Introduction to Cloud Computing : IaaS

Infrastructure as a Service (IaaS)
Hosted Infrastructures

Version M2 TI+DS+CCI+SRIV

Project Nebula

Manager : Jean-Patrick GELAS
Authors : JP. GELAS, G. KEMP
Université Claude Bernard - Lyon 1

Lifelong Learning Programme

NEBULA

Outline

- Introducing the IaaS model : Definition and Objectives
- Notions about *Green computing* and energy efficiency
- Architectural elements of a IaaS
  (Servers, Storage, Networks,…)
- Software elements of a IaaS
  (Virtualisation, Openstack platform deployment, Automated deployment and orchestration)
- Current solutions on market (non exhaustive)
- Conclusions

Cloud Computing

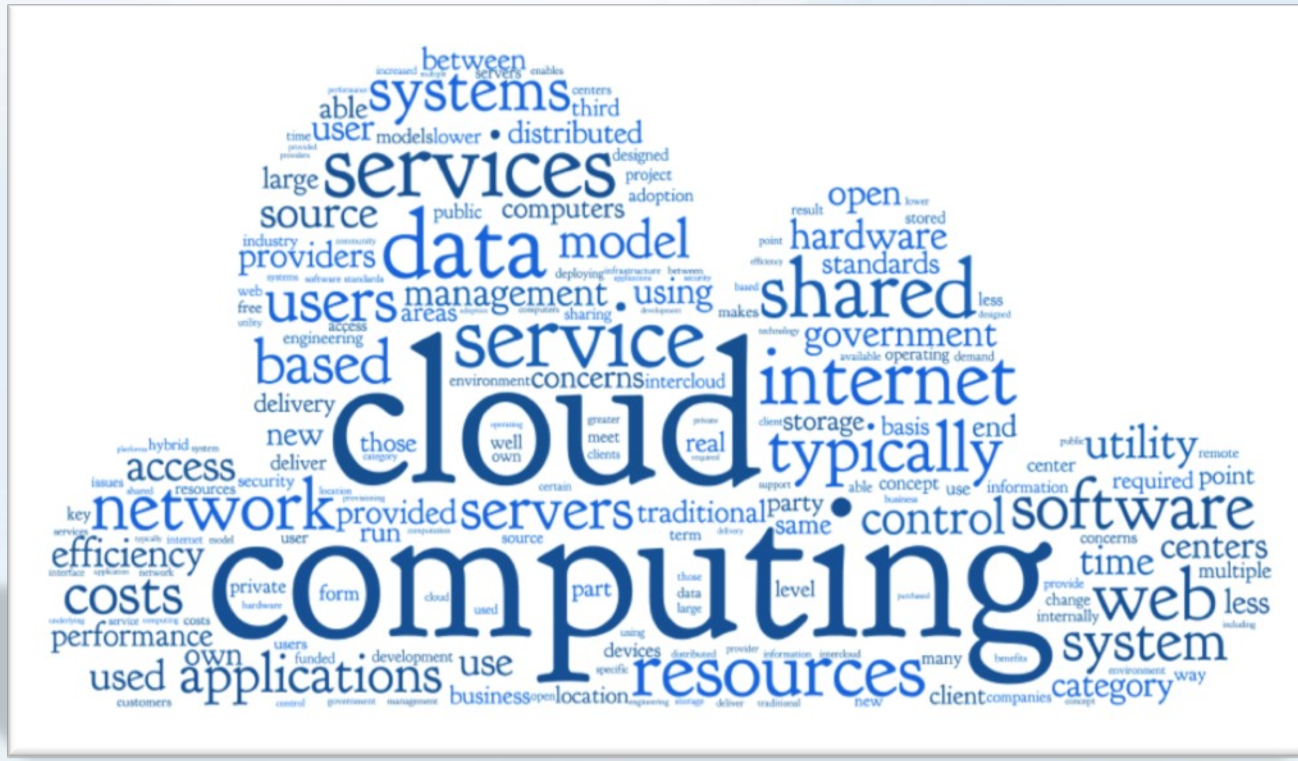

Définition

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.

Définition

The interesting thing about cloud computing is that we've redefined cloud computing to include everything that we already do. [...] **The computer industry is the only industry that is more fashion-driven than women's fashion.** Maybe I'm an idiot, but I have no idea what anyone is talking about. What is it? It's complete gibberish. It's insane. When is this idiocy going to stop?"

Larry Illison (Fondateur d'Oracle)

Définition

… faire largement appel au cloud computing est la pire des stupidités.

Richard Stallman

La fin du buzz ?

Recherche du mot Cloud Computing sur le net

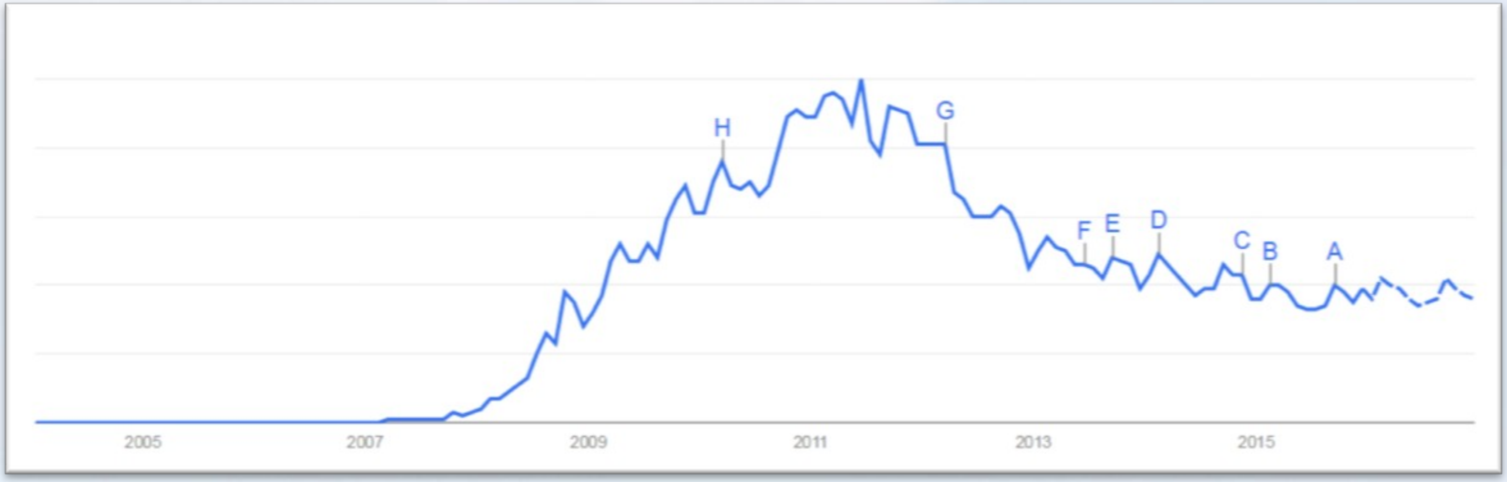

Source : Google Trends

Lifelong Learning Programme

NEBULA

Un marché en progression

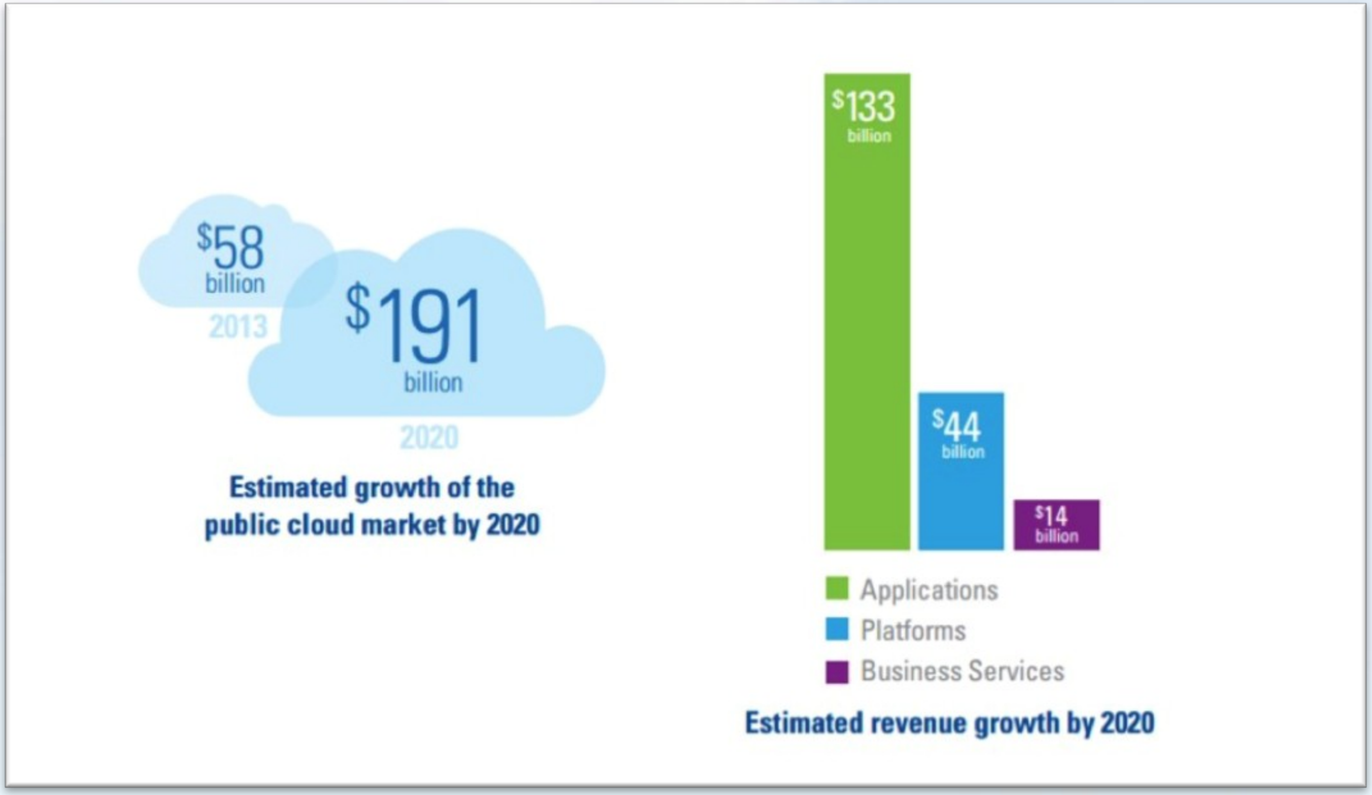

Source : Forrester

Part 1:

Introducing the IaaS model : Definition and Objectives

Introduction to Infrastructures as a Service
Another cloud computing service model

- **Infrastructure as a service** (**IaaS**) is a model of Cloud computing intended for companies.
- It is a low-level service that provides access to a virtualized computer equipment.
- The company is discharged from buying and managing the computer equipment.
- Virtual machines are provided on which the company can run the operating system of its choice (if available…).

It is a business model in which the company pays for a service. The computer infrastructure is physically located at the provider's place. This offers a solution to make economies to the customer company's computer department, mostly by transforming investments into renting contracts.

Lifelong Learning Programme

NEBULA

Introduction to infrastructure as a service
Who manages what?

- **The company** manages: Application softwares (executable, parameters, databases,…).

- **The Cloud provider** manages : **Everything else!**
  - Servers, operating systems (versions), virtualization layers, storage, backups, firewall, monitoring, internal networks, access/security model, but also…
  - Power supply, inverters, air conditioned, access control to the equipment, ...

Lifelong Learning Programme

NEBULA

Introduction to infrastructure as a service

In summary

- The **IaaS** model consists of migrating one's physical infrastructures to the "cloud".
- No more need for buying and maintaining your infrastructure on site.
- The company buys/rents the service similar to any other supply.
- This discharges the company from the need of specialized human competencies.

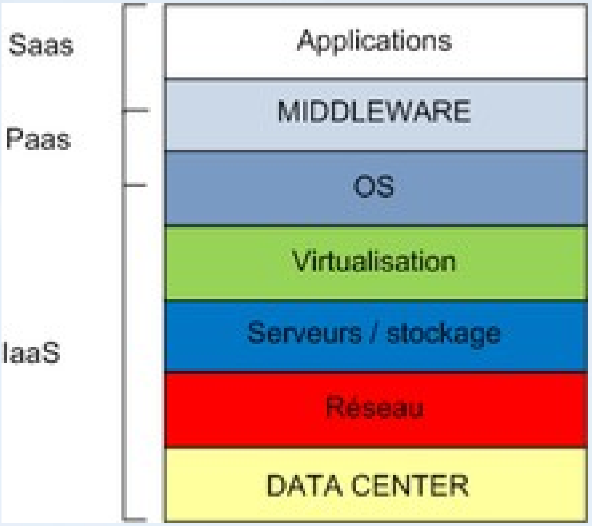

Positionnement IaaS, PaaS, SaaS (Figure Wikipedia)

Note: there exist models of service at an even lower level called MaaS (bare Metal as a Service). A MaaS provider provides physical servers to his clients (without virtualization) as well as necessary tools to deploy a system image, and run/stop the servers remotely.

Lifelong Learning Programme

NEBULA

Other benefits or side effects

- Due to the centralization of physical resources (e.g., servers) in the *datacenter* :
  - Allows rationalizing the hardware exploitation

  (resource pooling, *scalability*, elasticity, hardware maintenance).
  - A unique air conditioned system.
- Create a zone of high energy consumption.

Part 2:

Notions about Green computing and energy efficiency

Green IT

Some notions of *Green Computing*

- Aims at reducing ecological, economical, and societal footprint of Information Technologies (IT).
- Information, broadly speaking, consumes great quantities of electric energy, and contributes to the greenhouse effect. As an example, a recent estimation attributes 2% of greenhouse gas emission to IT.
- Data centers do care. The stakes are:
  - Economical (energetic costs)
  - Public image
- The resources of a datacenter are seldom used at the maximum of the datacenter's capacity. Management/visualization techniques aim at better exploiting the resources.

Power Usage Effectiveness

Some notions of *Green Computing (following)*

- A datacenter's Power Usage Effectiveness (PUE) is the ratio of total consumed energy divided by the energy effictively used by computers (2,5 to 2,7 for the worst, less than 1,3 for the most efficient).
- The datacenter's location has a major impact on green house gas emissions related to electric consumption.
- A datacenter with an excellent PUE but supplied by fossil energy (e.g., USA, China) has little ecological benefits. One might prefer a data center supplied by nuclear energy (e.g., France), or, even better, renewable energy (e.g., Island, Norway).

- Note: a new trend consists of cooling the datacenter with outside ambient air (freecooling).

Energy : 1st limiting factor for large scale systems (Datacenter, Grids, Clouds, Internet)?

- Future exascale/datacenters platforms -> systems from 20 to 100MW (current 1-18 MW)

Top500 : #1 : Tianhe-2 (China) (3M cores, 33 pflops, 17.8 MW, 1.8 Gflops/W);
…
#11 : GSIC center (japan) (74k cores, 2.8 pflops, 1,4 MW, 2 Gflops/W)

Green500 : #1: GSIC center (japan)

- How to build such systems and make them energy sustainable/responsible ?

- Hardware will help (component by component) :

« At a fixed computing load, the amount of battery you need will fall by a factor of two every year and a half. » -- Jonathan Koomey

– *Software must be adapted to be scalable but also more energy efficient*

– *Usage must be energy aware*

Lifelong Learning Programme

NEBULA

Power demand and *Green IT* explosion

- IT – 2-5% of CO_2 emissions / 10% electricity
- Green IT ➔ reducing electrical consumption of IT equipments - CO_2 impact depends on countries
- Focus on usage : fighting un-used/over-provisioned plugged resources
- Researchers are launched : GreenIT scientific events (tracks/issues – greenit-conferences.org)

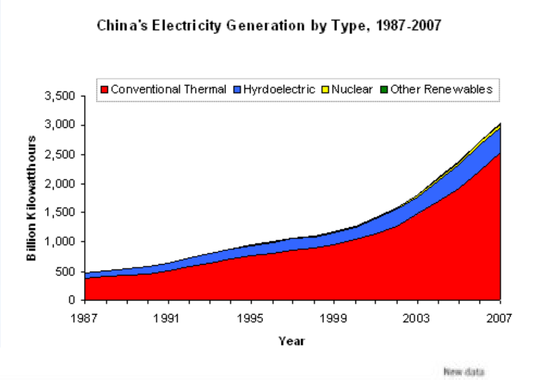

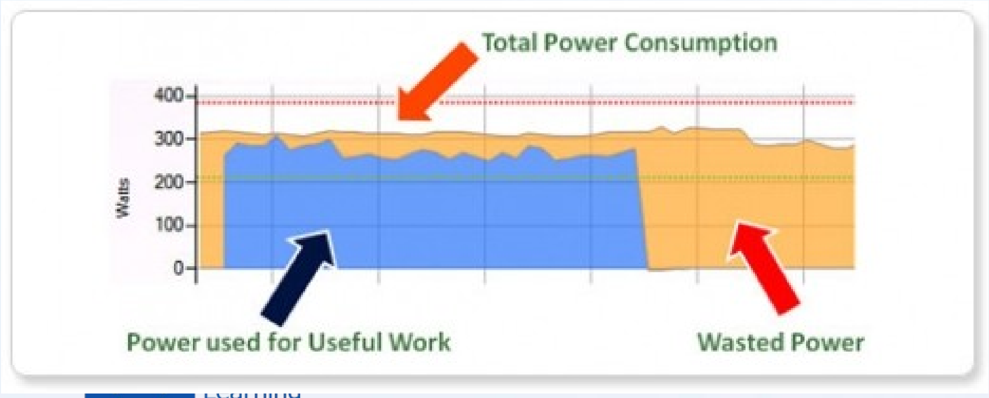

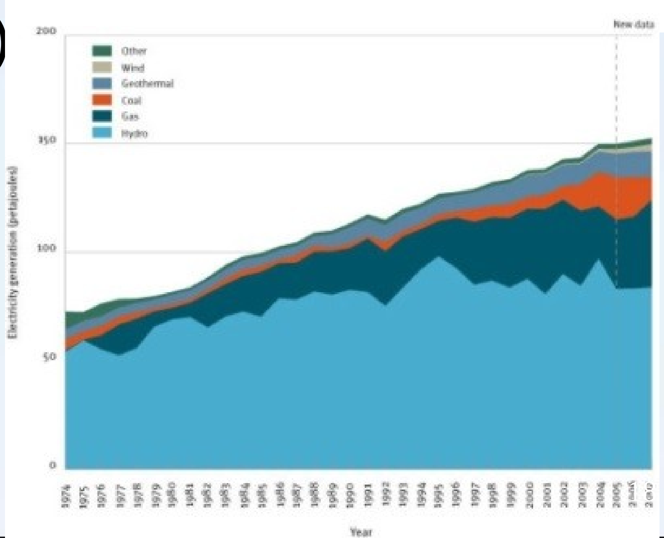

Programme

We are part of climate changing !

Or at least of enormous electricity usage

as IT users/designers

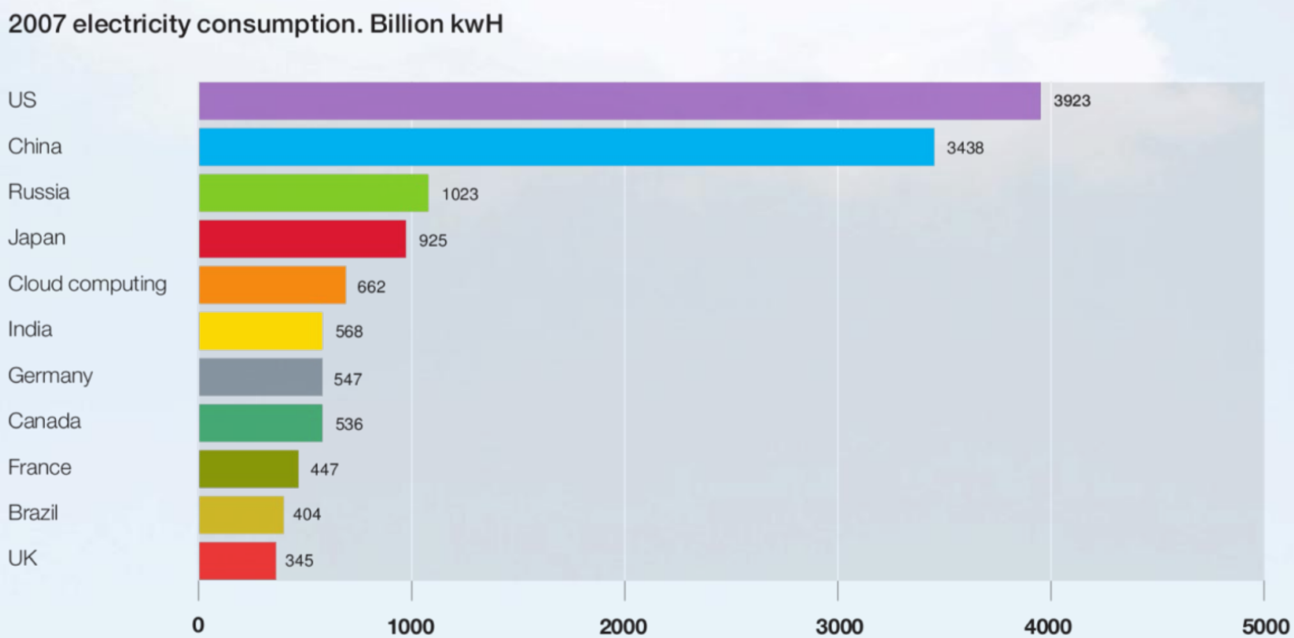

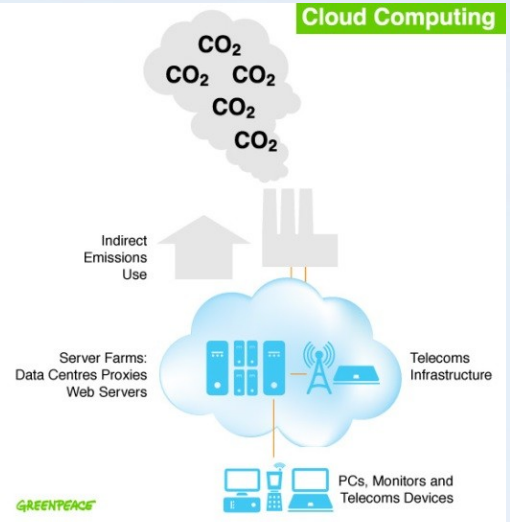

« Greenpeace reports 2010-2011 »

High Performance Datacenter has a bigger role to play

Lifelong Learning Programme

NEBULA

In the Wattmeters jungle…

Eaton, Schleifenbauer, OmegaWatt,
Dell iDrac6, WattsUp, Zimmer LMG450,…

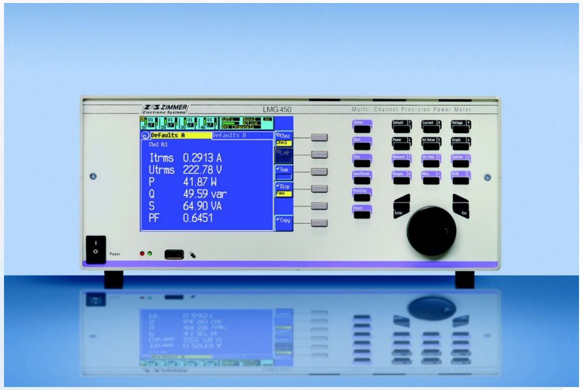

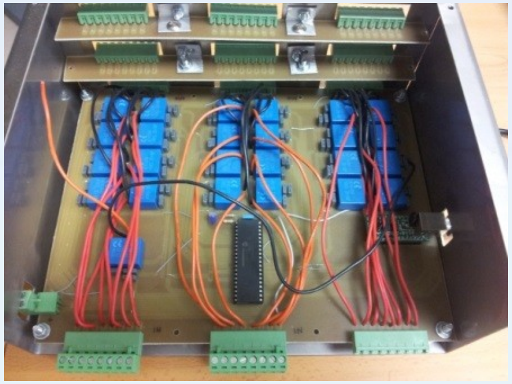

M. Diouri, M. Dolz, O. Glück, L. Lefevre, P. Alonso, S. Catalan, R. Mayo, E. Quintan-Orti. **"Solving some Mysteries in Power Monitoring of Servers: Take Care of your Wattmeters! "**, *EE-LSDS 2013 : Energy Efficiency in Large Scale Distributed Systems conference* , Vienna, Austria, April 22-24, 2013

Probes... with different specifications

(link(s), frequency, resolution, precision, packaging, price per plug,...

Name	Protocol(s)/Link(s)	Frequency	Resolution
Eaton	Serial or SNMP/Ethernet	5s	1 W
Schleifenbauer	SNMP/Ethernet	3s	0.1 W
OmegaWatt	"IrDA"/Serial	1s	0.125 W
Dell iDrac6	IPMI/Ethernet	5s	7 W
WattsUp?	*Legacy*/USB	1s	0.1 W
ZES LMG450	Serial	0.05s	0.01 W

« If you can not measure it, you can not improve it. » *Lord Kelvin*

If external monitoring is not enough : looking inside with dedicated equipment...

Example : Powermon 2 : from RENCI ilab: ilab.renci.org/powermon

Daniel Bedard, Min Yeol Lim, Robert Fowler, and Allan Porterfield

PowerMon: Fine-grained and Integrated Power Monitoring for Commodity Computer Systems.

8 measurement channels

connects directly to ATX supp

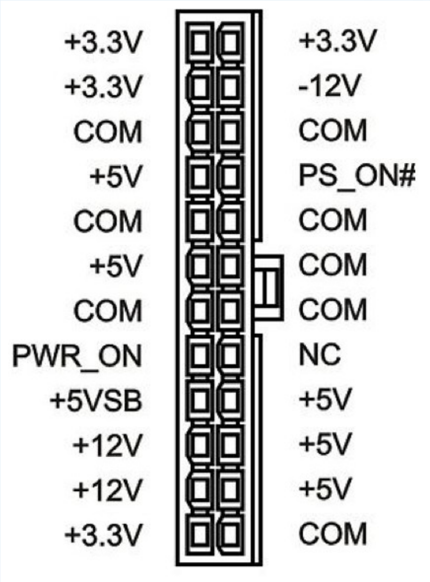

Limits :-(

Difficult to deploy at large scale !

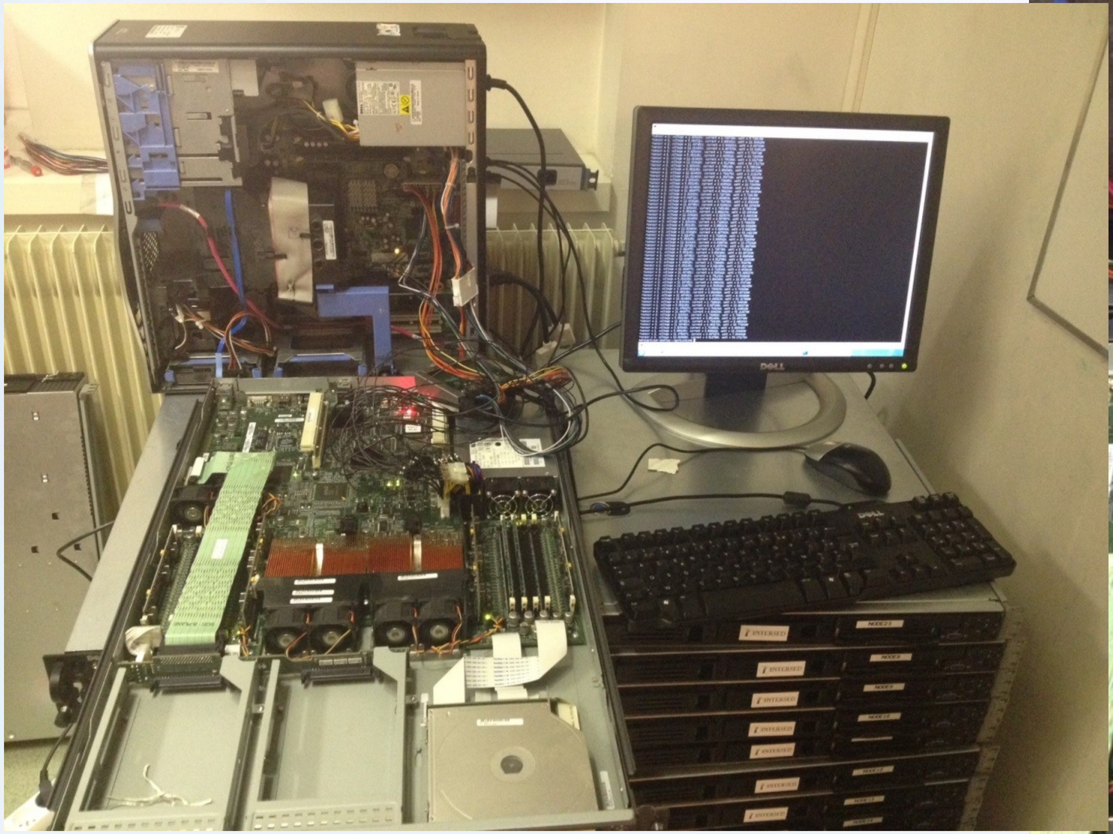

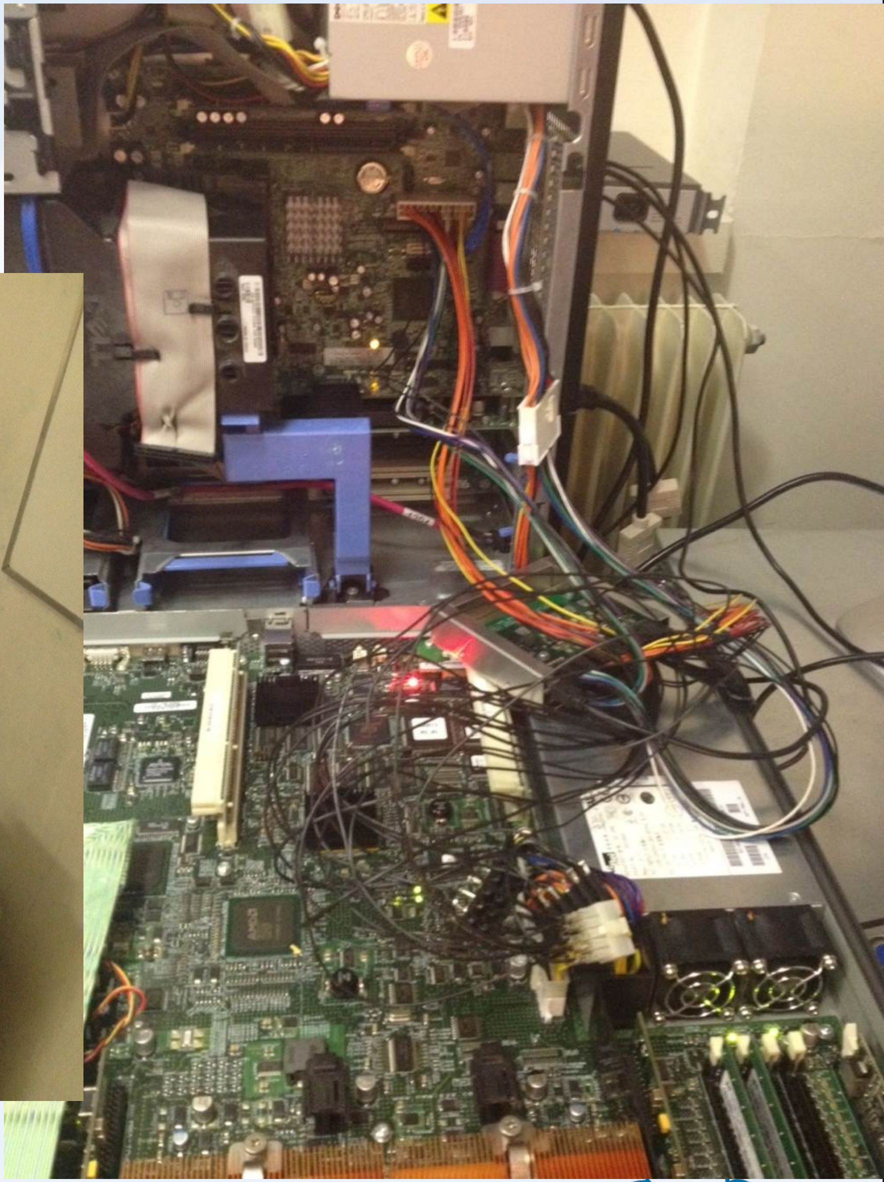

Profiling applications (old measures)

Profiling the energy consumption of applications

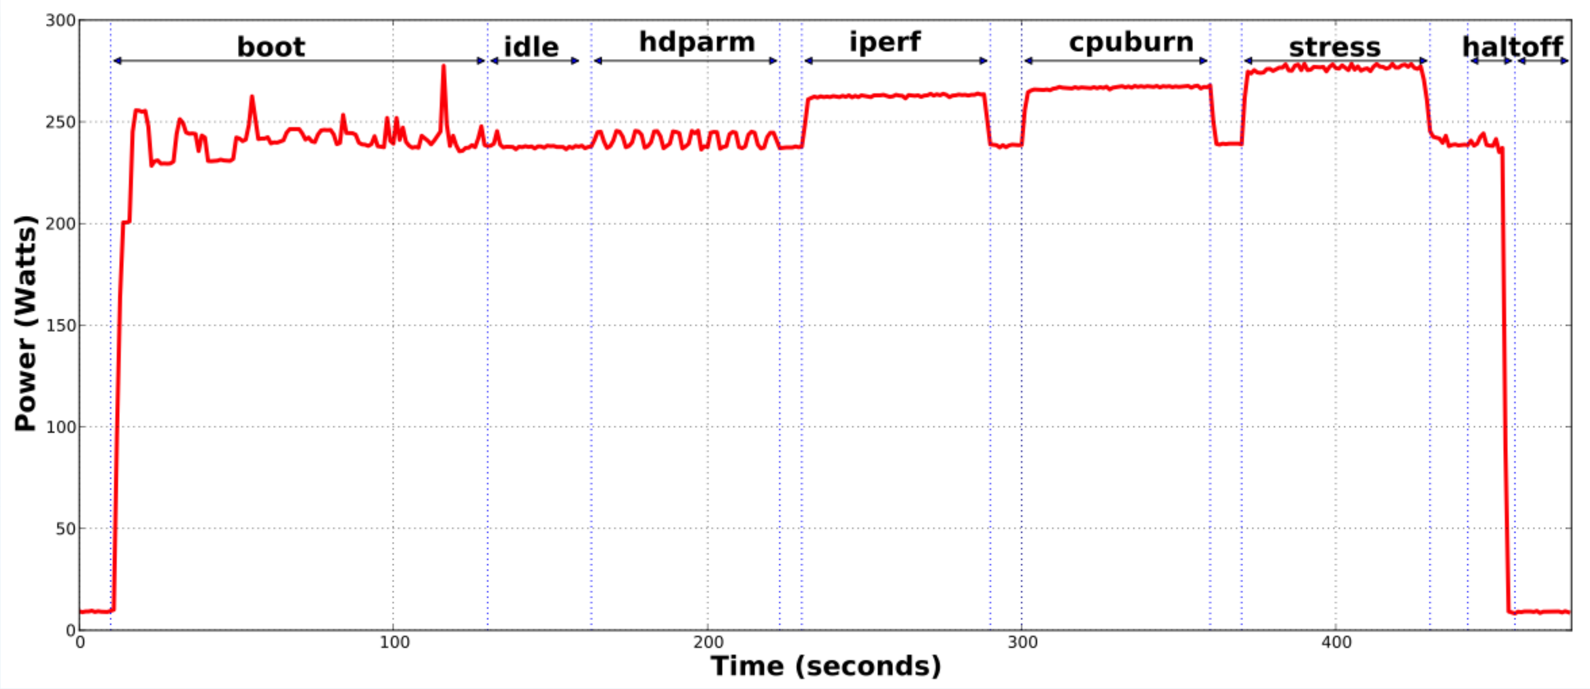

Profiling applications and Node’s power usage

Profiling the energy consumption of applications : more than watts ?

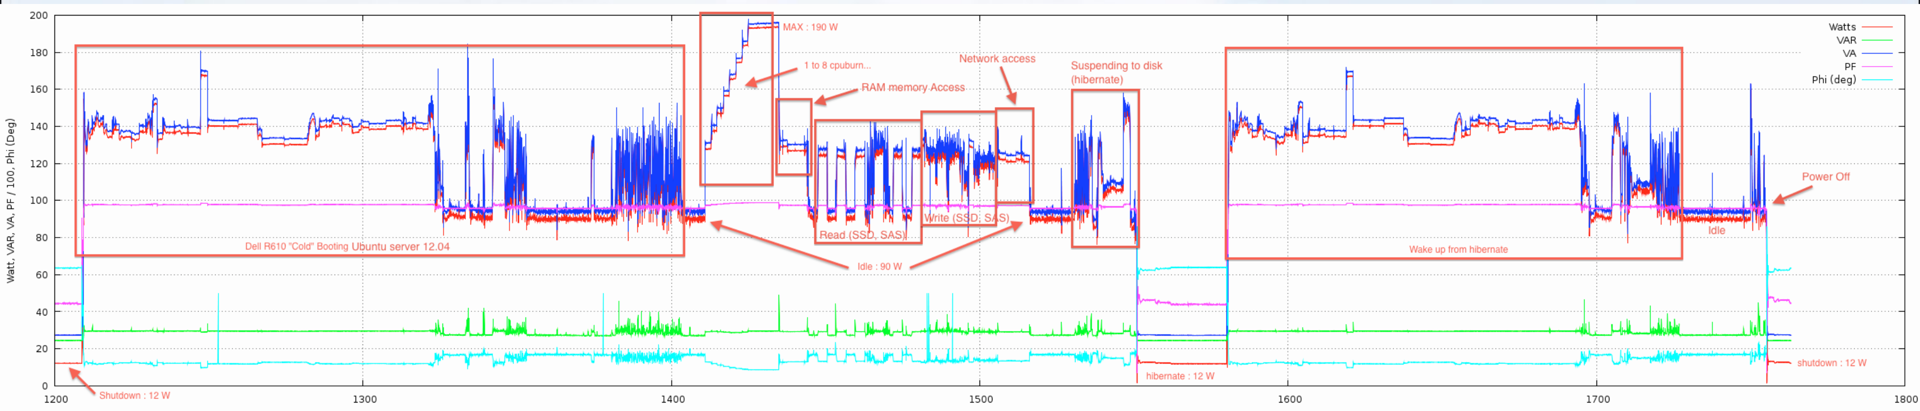

Dell R610 - Zimmer LMG450

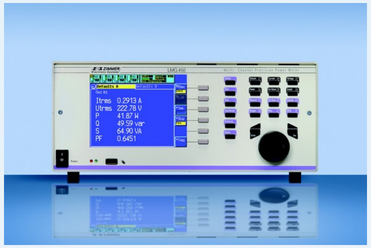

Lifelong Learning Programme

NEBULA

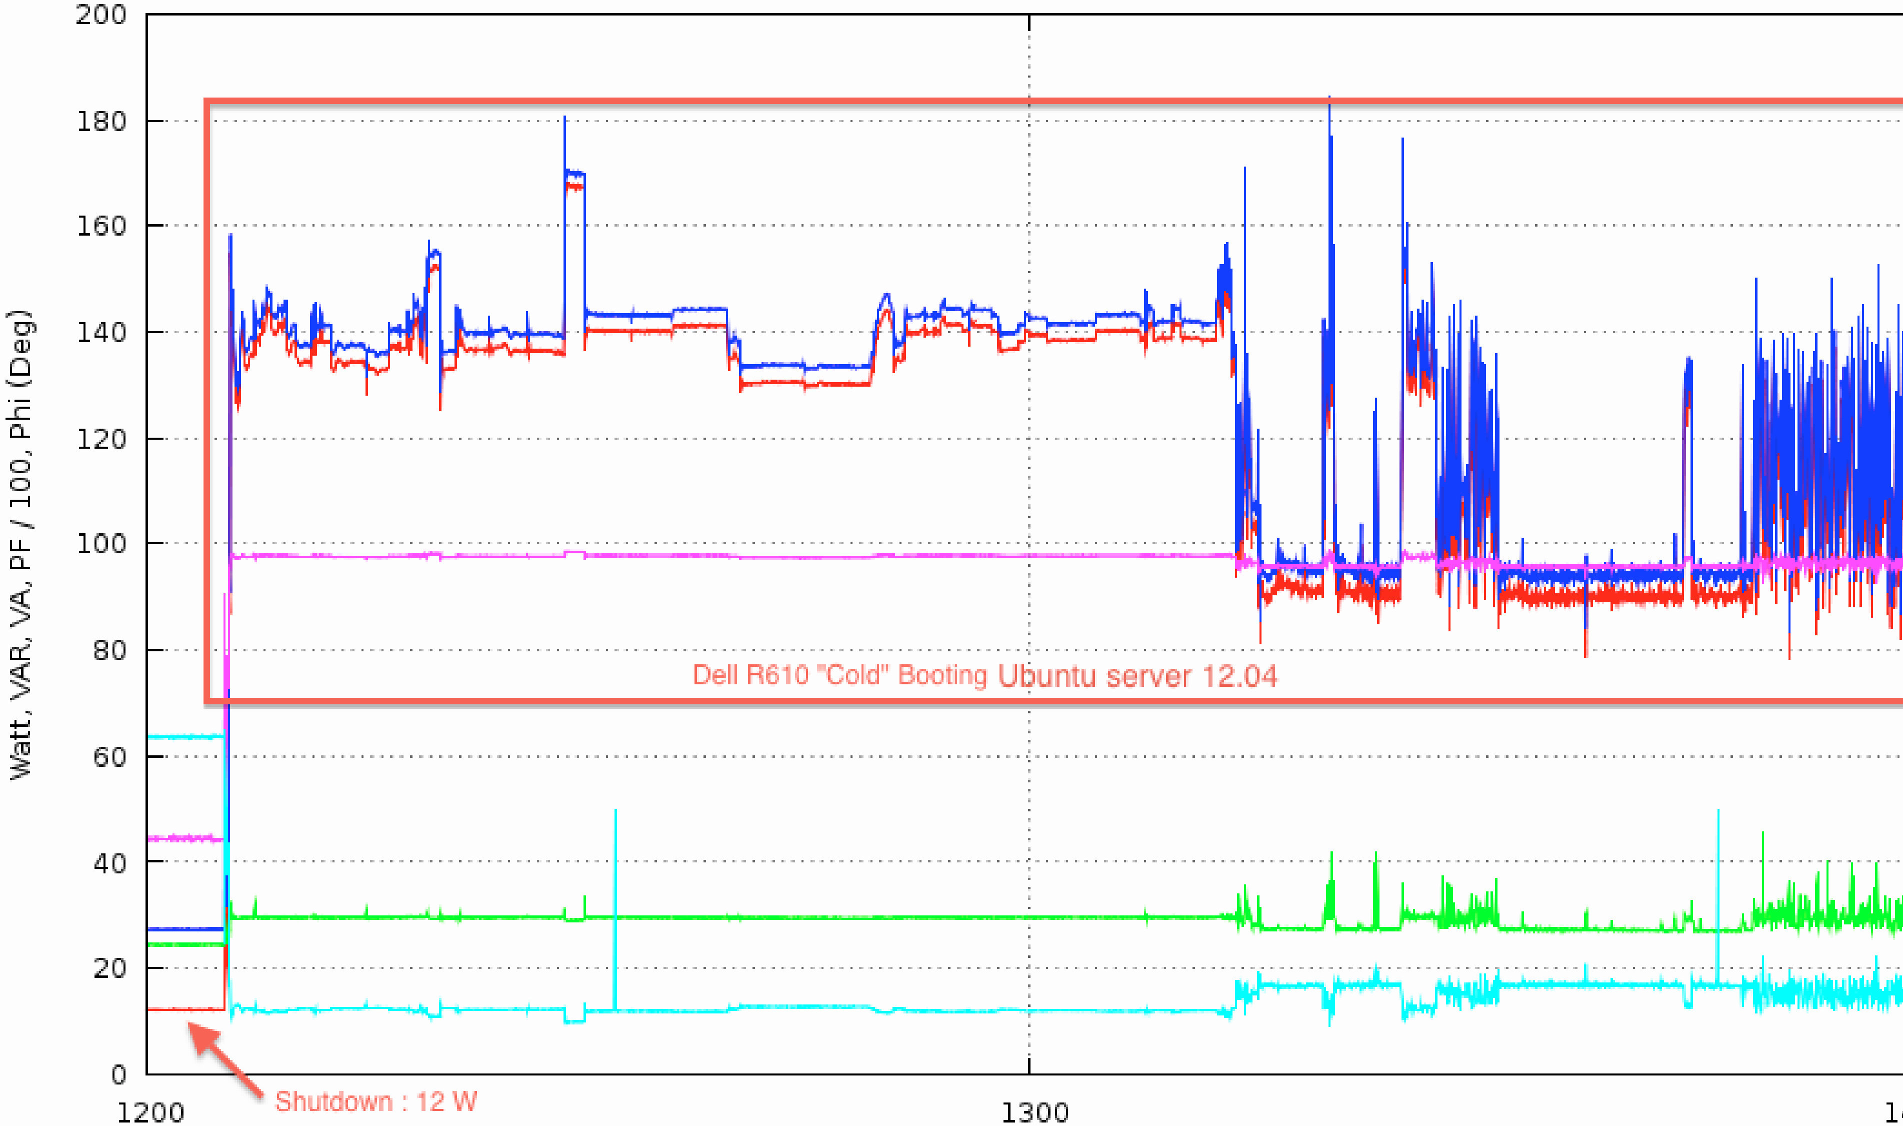

Lifelong Learning Programme

NEBULA

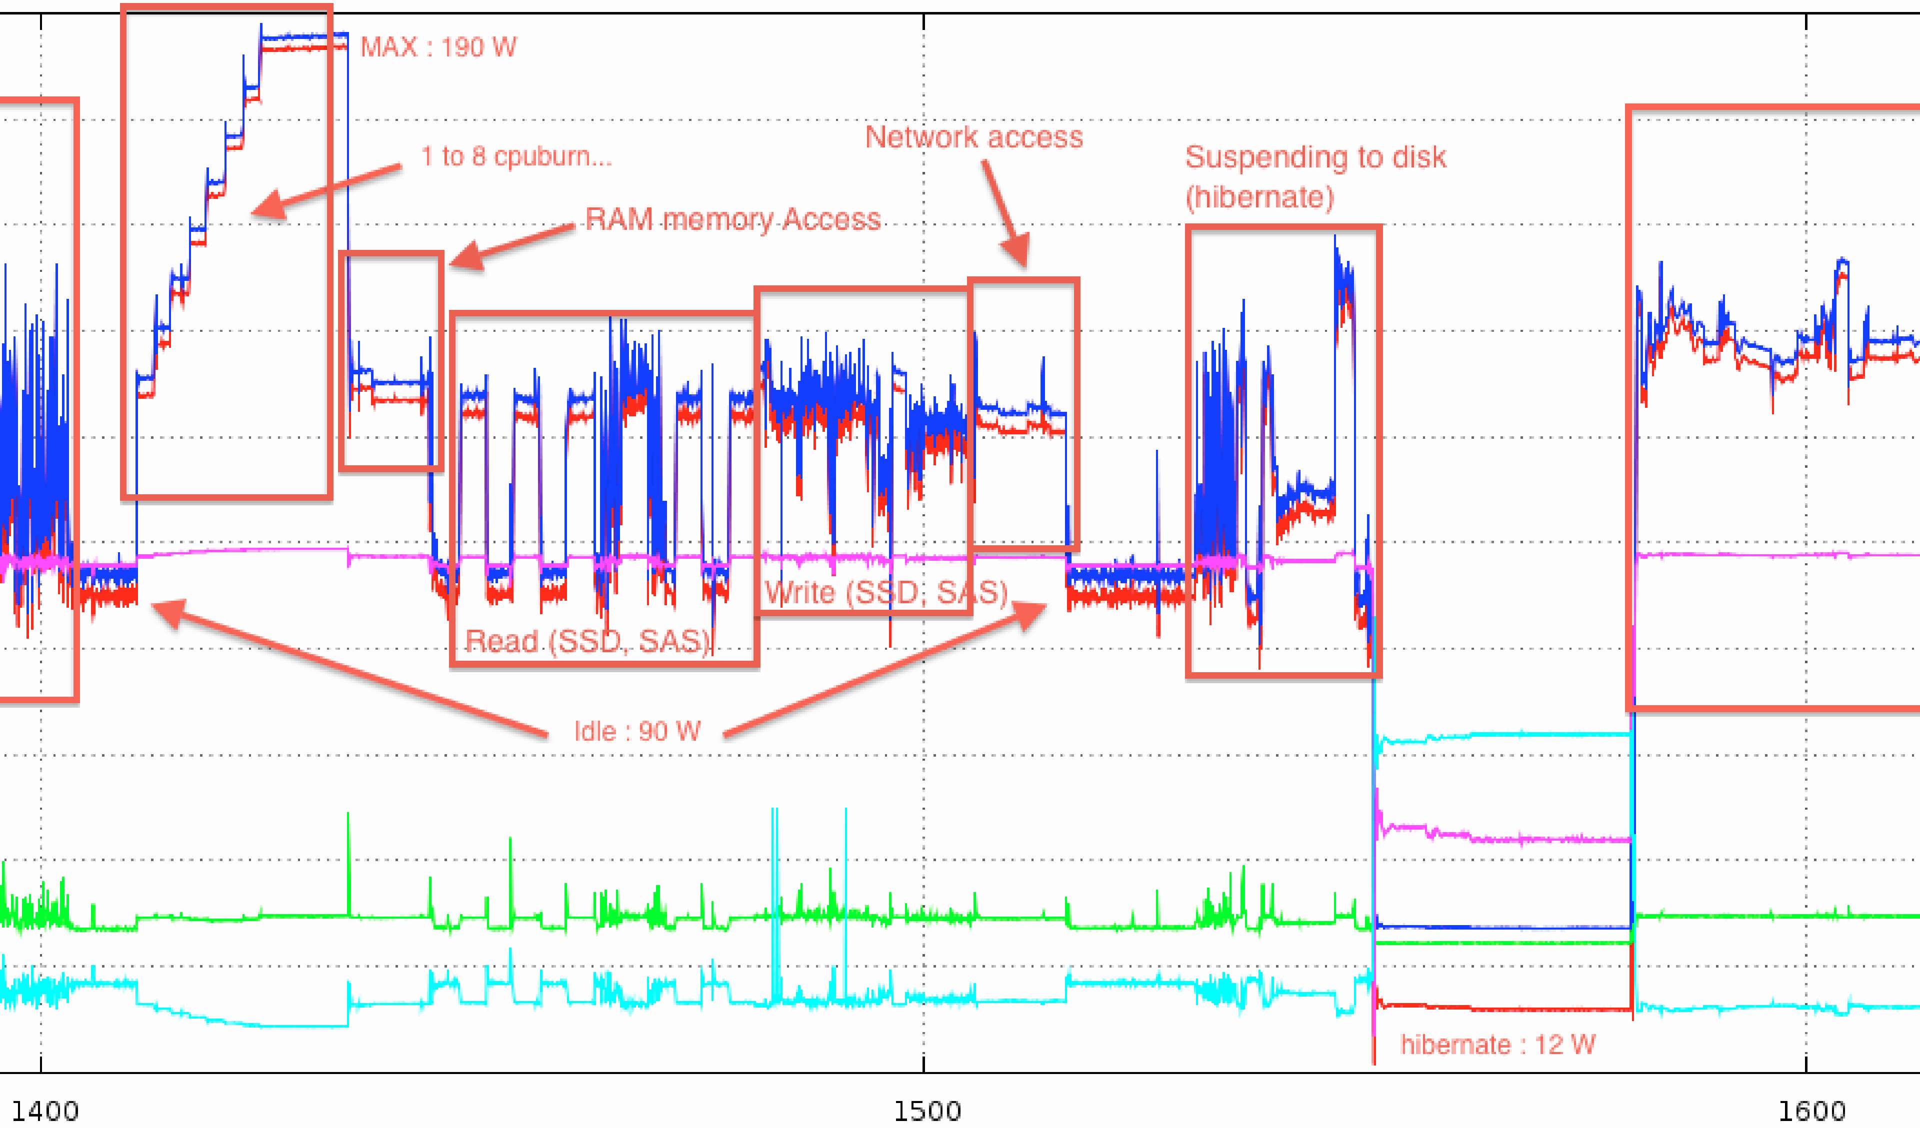
MAX : 190 W
1 to 8 cpuburn...
RAM memory Access
Network access
Suspending to disk (hibernate)
Read (SSD, SAS)
Write (SSD, SAS)
Idle : 90 W
hibernate : 12 W
1400
1500
1600

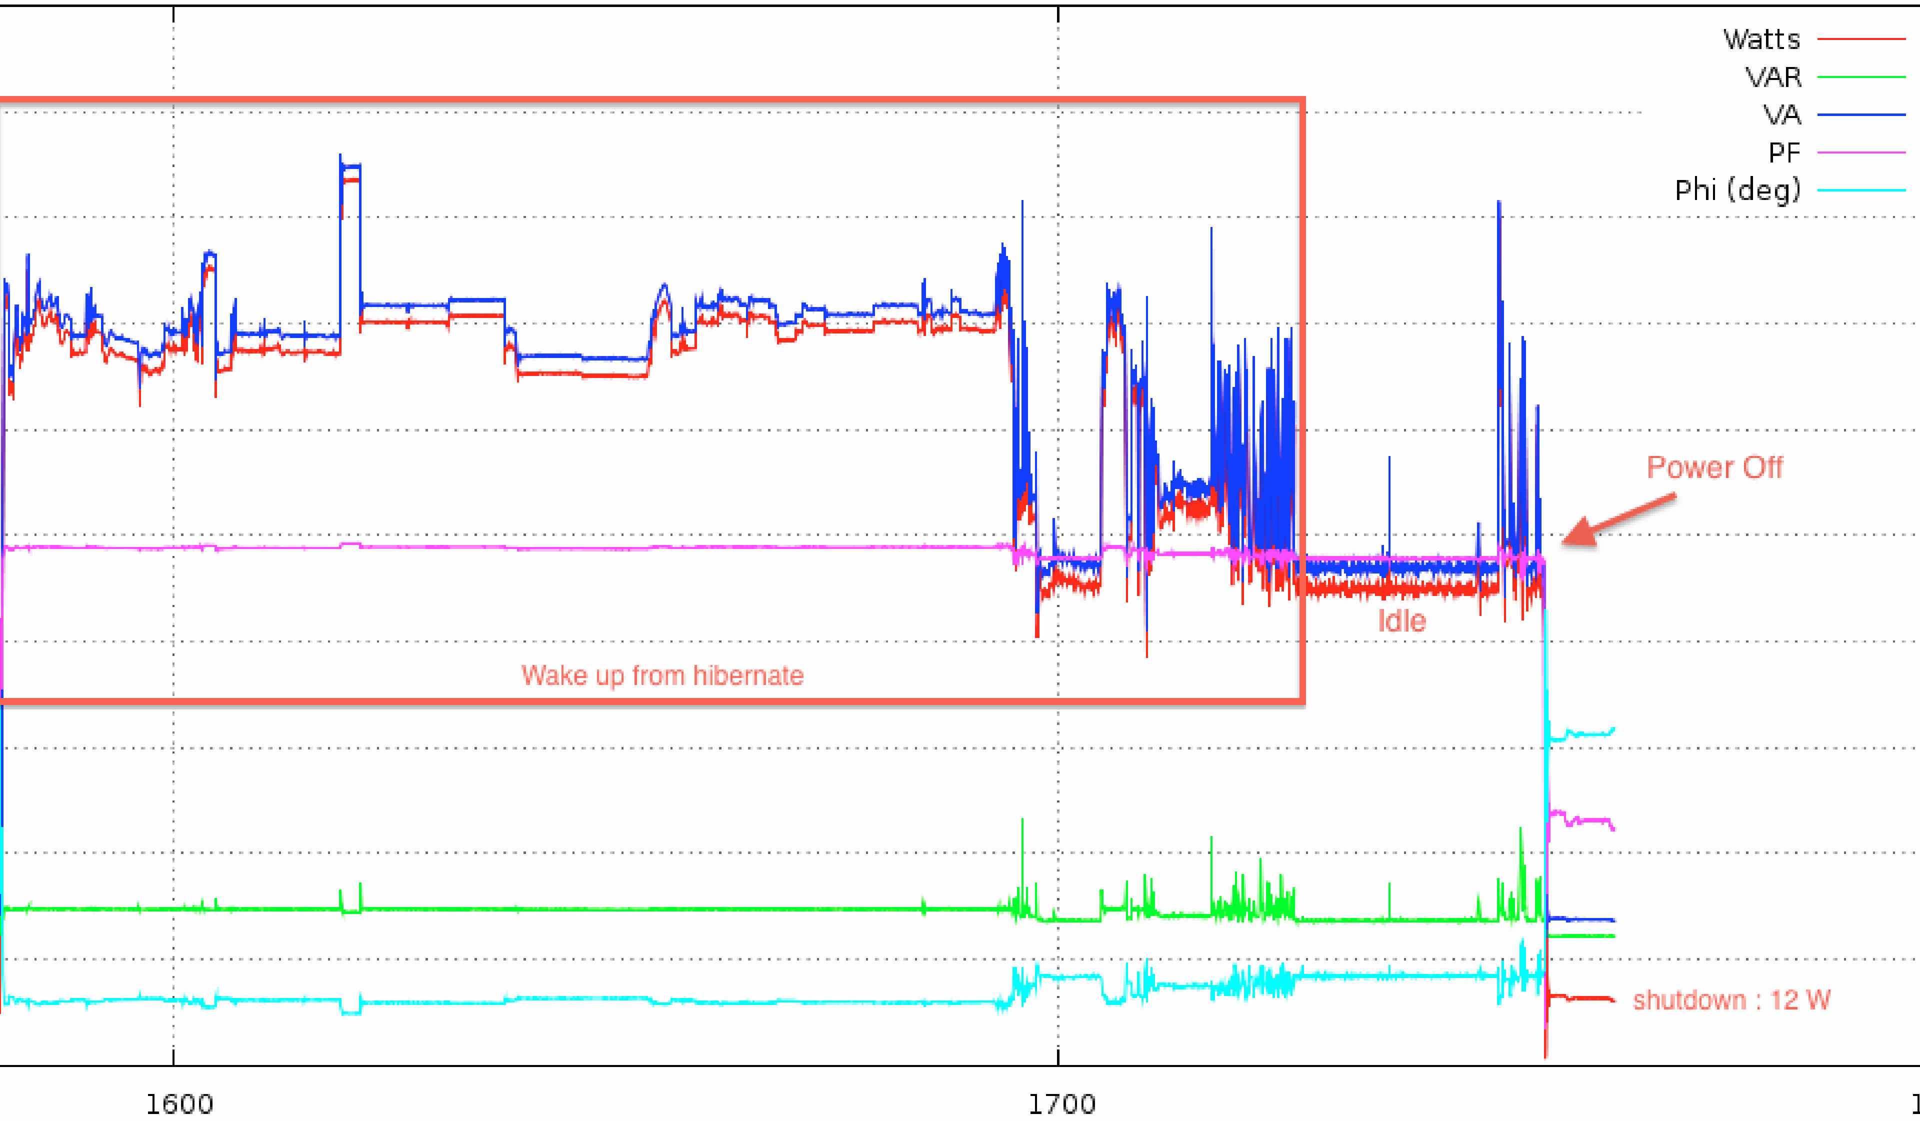
Watts
VAR
VA
PF
Phi (deg)
Wake up from hibernate
Idle
Power Off
shutdown : 12 W
1600
1700

Homogeneity (in energy consumption) does not exist!

- Depends on technology
- Same flops but not same flops per watt
- Idle / static cost
- CPU : main responsible

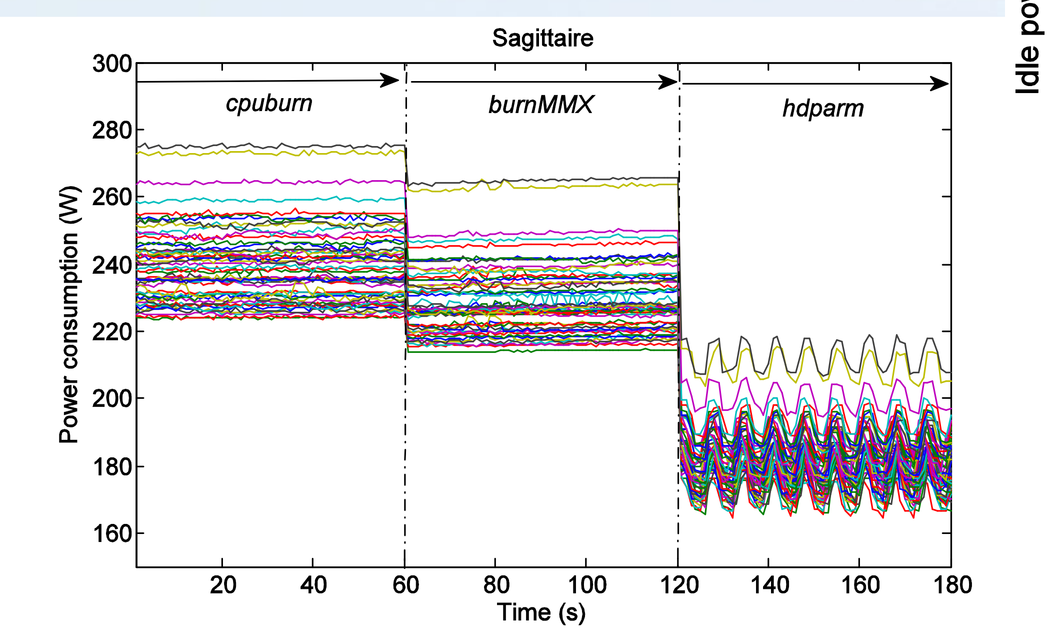

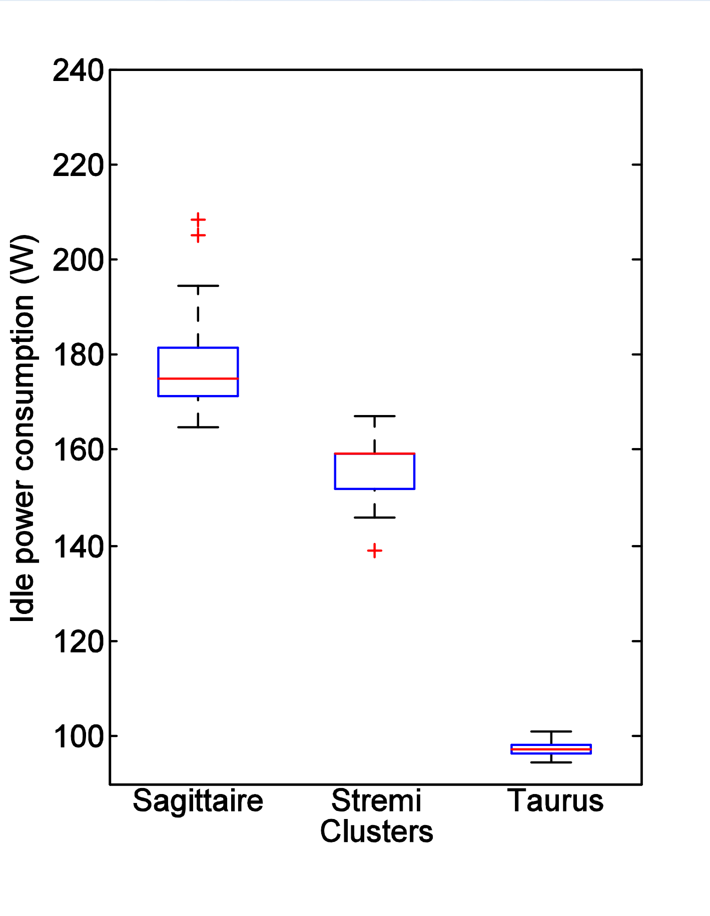

Mohammed el Mehdi Diouri, Olivier Gluck, Laurent Lefevre and Jean-Christophe Mignot. **"Your Cluster is not Power Homogeneous: Take Care when Designing Green Schedulers!"**, *IGCC2013 : International Green Computing Conference*, Arlington, USA, June 27-29,

Part 3:

Architectural and software elements of an Infrastructure as a service

Servers

- Servers are equipment that provides computational resources.
- Computational resources are shared through the virtualization of physical resources principle (CPUs, RAM, Networks, …)
- Virtualization can be achieved through material embedded in processors (Virtualized hardware (VT-x/AMD-V, NPT/EPT)) or purely software.

Storage

- Storage equipment are used to store :
  - System images
  - Virtual machine snapshots
  - Storage drives
  - …

- Depending on usages, there exist different storage solutions (SAN, NAS, block storage, keyed storage,...)

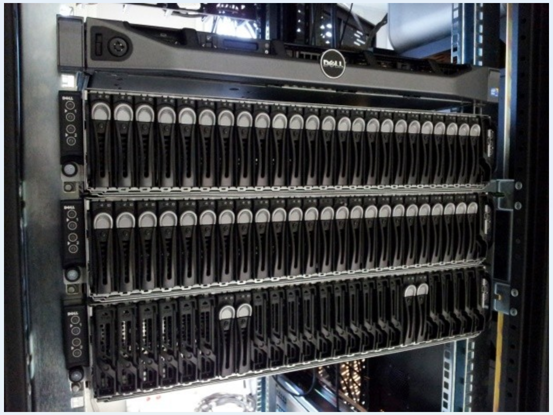

Stockage Cloud

- Stockage block :
  - Accès direct à des raw devices comme /dev/sdb
  - Choix du système de fichiers.
- Stockage Objet :
  - Pousser et retirer des objets dans des containers.
  - Pas de hiérarchie des données.
  - Accès par URL : GET, PUT, DELETE….

Networks

- It’s the most important part!

- Without network (and a suitable speed and latency) there is no cloud computing possible, whatever the model (IaaS, PaaS, SaaS, …).

- Network technology:
  - Gigabit Ethernet,
  - FiberChannel (FC),
  - iSCSI,
  - …

IaaS : Software elements

Technologies and Tools

- Virtualization technologies: virtual machines or containers
  - Linux KVM and Xen, Vmware, OpenVZ, …
- Openstack (see next slide), CloudStack, Eucalyptus, OpenNebula,…
- Management, deployment, and orchestration tools
  - Fuel, Puppet, Chef, Foreman, Docker,…

Openstack

- Characteristics:
  - Started 2010
  - Run by companies like : HP cloud, IBM cloud computing, Rackspace, Intel,… (http://www.openstack.org/user-stories/)
  - OS depending on provider
- pros:
  - Open source IaaS
  - provides REST APIs
  - Compatible with Amazon REST APIs
  - Growing list of services (nova, swift, keystone, horizon,…)
- cons:
  - lack some services (ex : messaging tool (in development))

Part 4:
SOLUTIONS AVAILABLE ON THE MARKET

Some providers on the market

- Cloud Power,
- CloudSystem,
- Desktone,
- Infoserv,
- Provectio,
- DotRiver,
- **Amazon AWS,**
- **Windows Azure,**
- Rackspace,
- OVH
- **Google Compute Engine**
- …
- …
- …

Amazon AWS

Amazon Web Service

- Characteristics:
  - started 2006 by Amazon
  - one of the 3 biggest providers
  - runs Ubuntu, CentOS and Windows Server (more expensive)
- pros:
  - preinstalled PaaS
  - provides REST APIs and SOAP APIs
  - largest range of compute services: mapreduce, object storage, databases (SQL and NoSQL), GPU clusters, …
- cons:
  - complex to use

Windows Azure

- Characteristics:
  - started 2010
  - part of the 3 biggest providers
  - Runs Linux virtual machines and windows Servers
- pros:
  - preinstalled PaaS
  - REST APIs
  - good range of compute service
  - easy administration for windows administrators

Google Compute Engine

- Characteristics:
  - started 2012
  - part of the 3 biggest providers
  - runs Ubuntu and CentOS
- pros:
  - runs on Google's infrastructure : good scalability
  - REST APIs
  - access to Google services : Google big query, Google clouds storage, Google cloud SQL.
- cons:
  - Lack extra services like mapreduce
  - No preinstalled PaaS

Part 5:
Conclusions

Conclusions

- Without a good quality network (reliable end to end, from the customer to the remote infrastructures), the cloud looses all of its interest.

- The hosting location (area in the case of geo-replication) may matter for two reasons:
  - Legislation (e.g., Patriot Act)
  - Energy (e.g., sustainable energy supply)

- The vast majority of IaaS rely on OpenStack software solutions supported by the major constructors.

- Overabundance of IaaS providers, meanwhile only three dominate the market (with small providers delivering more specific services).

- Utilization of IaaS models require more technical competences that PaaS and SaaS models but brings more freedom.

Acknowledgements

- Prof. Frédérique Biennier, INSA de Lyon
- Dr. Olivier Georgeon, Université Claude Bernard Lyon 1.
- NEBULA project members
- Université Claude Bernard Lyon 1, France

Sources and useful links

http://searchcloudcomputing.techtarget.com/definition/Infrastructure-as-a-Service-IaaS

http://www.businessnewsdaily.com/4862-what-is-iaas.html

http://www.ibm.com/developerworks/cloud/library/cl-cloudservices1iaas/

http://www.techrepublic.com/blog/the-enterprise-cloud/top-cloud-iaas-providers-compared/

http://www.theresearchpedia.com/research-articles/top-benefits-of-iaas

http://www.elastichosts.com/cloud-hosting/infrastructure/

http://fr.wikipedia.org/wiki/Infrastructure_as_a_service

https://www.youtube.com/watch?v=GCGLYMeh75Y

http://www.cloud-experience.fr/le-cloud-est-il-vert/

http://www.tomsitpro.com/articles/iaas-providers,1-1560.html

Proposed exercise work

- Work 1 : Deployment of system image on an Openstack platform
- Work 2: Administering an Openstack platform
- Work 3: Creation and management of storage volume.