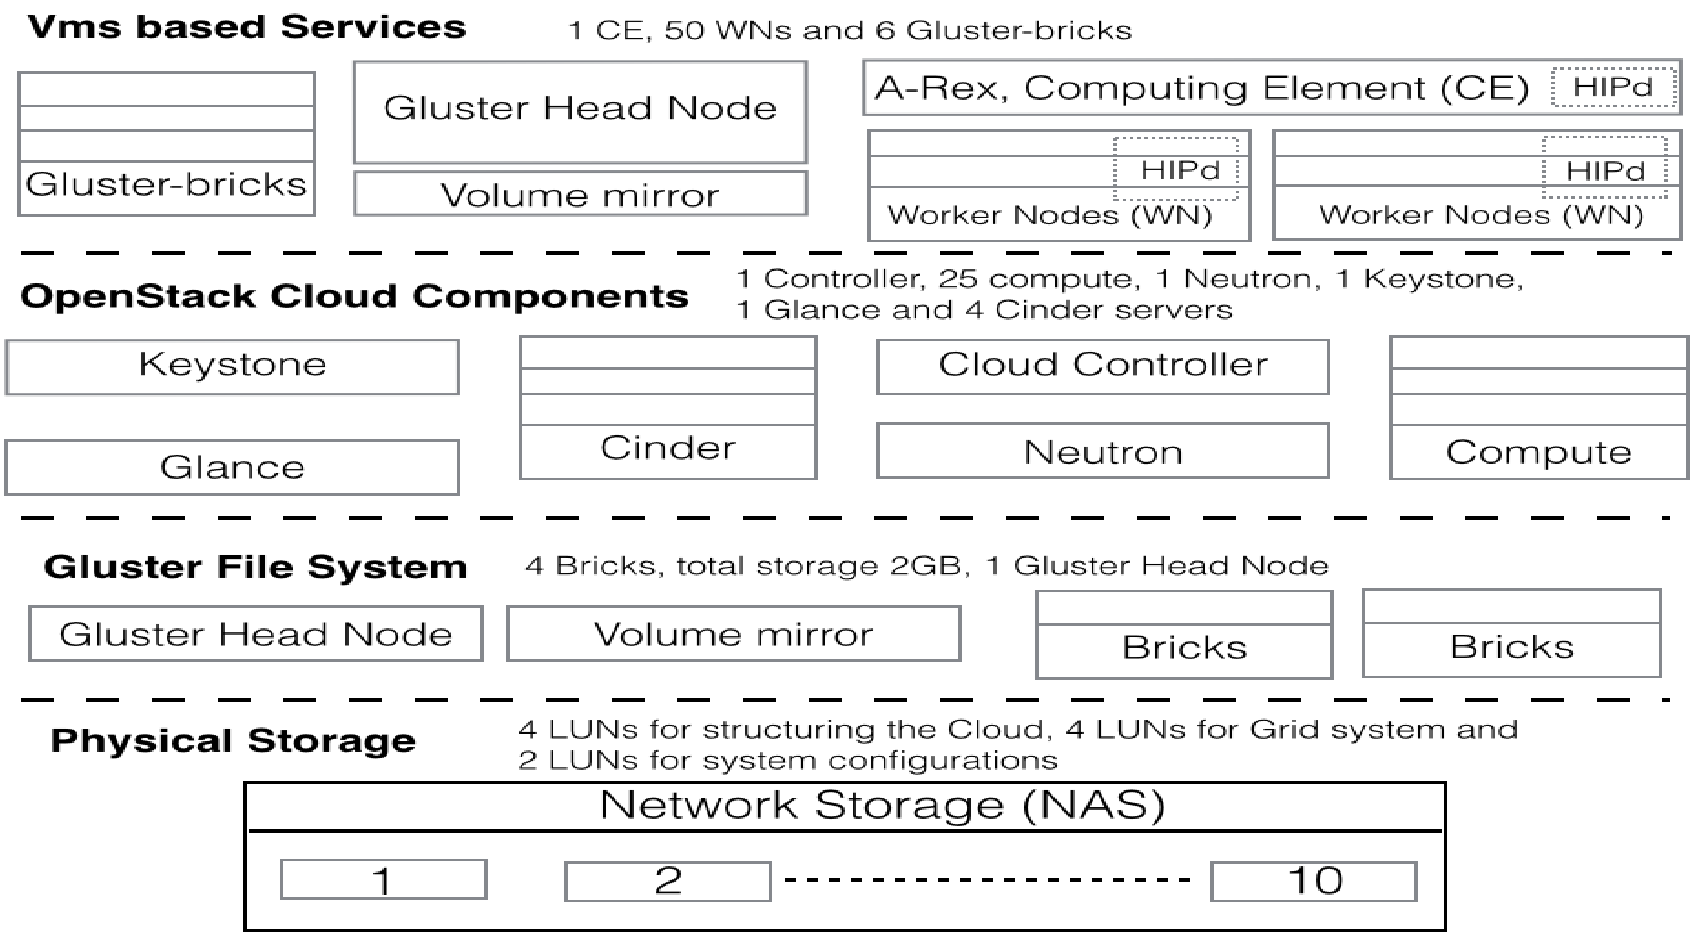

Secure cloud setup

HELSINKI INSTITUTE OF PHYSICS

UNIVERSITY OF HELSINKI
DEPARTMENT OF COMPUTER SCIENCE

We have used Host Identity Protocol (HIP) for structuring the secure cloud

Host Identity Protocol (HIP)

- Designed for mobile networks
- Provides a secure mechanism for IP multihoming and mobility (VM migration)
- HIP separates the end-point identifier and locator roles of IP addresses
- Provides persistent cryptographic identifiers
- Supports both IPv4 and IPv6 addressing
- The Host Identifiers (HI) are not routable, so they are translated into routable addresses (locators) between network and transport layer
- The HIP connections are typically protected with IPSec

HIP -> http://infrahip.hiit.fi

Secure hybrid cloud

Tenant driven solution:
Securing intra and inter cloud VM communication with the Host Identity Protocol (HIP)

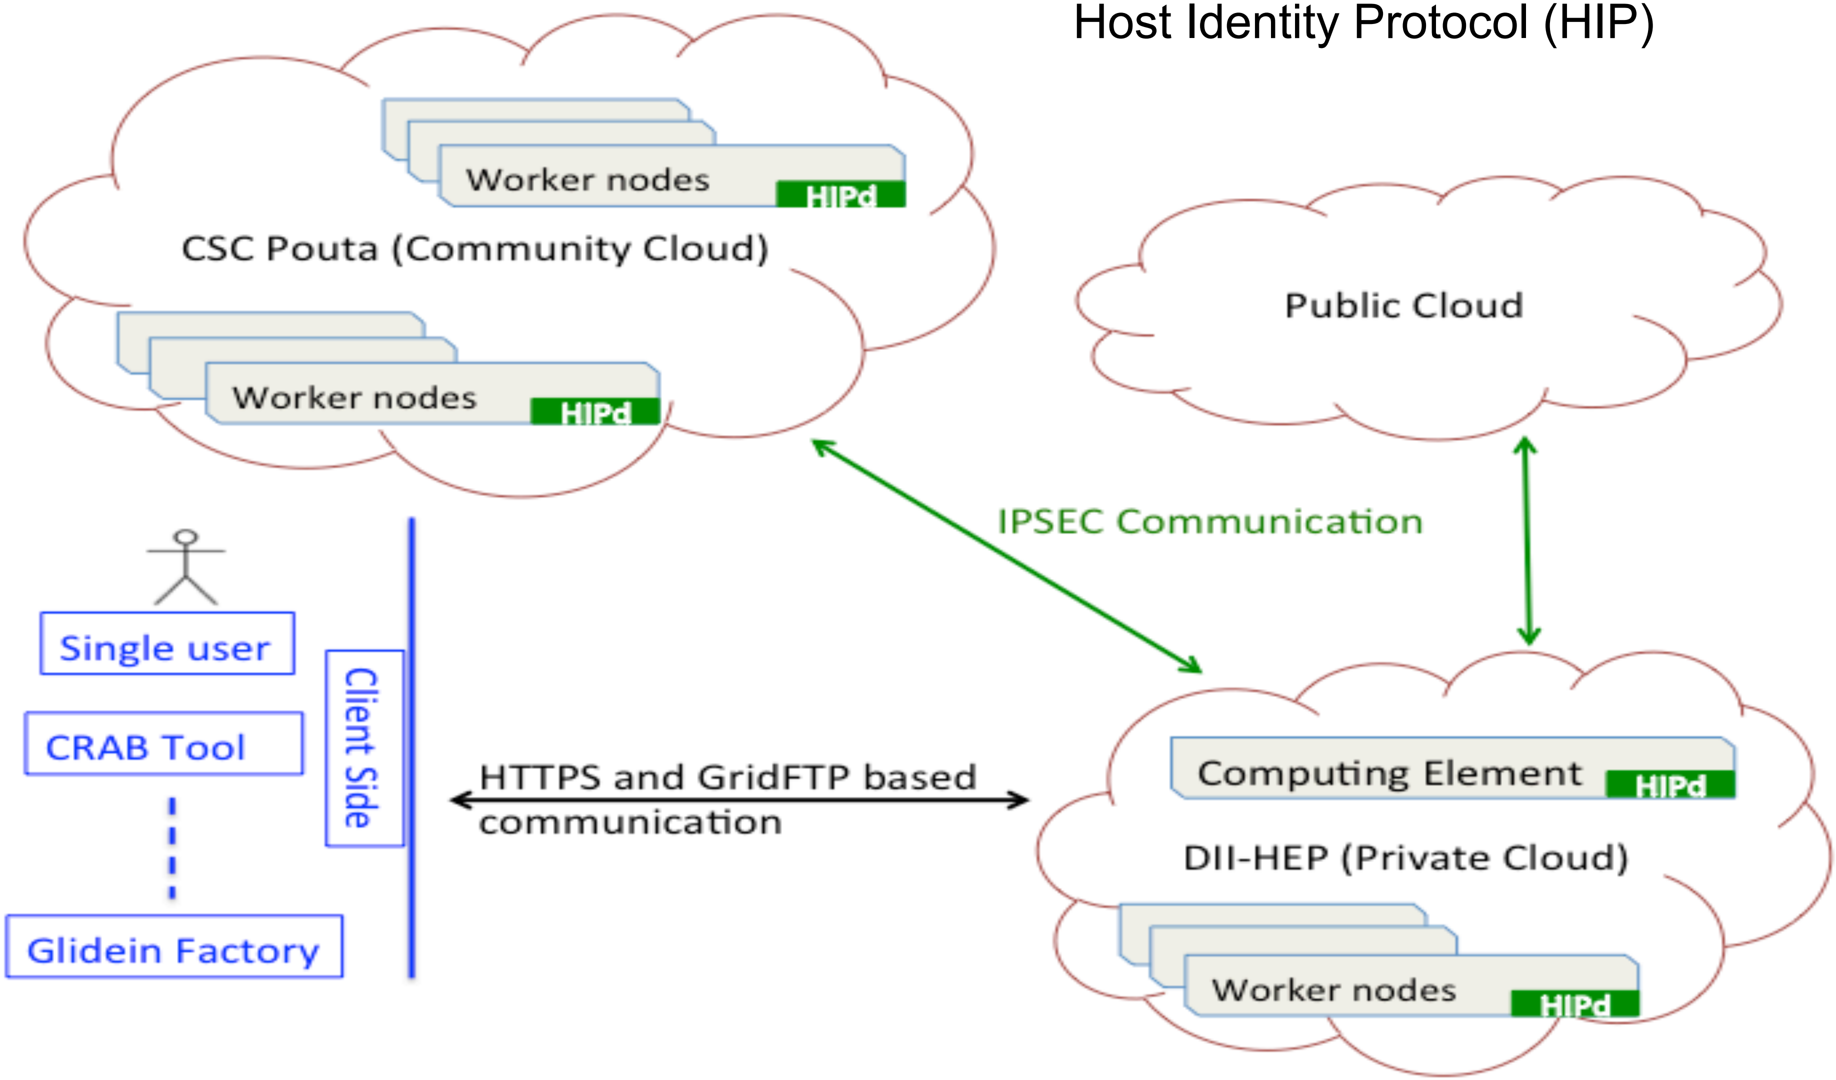

Performance analysis: Impact of Running CMS in the Cloud

Application Level Performance

- 4% performance loss evaluated with the HEPSPEC-2006 (Thanks to Ulf Tigerstedt, CSC for help with HEPSPEC tests)

System Level Performance

- VM boot response both at local vs GlusterFS based setup
- GlusterFS gives an acceptable VM startup time compared to local disk

Security performance of HIP is comparable to other VPN solutions

- Large performance penalty on throughput, negligible impact on latency
- Our results indicate that SSL/TLS or HIP as a security measure do not drastically impact the performance footprint in production environment

Related publications

L. Osmani, S. Toor, M. Komu, M. J. Kortelainen, T. Lindén, J. White, R. Khan, P. Eerola, S. Tarkoma. **Secure Cloud Connectivity for Scientific Applications**. IEEE Transactions on Services Computing, 2015.

J. White, S. Toor, P. Eerola, T. Lindén, O. Kraemer, L. Osmani, S. Tarkoma, **Dynamic Provisioning of Resources in a Hybrid Infrastructure**, PoS(ISGC2014)019.

S. Toor, L. Osmani, P. Eerola, O. Kraemer, T. Lindén, S. Tarkoma, J. White. **A scalable infrastructure for CMS data analysis based on OpenStack Cloud and Gluster file system**. Journal of Physics: Conference Series 513 062047 doi:10.1088/1742-6596/513/6/062047.

4G and 5G Core Network

5^{th} Generation Mobile Networks (5G)

- Beyond 4G for early 2020s
- Significant improvements in wireless communication
  - Smart radios and spectrum sharing
  - 1000 times higher spectral efficiency
  - Cooperative relays and femtocells
- Device-to-Device communication
- Support for Internet of Things and Machine-to-Machine
- World Wide Wireless Web
- SDN and cloud for the core network

Virtualization of Resources

Open Networking Foundation (ONF) has identified SDN in future mobile networks where inter-cell interference management and mobile traffic management are the key use cases

Network Function Virtualization (NFV) ETSI Industry Specification Group

Started in 2013

Network Nodes as software on top of COTS Hardware

Separation between hardware and software

Main Functions of the LTE Core

Route packets

Control network access

Support client mobility

Provide network security

Manage network functions

4G LTE Evolved Packet Core (EPC)

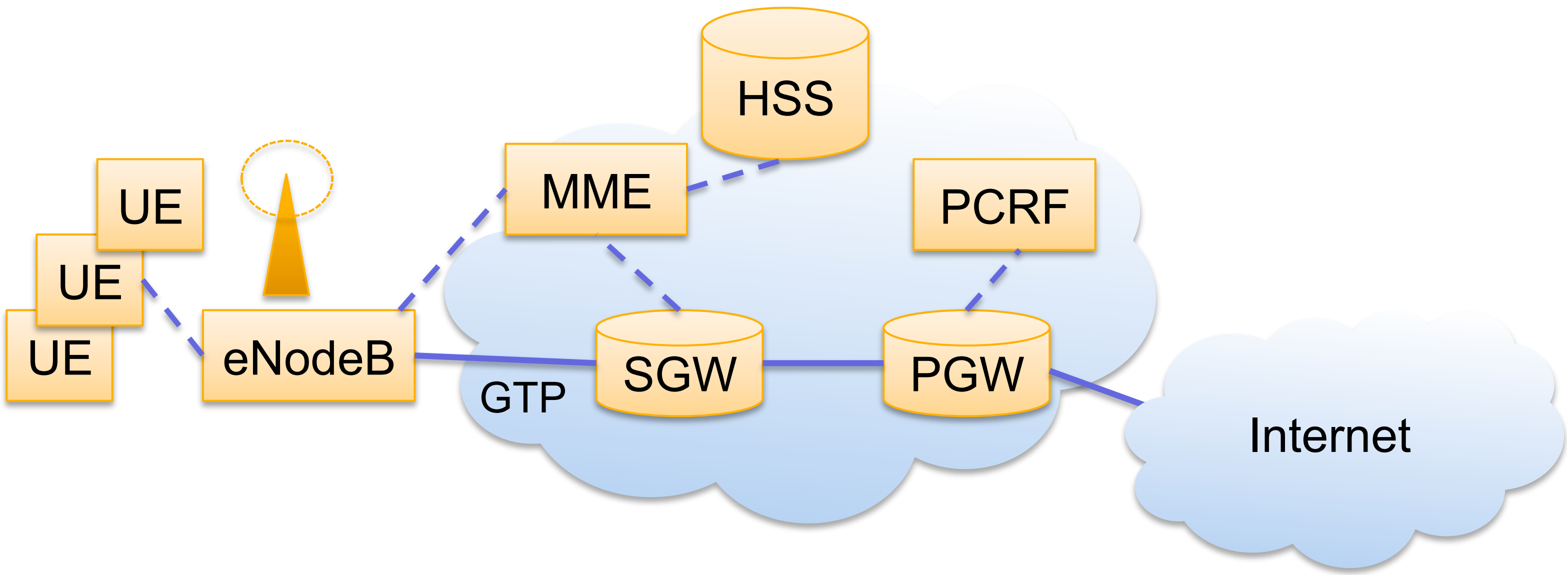

GTP: GPRS Tunneling Protocol for IP-over-UDP control and data
MME: Mobile Management Entity
SGW: Serving Gateway: forwards user traffic and mobility anchor
PGW: Packet Data Network Gateway: external networks and billing
HSS: Home Subscriber Service
PCRF: Policy Charging and Rules Function

LTE EPC with SDN

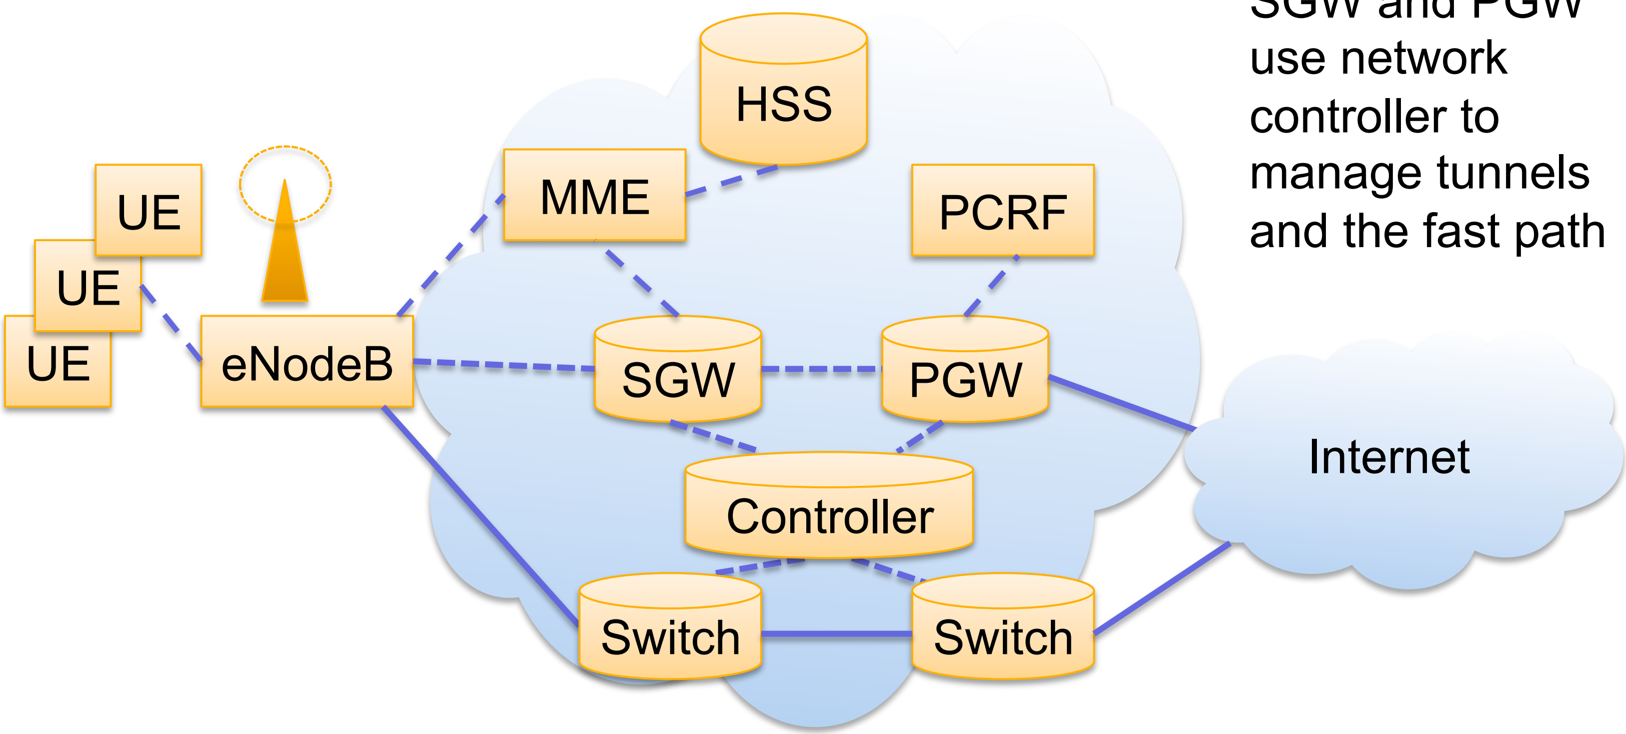

SGW and PGW use network controller to manage tunnels and the fast path

LTE EPC with SDN and Cloud

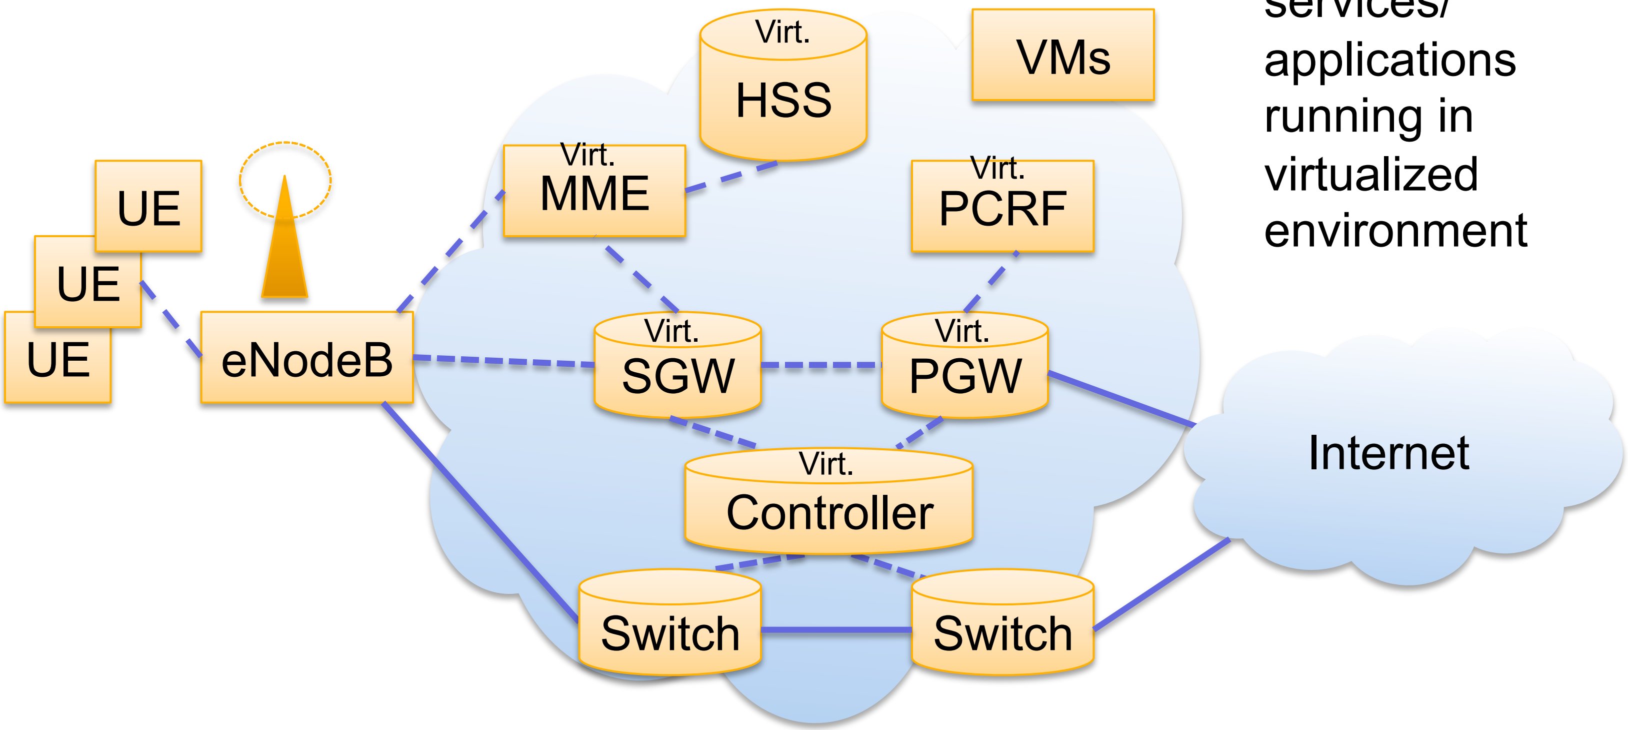

5G elements as services/ applications running in virtualized environment

LTE RAN and EPC with SDN and Cloud

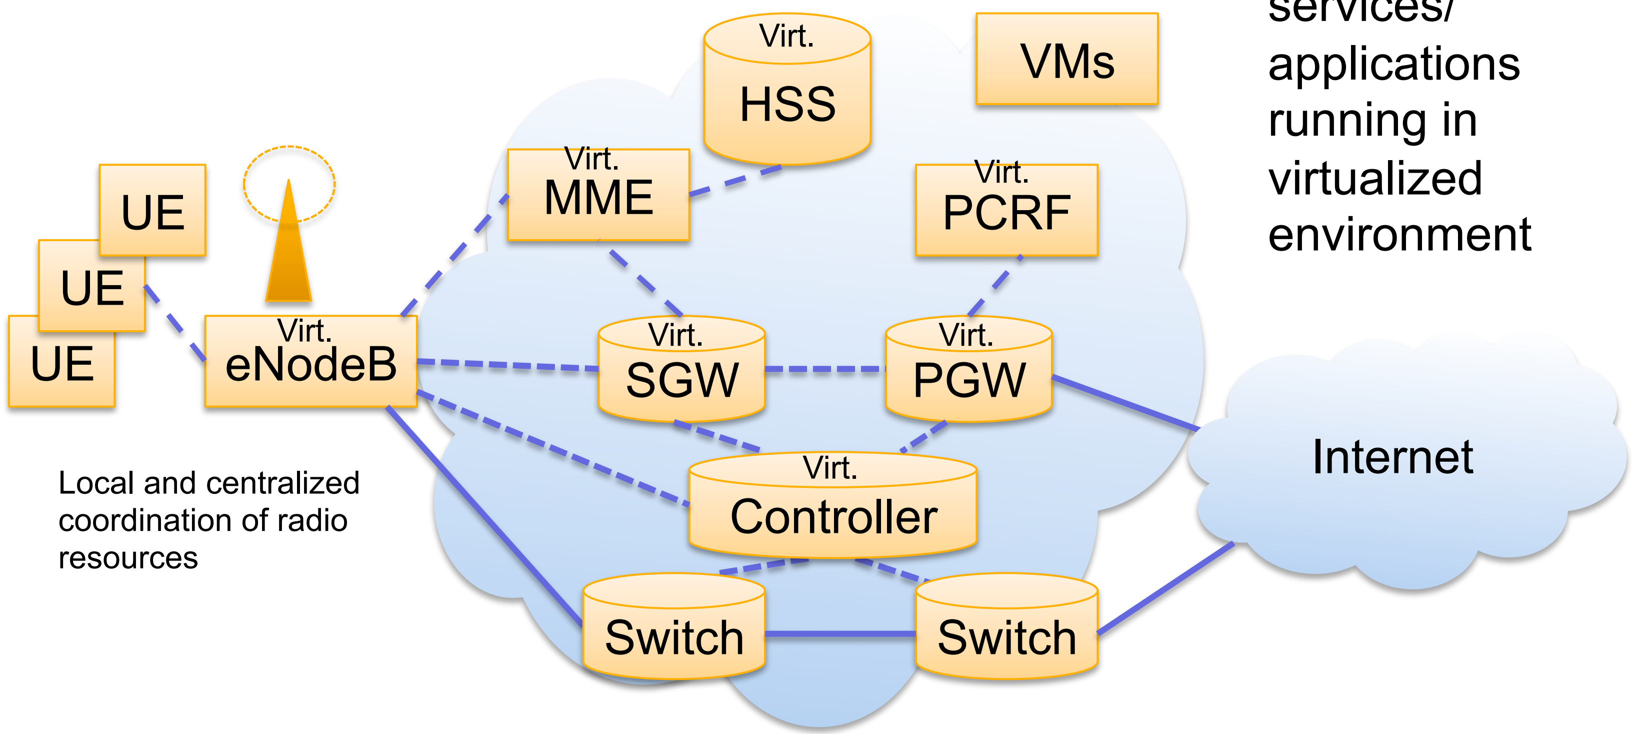

5G elements as services/ applications running in virtualized environment

Mobile Edge Computing

The cloudlet architecture from CMU consists of customized ephemeral virtual machines with soft state, and a platform for running them

Nokia Liquid Applications run on base stations

Deploy applications near the users to avoid latency and bandwidth problems

Facilitates elastic and mobile execution of network components and application logic in base stations

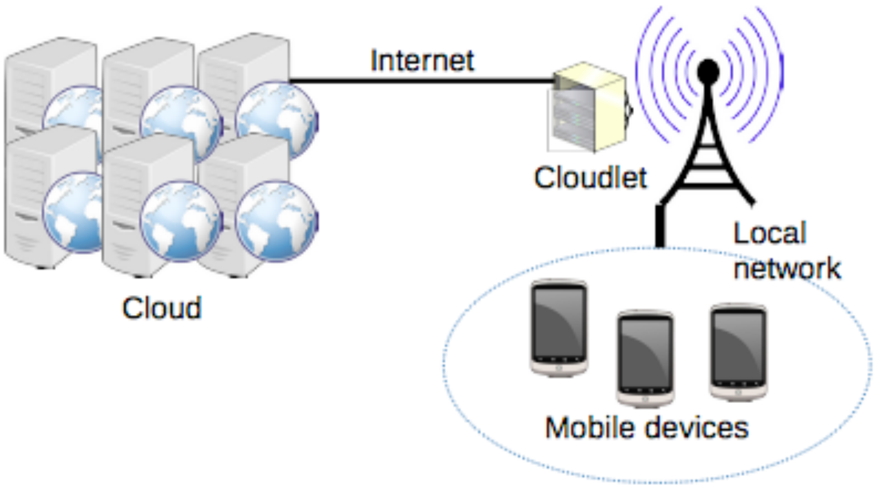

Related publications

Aaron Yi Ding, Jon Crowcroft, Sasu Tarkoma, Hannu Flinck: **Software defined networking for security enhancement in wireless mobile networks**. Computer Networks 66: 94-101 (2014).

Heikki Lindholm, Lirim Osmani, Hannu Flinck, Sasu Tarkoma, Ashwin Rao. **State Space Analysis to Refactor the Mobile Core**. AllThingsCellular workshop in conjunction with ACM Sigcomm, August 17, 2015, London, United Kingdom.

Jose Costa-Requena, Jukka Manner, Raimo Kantola, Aaron Yi Ding, Sasu Tarkoma. **Software Defined 5G Mobile Backhaul**. 5GU 2014.

Huber Flores, Pan Hui, Sasu Tarkoma, Yong Li, Satish Narayana Srirama, Rajkumar Buyya: **Mobile code offloading: from concept to practice and beyond**. IEEE Communications Magazine 53(3): 80-88 (2015).

Conclusions

Conclusions

The current and emerging digital infrastructure builds on cloud technology and virtualization

Cloud technology offers flexibility and elasticity as well as management capability

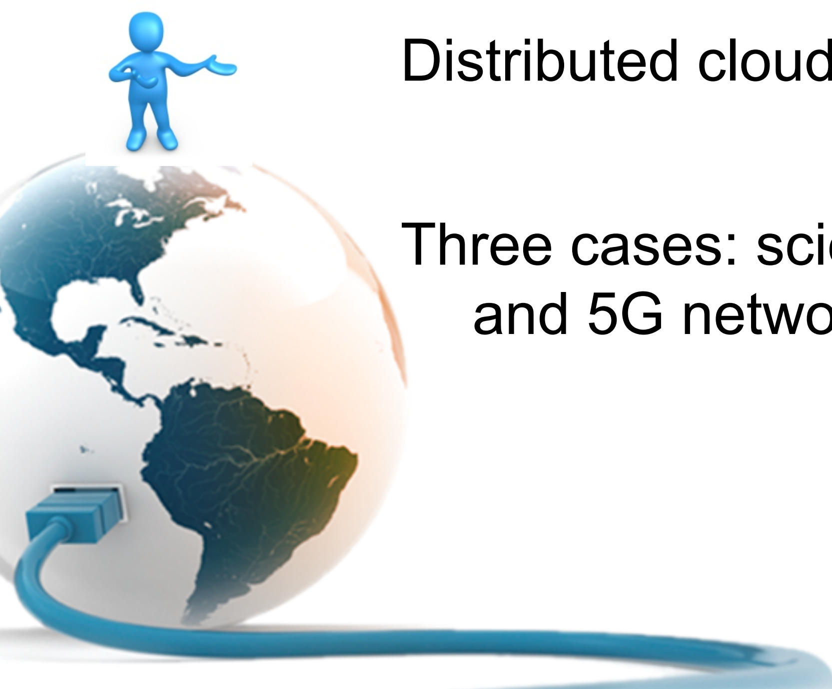

Distributed cloud and mobile edge computing

Three cases: scientific clouds, smartphone analysis, and 5G networks