Des Grilles vertes aux nuages verts : vers des systèmes efficaces en énergie

Laurent Lefèvre
INRIA – Université de Lyon - LIP

A.-C. Orgerie, M. Dias, J.-P. Gelas

laurent.lefevre@inria.fr

Grids2Clouds Day, Lyon, December 13, 2010

ENS DE LYON

INRIA RHÔNE-ALPES

Lyon 1

Lip

cnrs

Energy : 1st challenge for large scale systems ?

- Future exascale platforms -> systems from 20 to 100MW (current 4-6 MW)
- How to build such systems and make them energy sustainable/responsible ?
  - Hardware can help (component by component)
  - Software must be adapted to be scalable but also more energy efficient
  - Usage must be energy aware

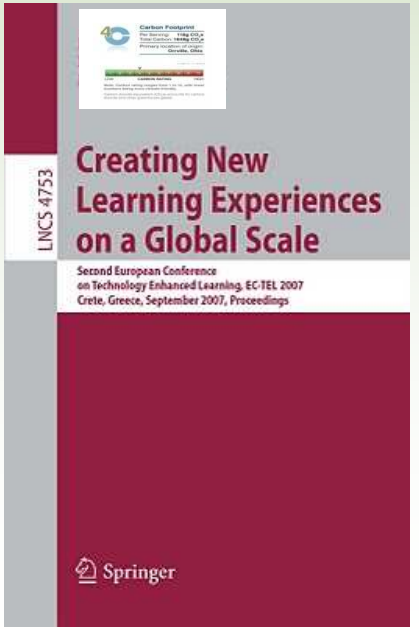

Power demand and Green IT explosion

- IT – 2-5% of CO2 emissions
- Green It ➔ reducing electrical consumption of IT equipments - CO2 impact depends on countries
- Focus on usage : fighting un-used/over-provisioned plugged resources
- Problem : grey energy (ecoinfo.org)
- GreenIT scientific events (12/2010 – 15/2011 + tracks/issues – greenit-conferences.org)

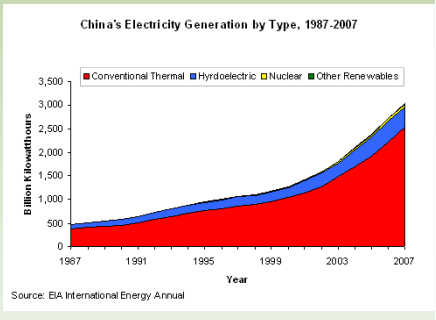

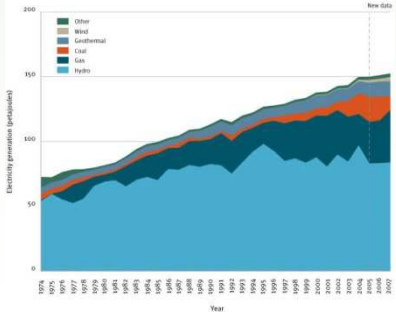

Towards Energy Aware Large Scale Systems : open questions

How to decrease the energy consumption of Grids & Clouds without impacting the performances?

- How to understand and to analyze the usage and energy consumption of large scale platforms?
- How to monitor lively such usage from pico to large scale views?
- How to design energy aware software frameworks ?
- How to help users to express theirs Green concerns and to express tradeoffs between performance and energy efficiency ?

Green-IT Leverages

- **Shutdown** : reducing the amount of powered unused resources
- **Slowdown** : adapting the speed of resources to real usage
- **Optimizing** : improving hardware and software for energy reduction purpose
- **Coordinating :** using large scale approaches to enhance green leverages

Explosion of initiatives

For each domain

- **Data centers/HPC : Green500 (1Gflops/W –> 20 Mflops/W), EU CoC**
- **Grids : The Green Grid (metrics) / Open Grid Forum**
- **Storage : SNIA (ON/Off disks, metrics)**
- **Networks : Green Touch (x1000 factor) / EEE (LPI)**

Methodology

- Proposing a generic energy aware model able to be derivated onto different scenario (Grids, Clouds, Networks)
- Designing software solutions for infrastructures
- Simulating and Validating at medium and large scale

General approach

Everything is a resource reservation :

- Reserving CPU in HPC and Grids
- Reserving Virtual machines in Clouds
- Reserving Bandwidth in large transport of data

- Leverages:
  - Finding and powering the optimal number of resources in front of needs of applications
    - HPC and Grids : switching on/off resources
    - Clouds : migrating -> switching on/off VMs
    - Networks : lighting or switching off paths, interfaces, links, routers
  - Adapting « speed » (and consumption) to the need of appliciations/users
    - HPC, Grids : dvfs
    - Clouds : tuning, capping
    - Networks : adaptive link rate, LPI

The ERIDIS approach

- Energy-efficient Reservation Infrastructure for large-scale Distributed Systems
- Systems Collecting and exposing : usage, energy profiling of applications and infrastructures
- Expressing and Proposing : to deal with tradeoffs between perf and energy, Green Policies
- Agregating resources reservations and usage
- Enforcing Green leverages : shutdown or adapt performance
- Predicting usage of infrastructures

ERIDIS

Multi-View Understanding of Large Scale Systems Usage

Monitoring and Analyzing Energy Information

Designing Energy Efficient Frameworks

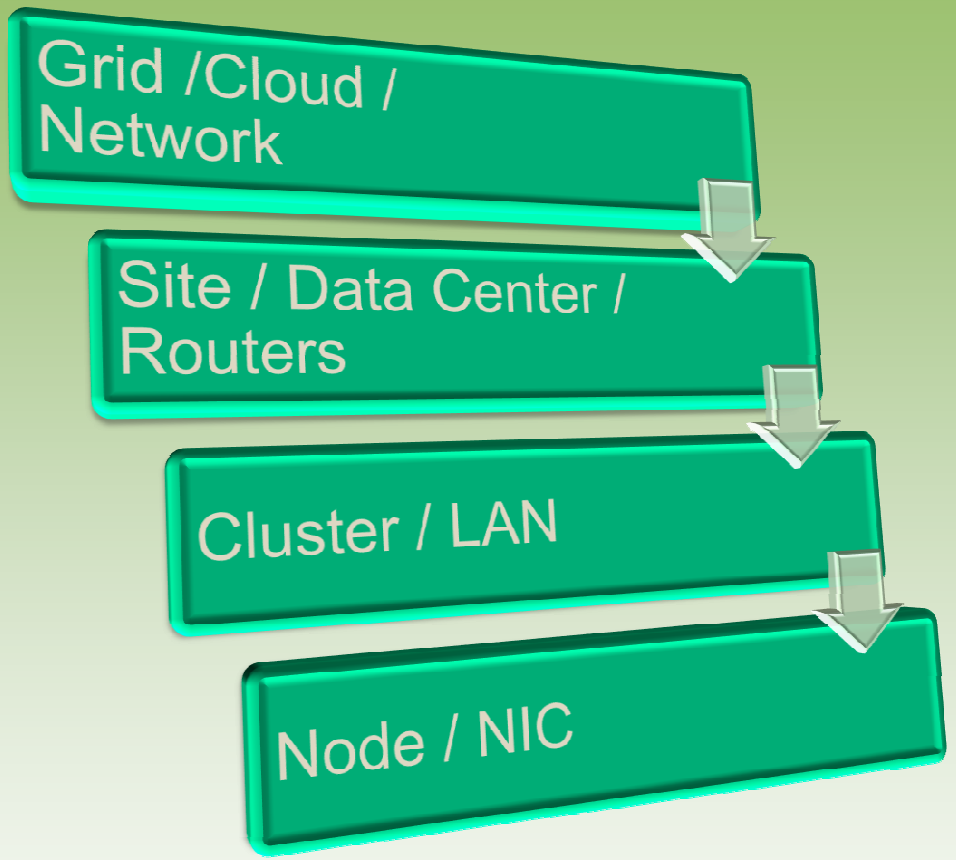

ERIDIS

Multi-View Understanding of Large Scale Systems Usage

Monitoring and Analyzing Energy Information

Designing Energy Efficient Frameworks

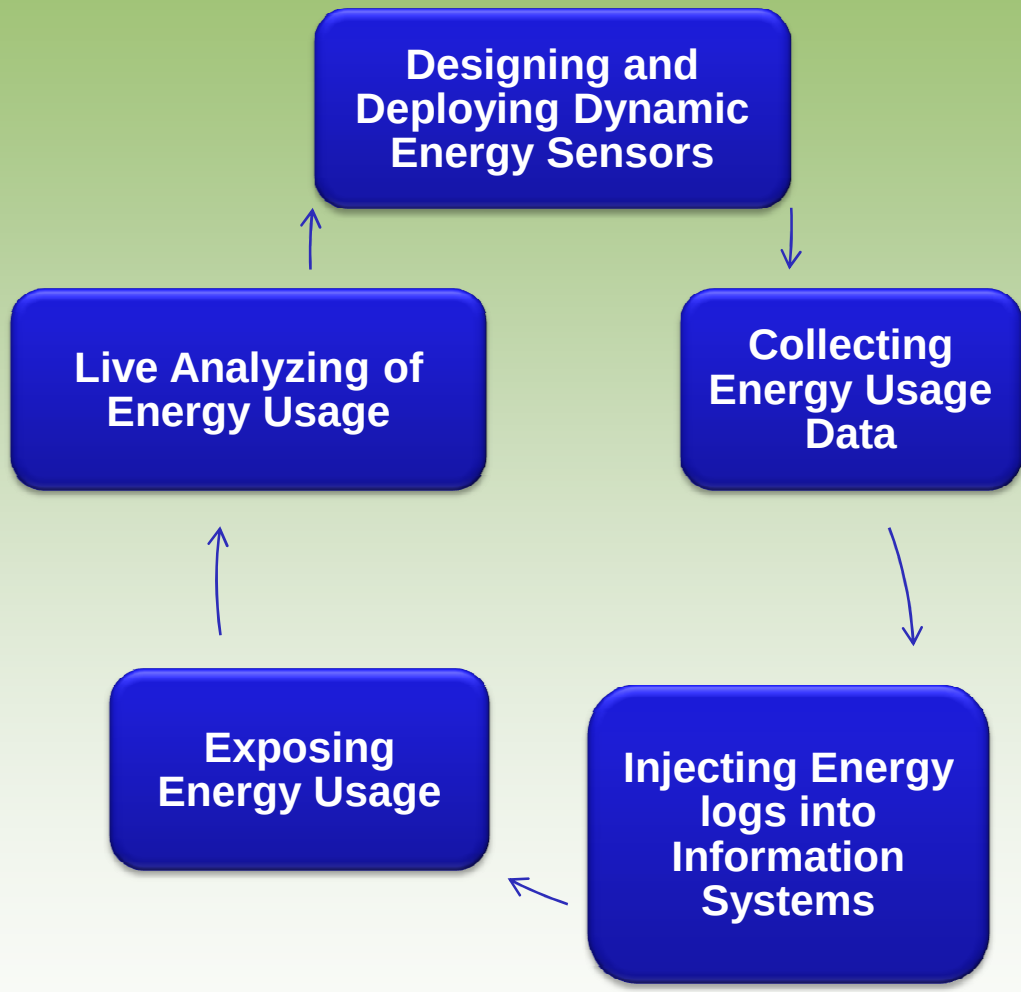

ERIDIS

Multi-View Understanding of Large Scale Systems Usage

Monitoring and Analyzing Energy Information

Designing Energy Efficient Frameworks

Prediction Systems

Adapted Schedulers and Resource Managers

Node and Services Virtualization

Delegated Trust and Network Presence

Green Policies Support

Node Energy Controlers

The ERIDIS Framework for Grids

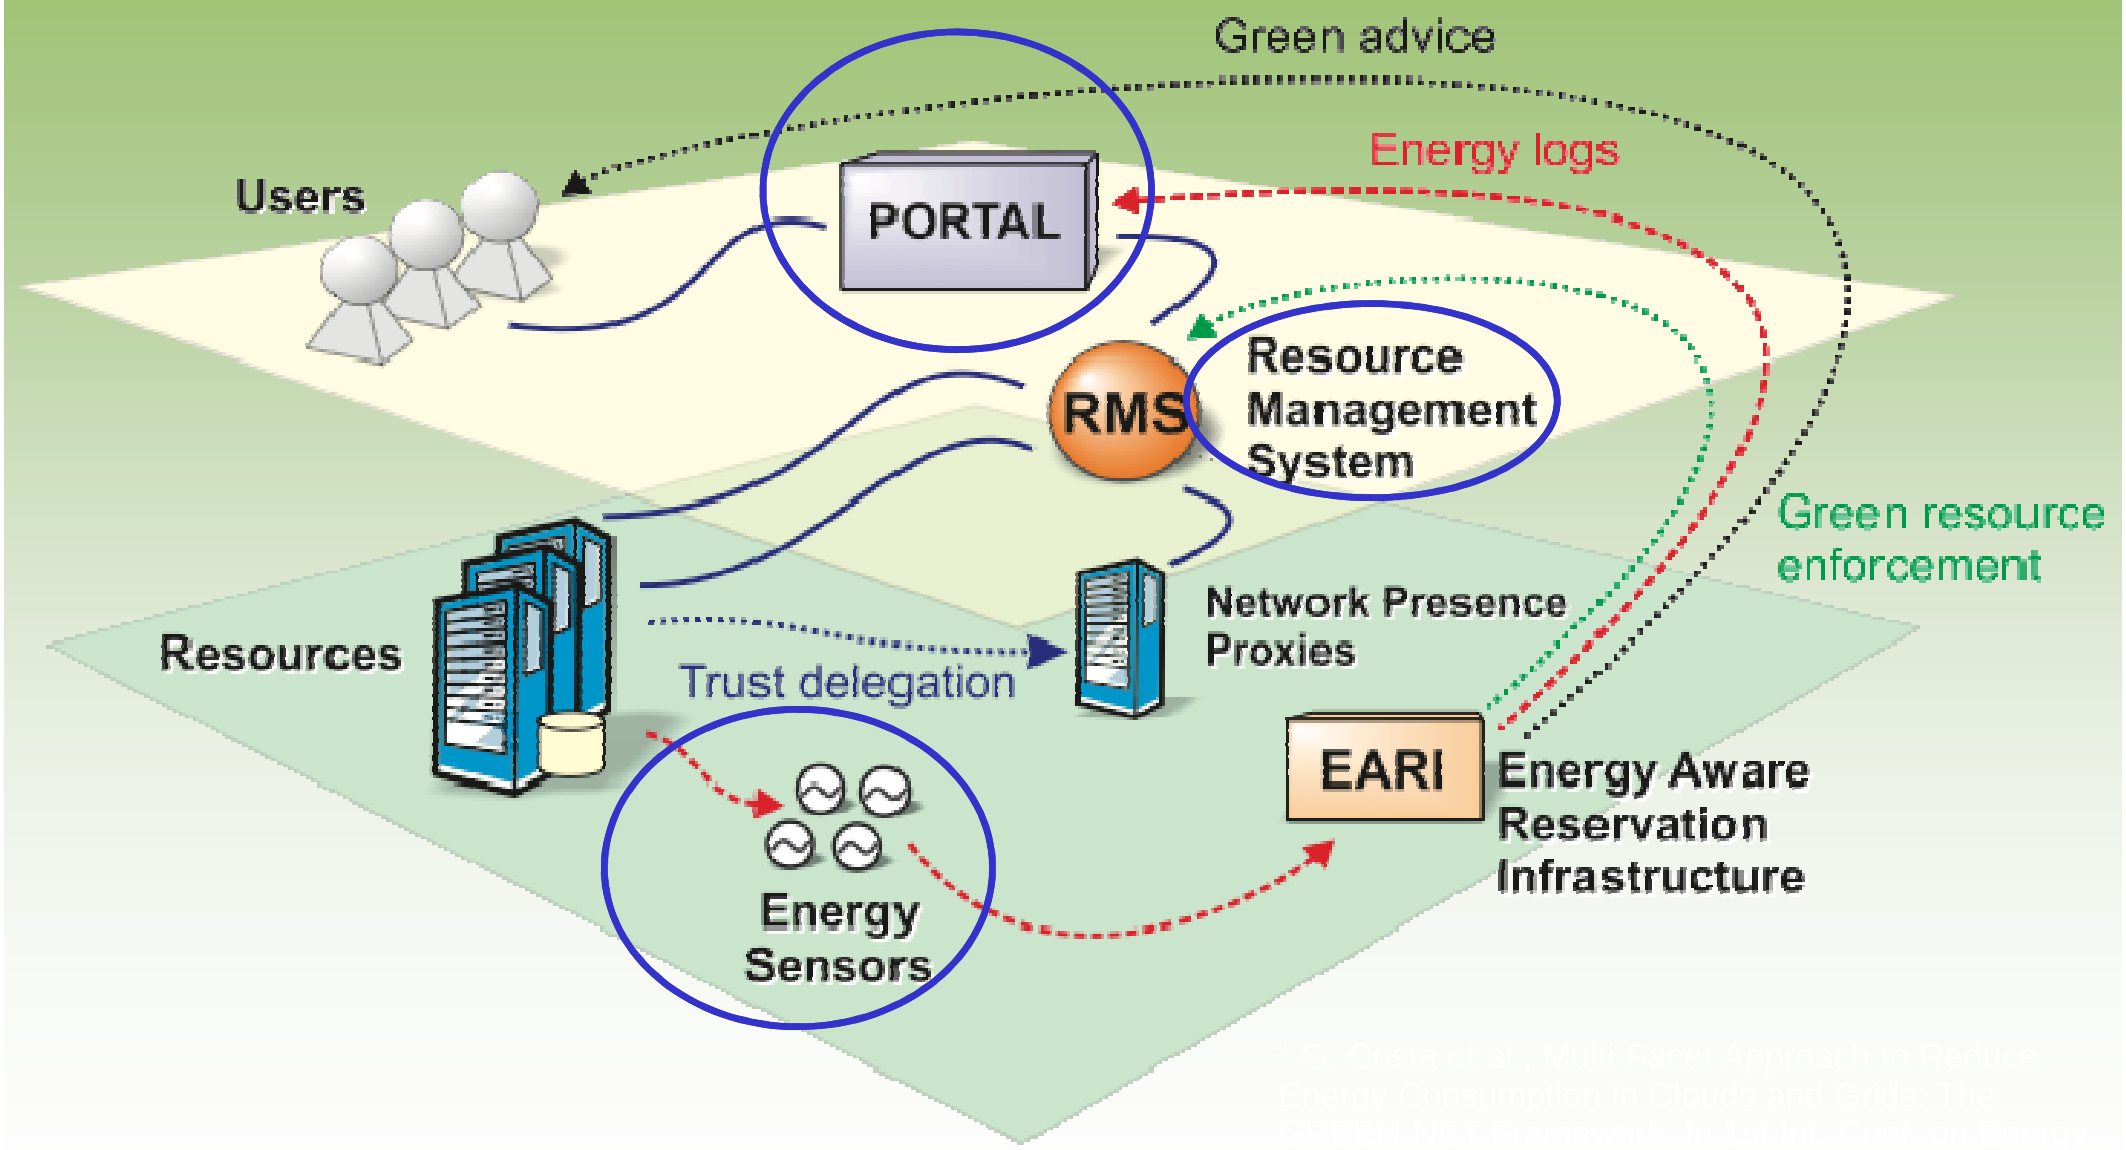

* G. Costa et al., Multi-Facet Approach to Reduce Energy Consumption in Clouds and Grids: The GreenNet Framework, the 1st Int. Conf. on Energy-Efficient Computing and Networking

1st focus : Collecting and exposing

- Grid'5000
  - French experimental testbed
  - 5000 cores
  - 10 sites

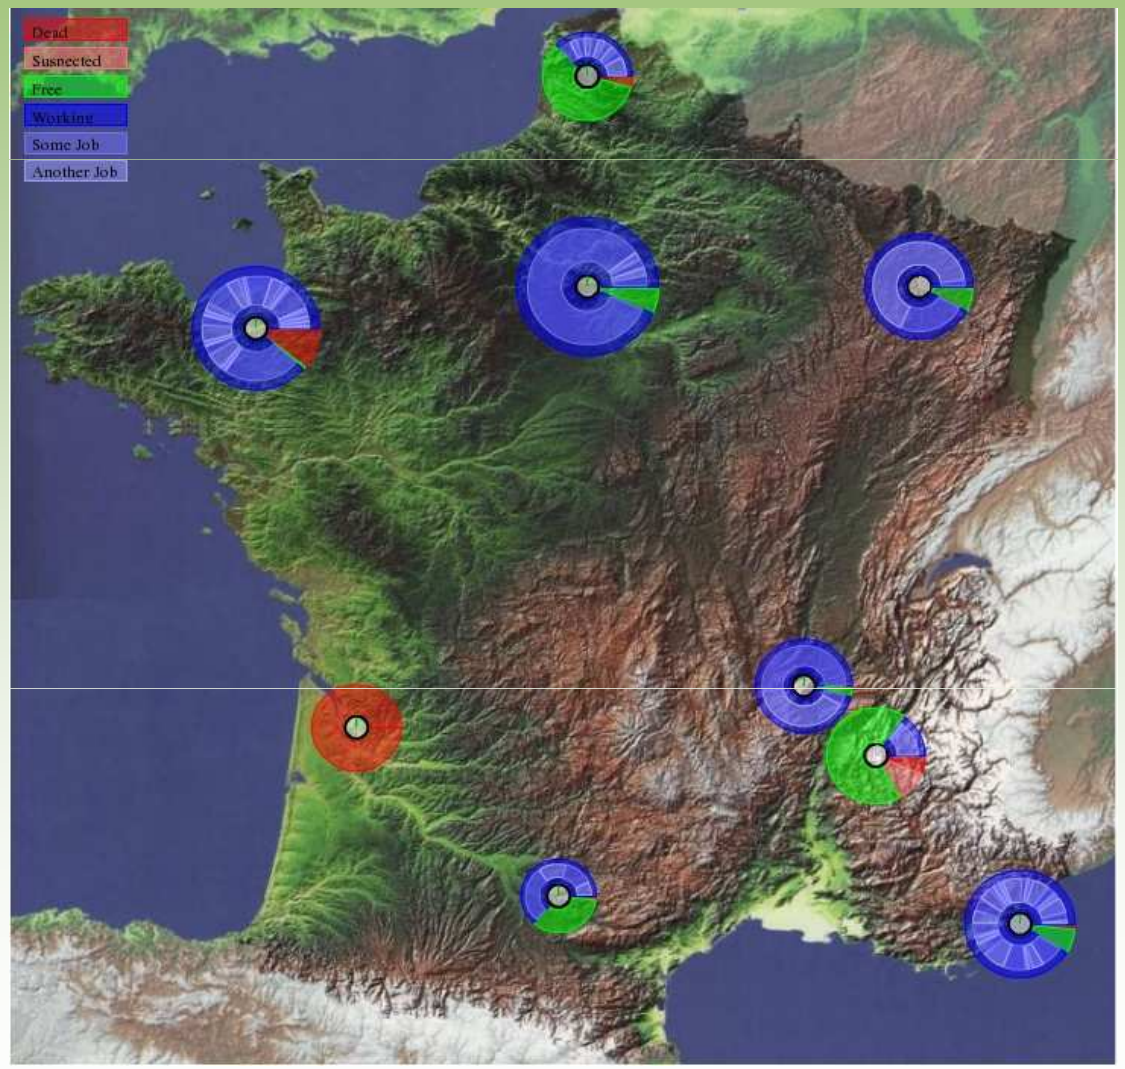

The Green Grid5000

Energy sensors

- 6 or 48 ports wattmeters boxes / PDUs
- Deployed on three sites of Grid'5000 (Lyon, Grenoble, Toulouse)
- Library for **interfacing** with **energy sensors**
- **Client-side** applications to **obtain** and **store** the energy consumption data

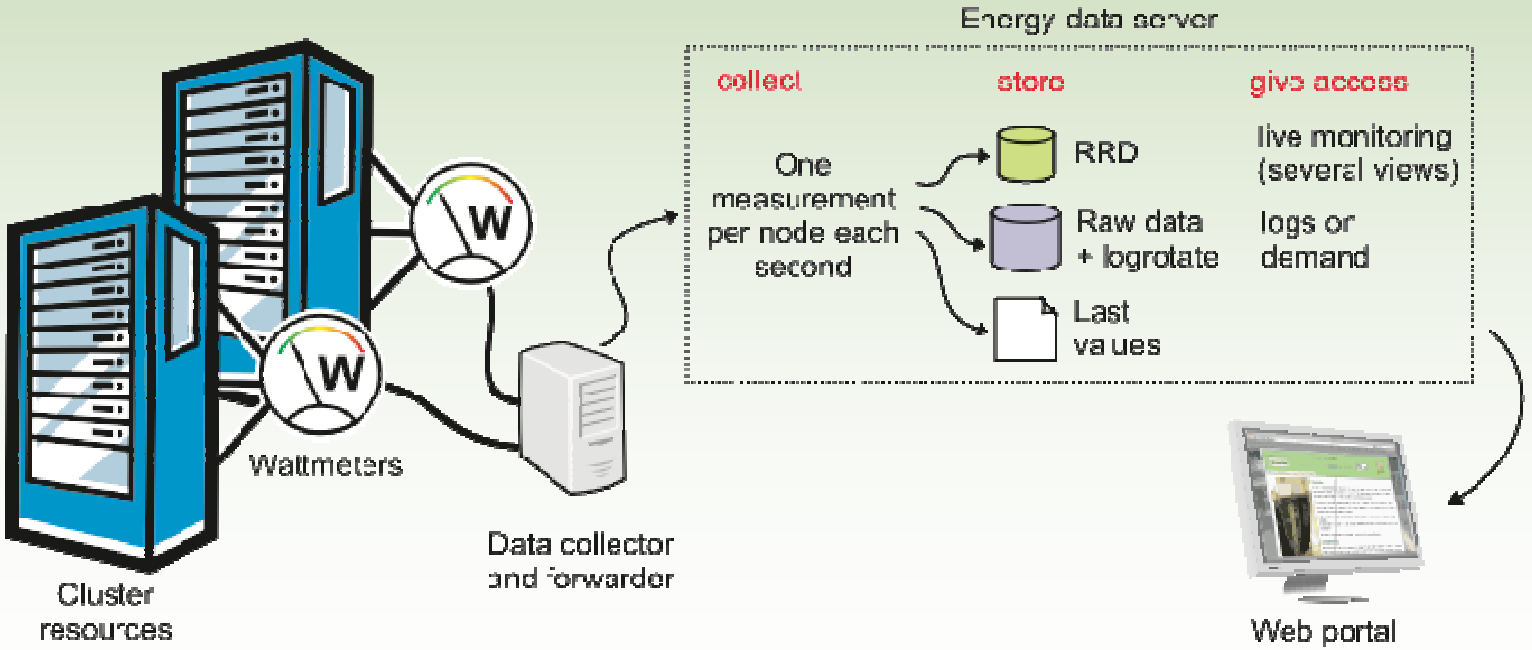

Electrical consumption / Usage

Periodicity of **energy measurements**:

- Application oriented
- One measurement per **second** for each equipment

*

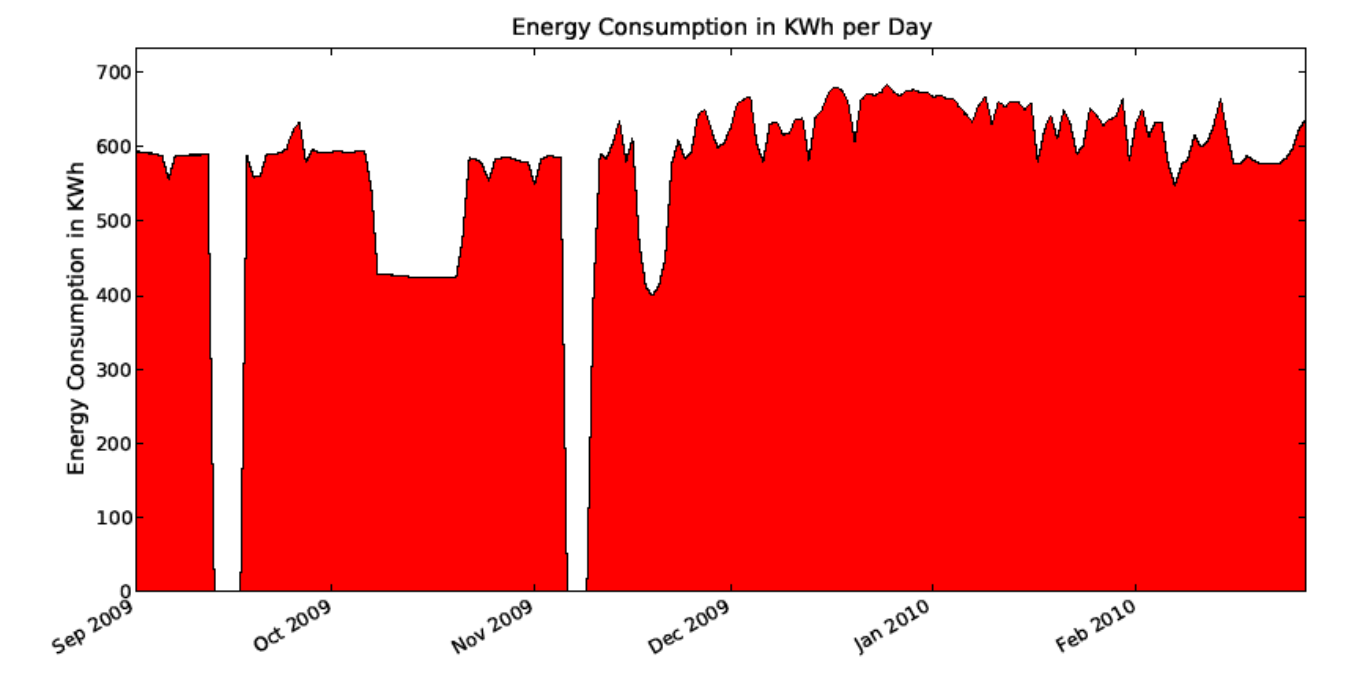

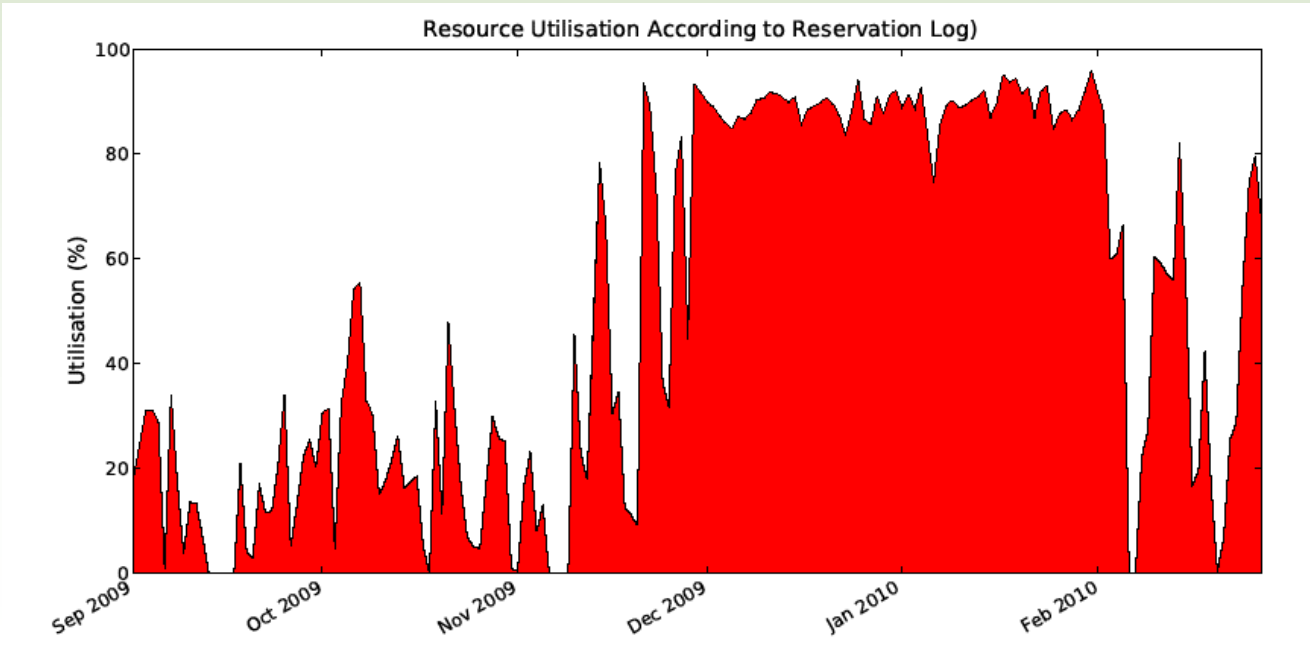

Example I : Profiling applications

Profiling the energy consumption of applications

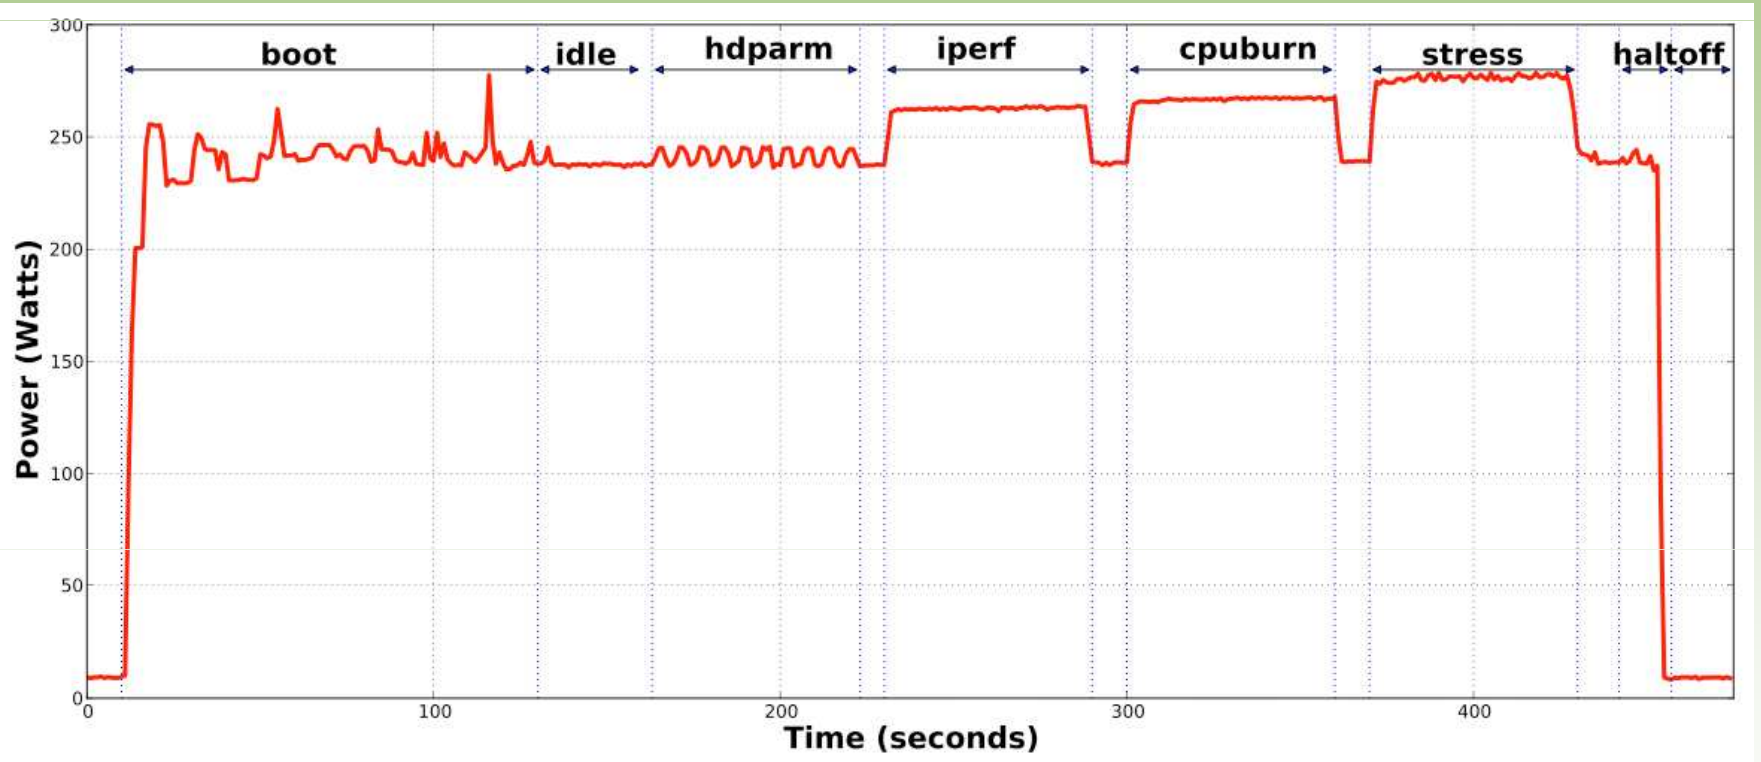

Example II : detecting anomalies

Improving frameworks/middleware and policies

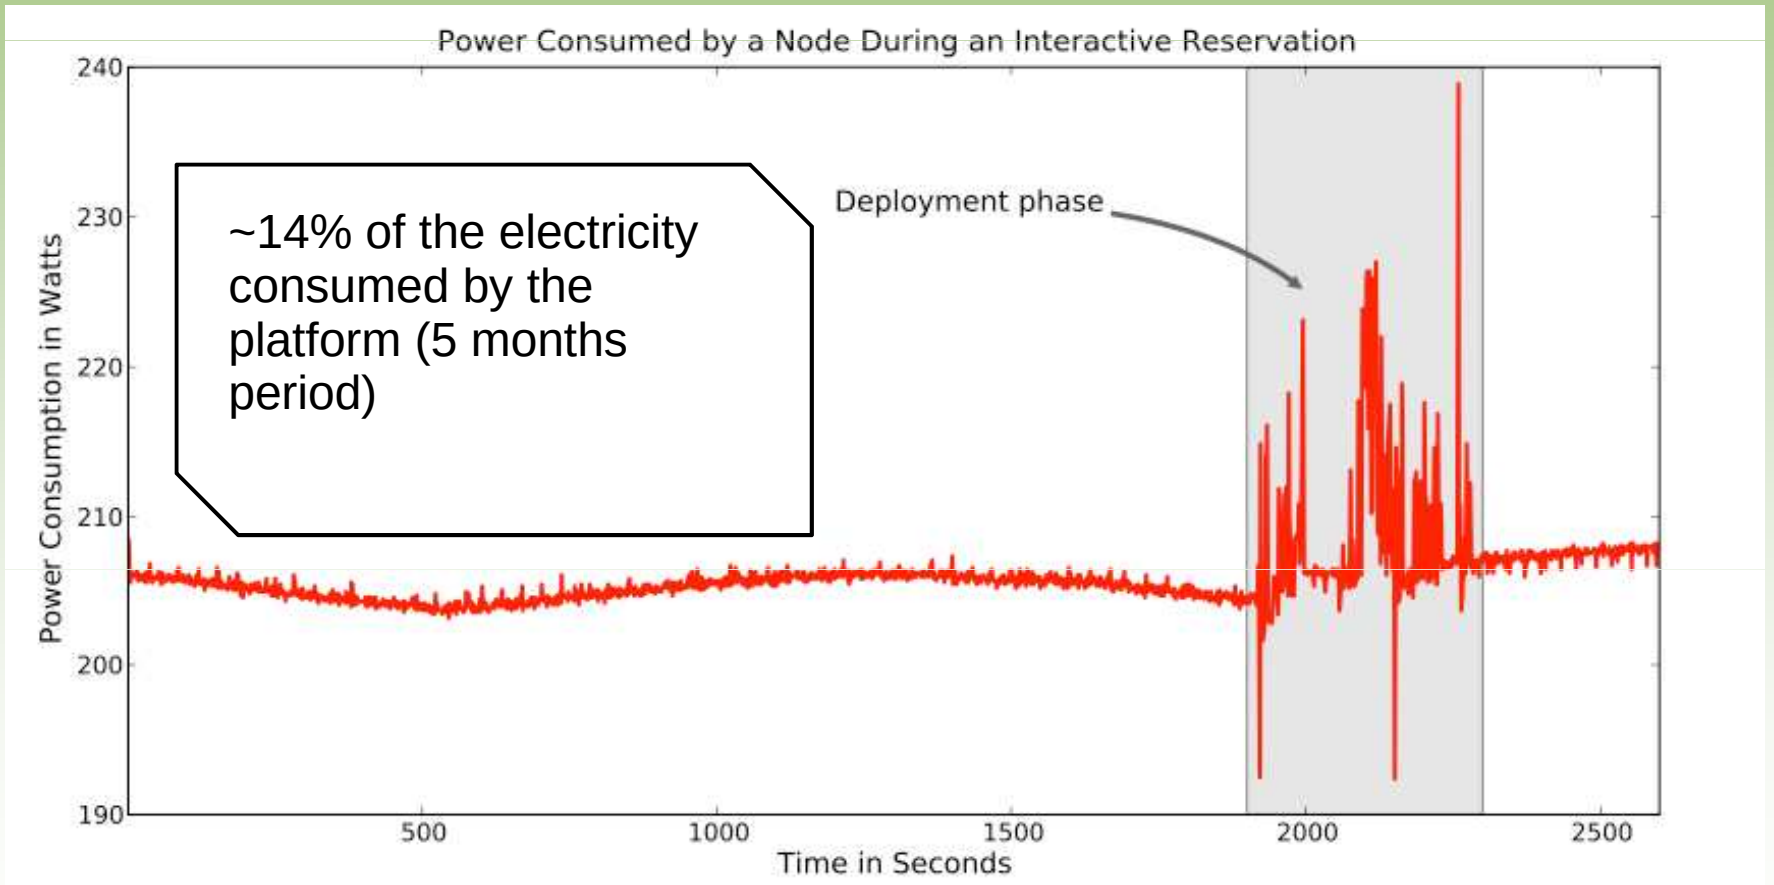

Example III : providing global views

Understanding the overall infrastructure

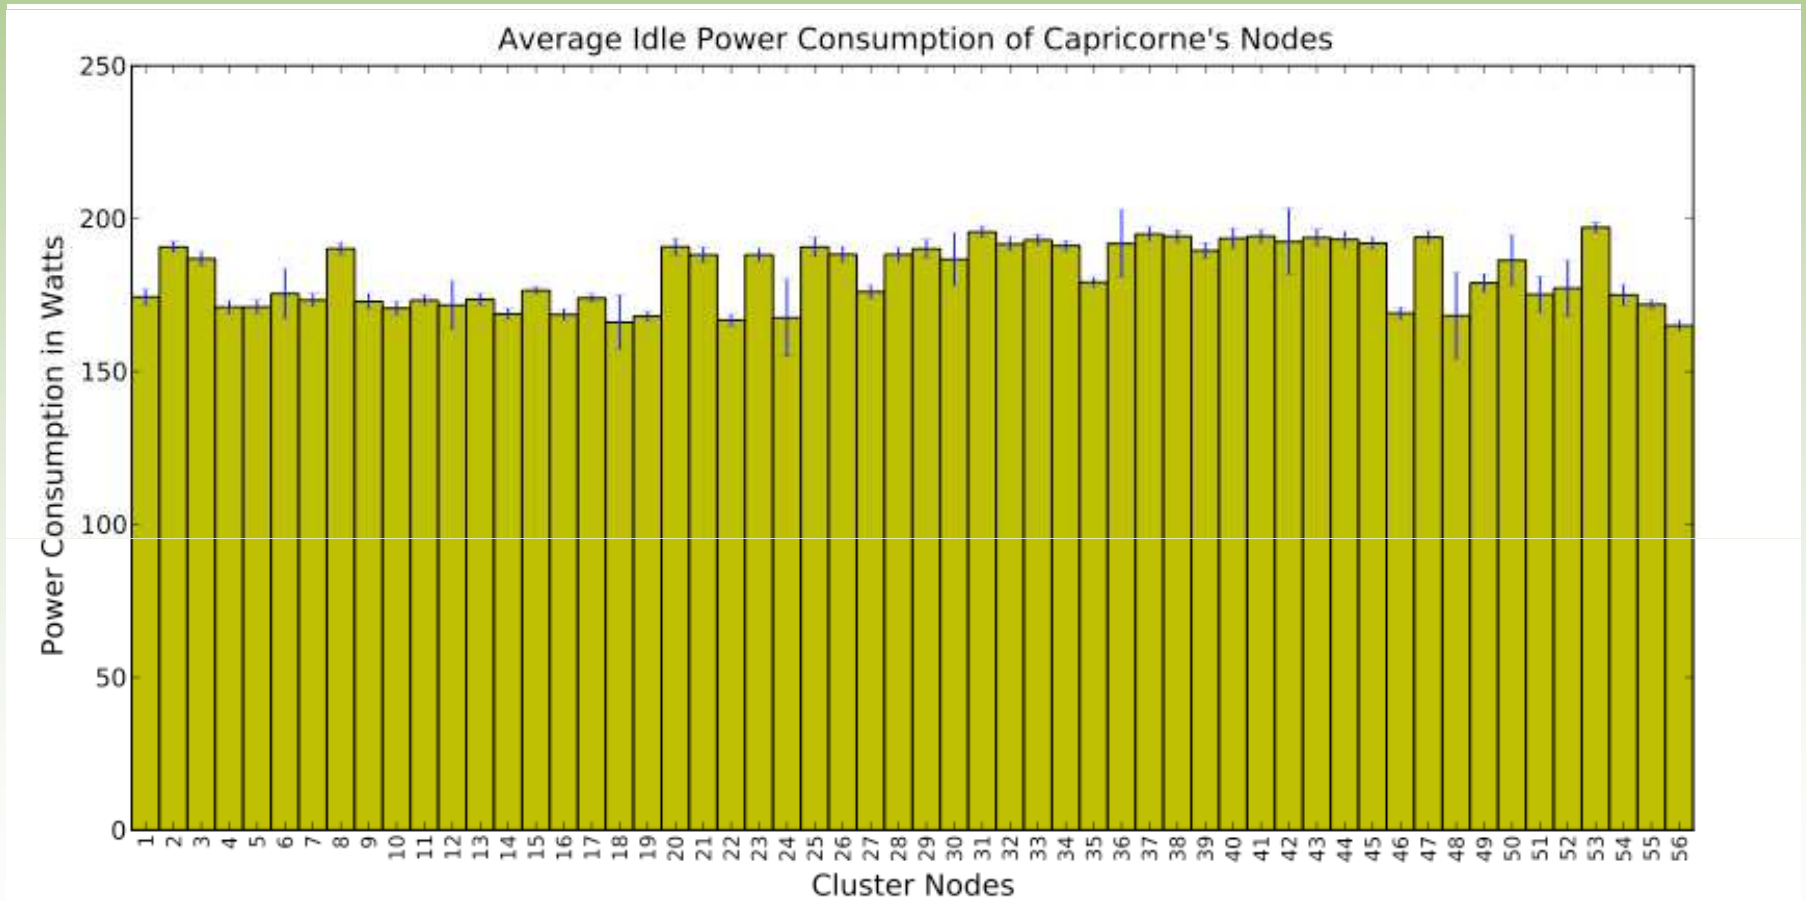

Large scale energy exposing

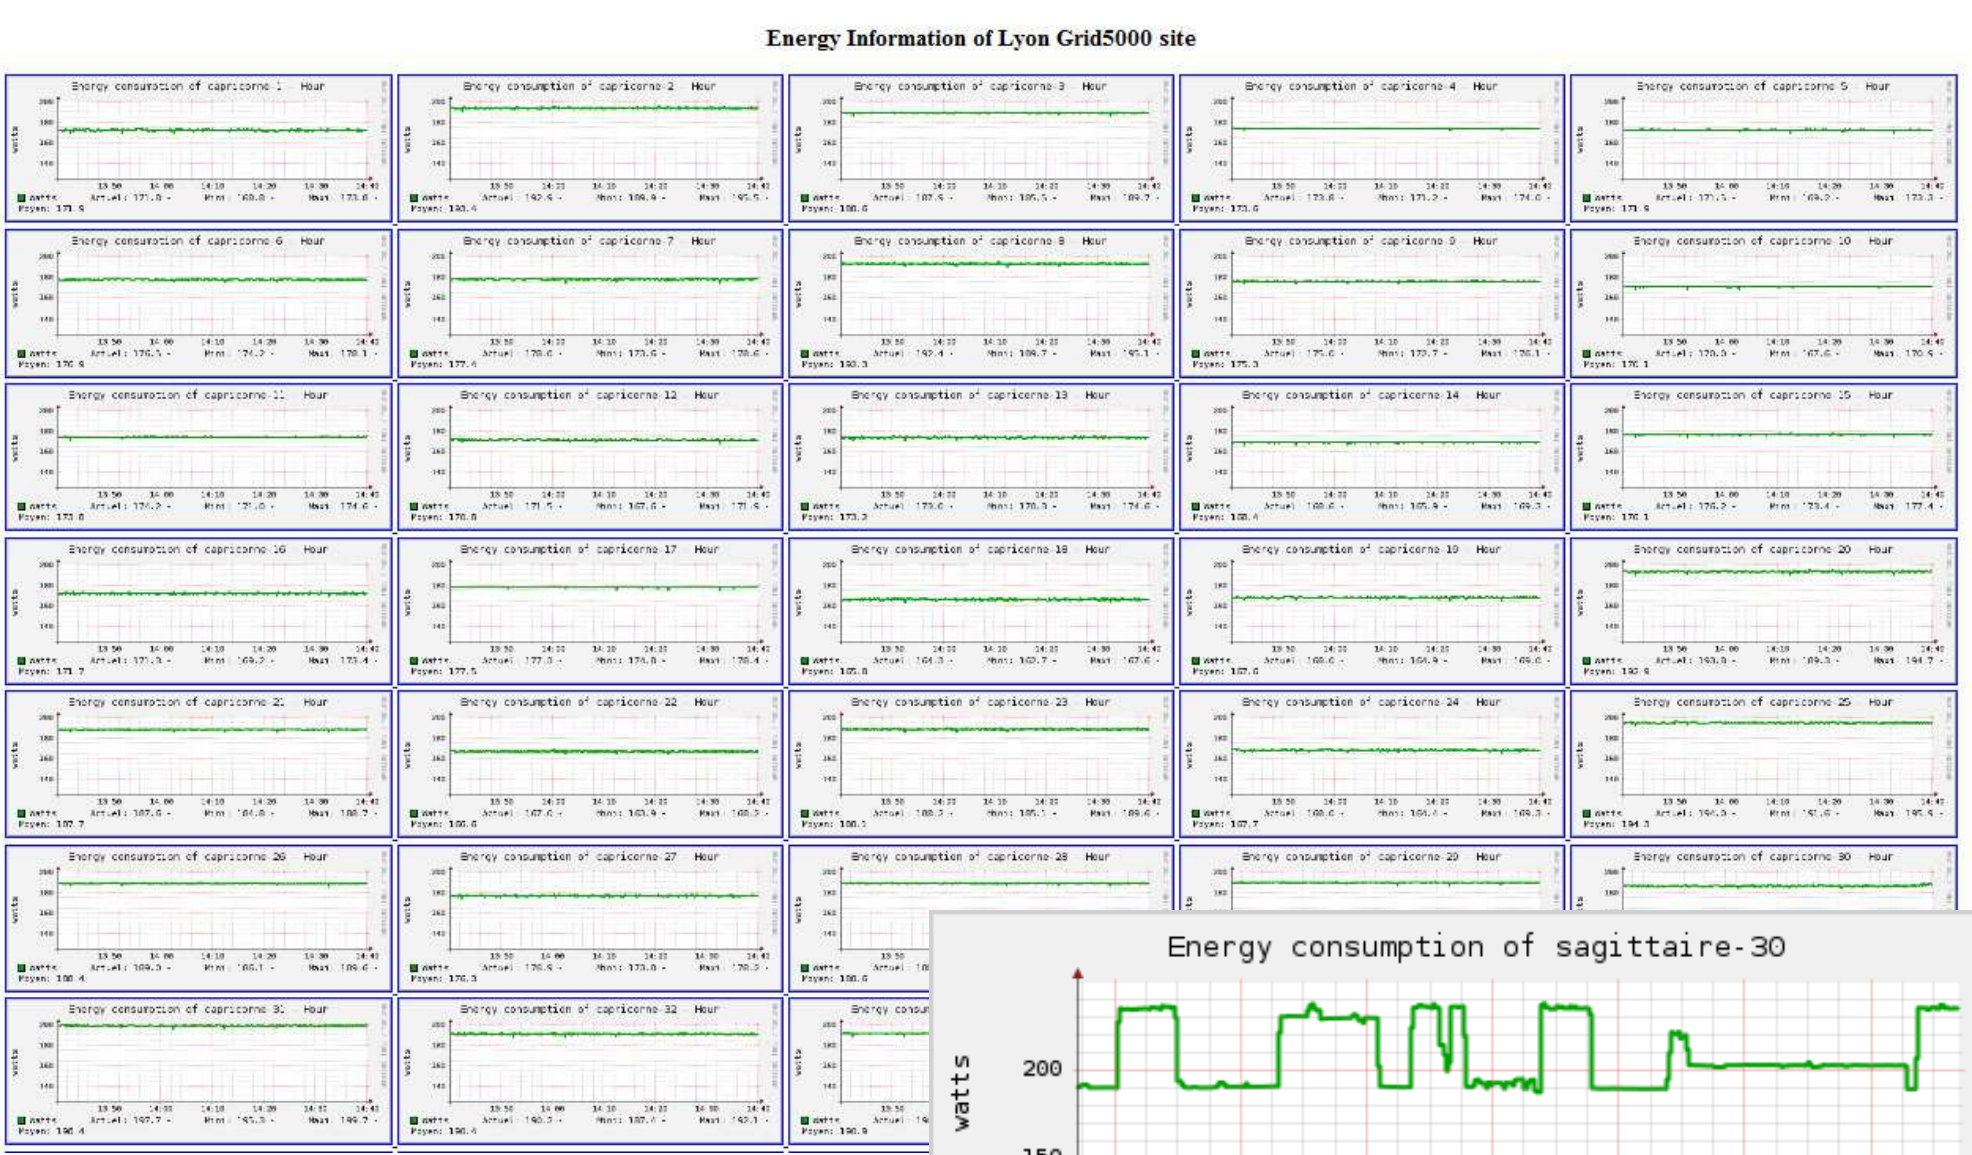

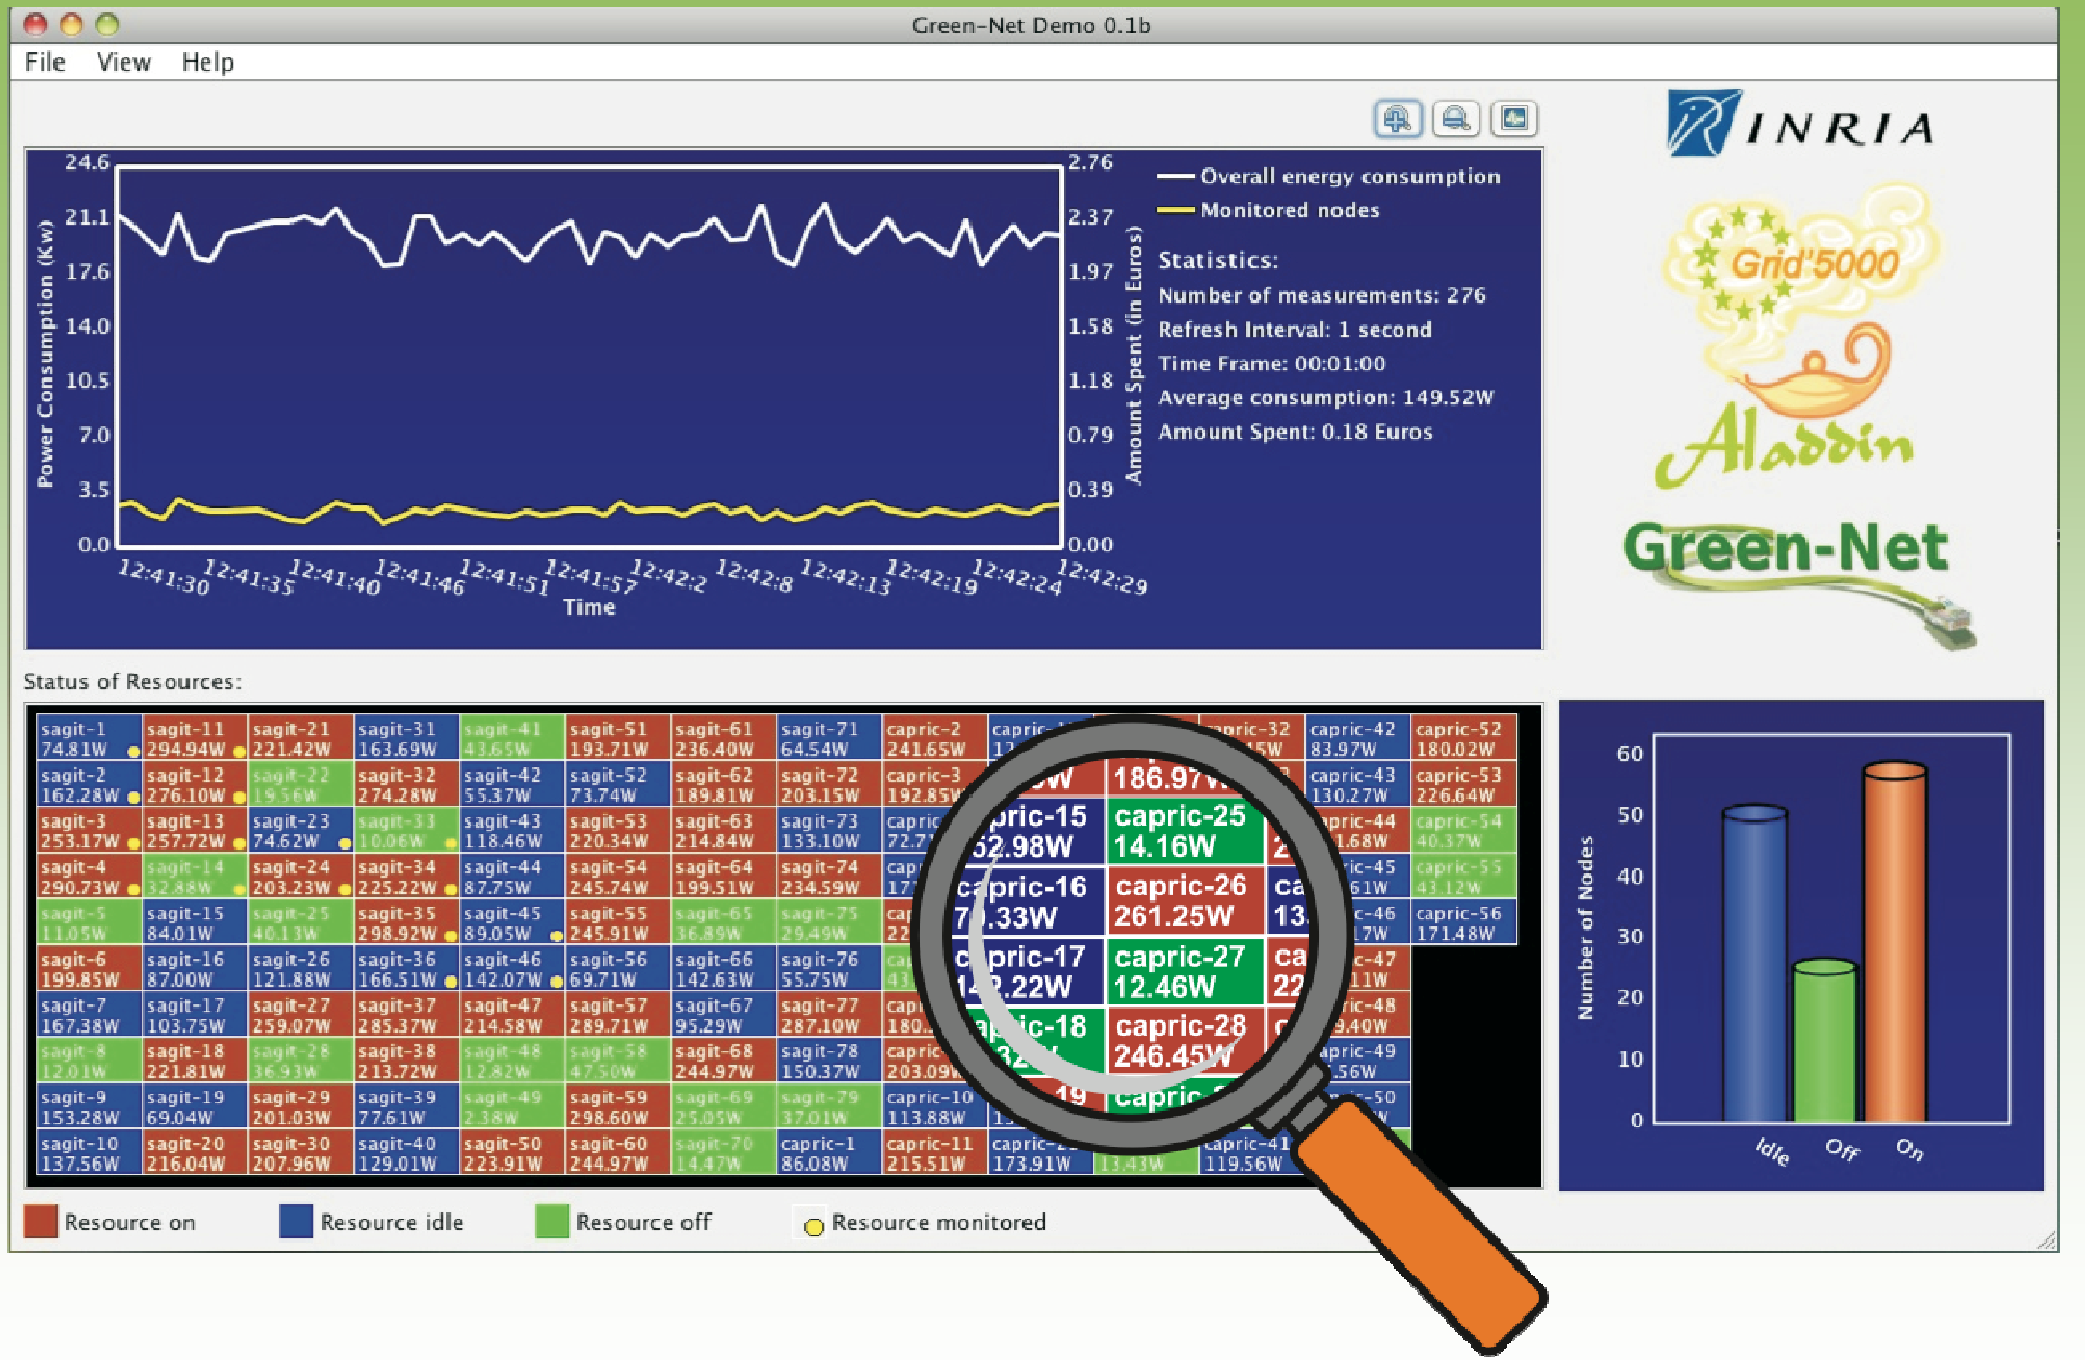
Green-Net Demo 0.1b
File View Help
INRIA
Grid'5000
Aladdin
Green-Net
Power Consumption (Kw)
Amount Spent (in Euros)
Time
Overall energy consumption
Monitored nodes
Statistics:
Number of measurements: 276
Refresh Interval: 1 second
Time Frame: 00:01:00
Average consumption: 149.52W
Amount Spent: 0.18 Euros
Status of Resources:
Resource on
Resource idle
Resource off
Resource monitored
Number of Nodes
Idle
Off
On

www.cost804.org

Action IC0804

ICT-Energy-Logs4all !

- Let's put together and exchange collections of energy traces for large scale distributed systems
- Goal : having an open repository of workload, energy and power traces in order to validate energy models on distributed scenario
- Proposal of a shared trace format
- First traces available : workload and power traces from the Grid5000 platform site of Lyon for a 6 months period : 6 Gbytes of data

  http://www.ens-lyon.fr/LIP/RESO/ict-energy-logs/

- Contact us : M. Dias, A.-C. Orgerie, L. Lefèvre

Experimental Validation of EARI

- Real traces of an experimental Grid: Grid'5000
- 4 sites, one year period, different policies

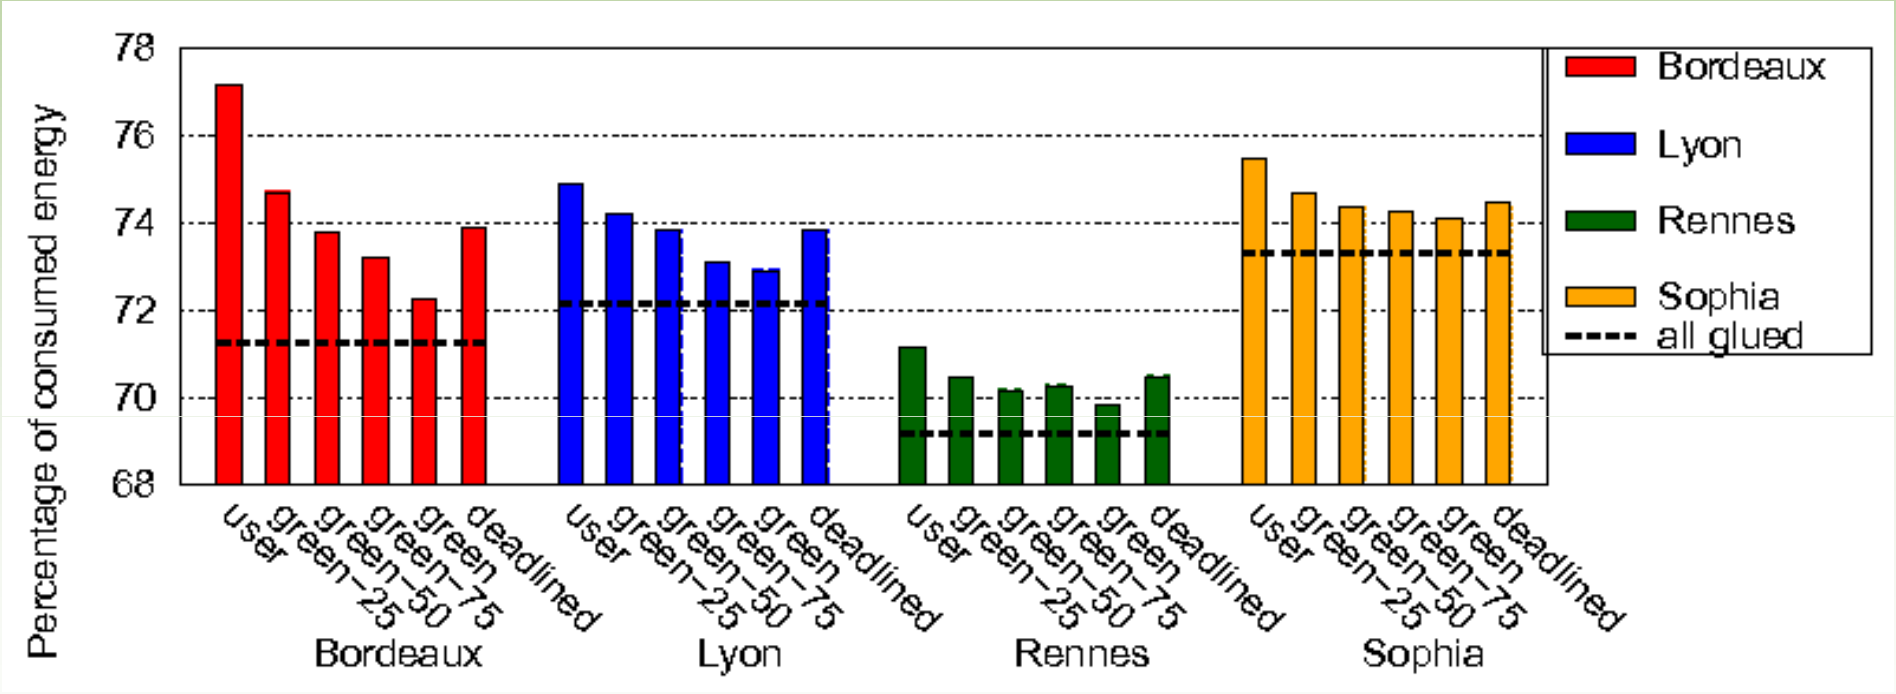

Green / Cloud links

- Aggegrating factor
- Flexibility
- Accounting/SLAs

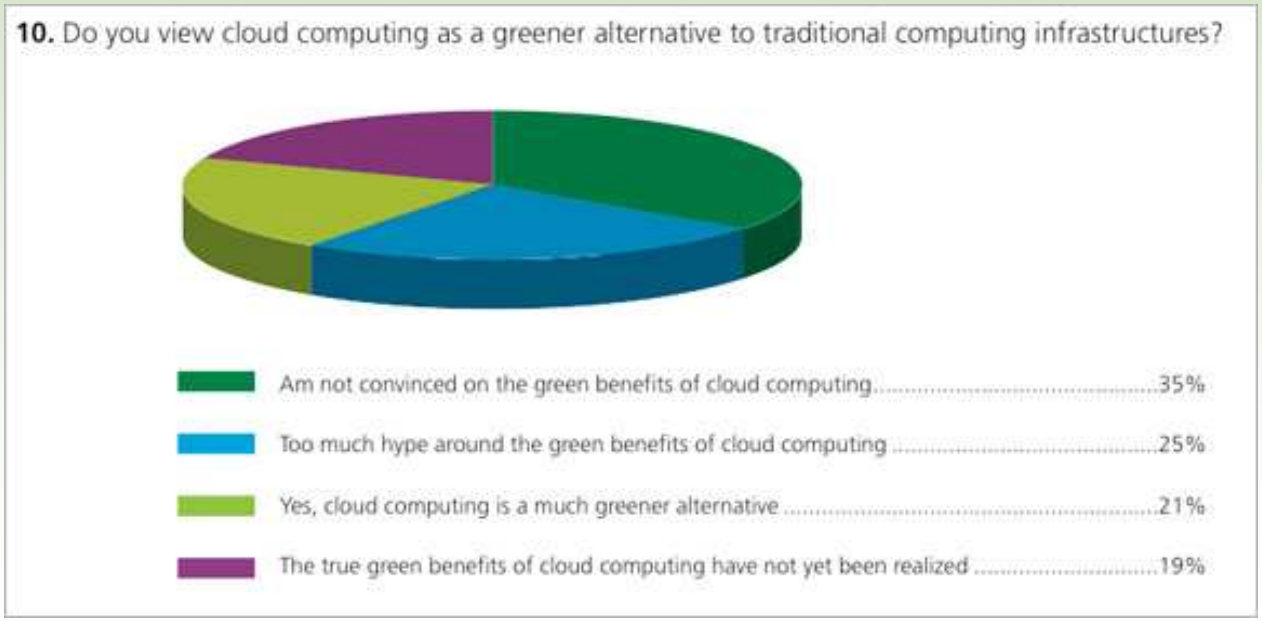

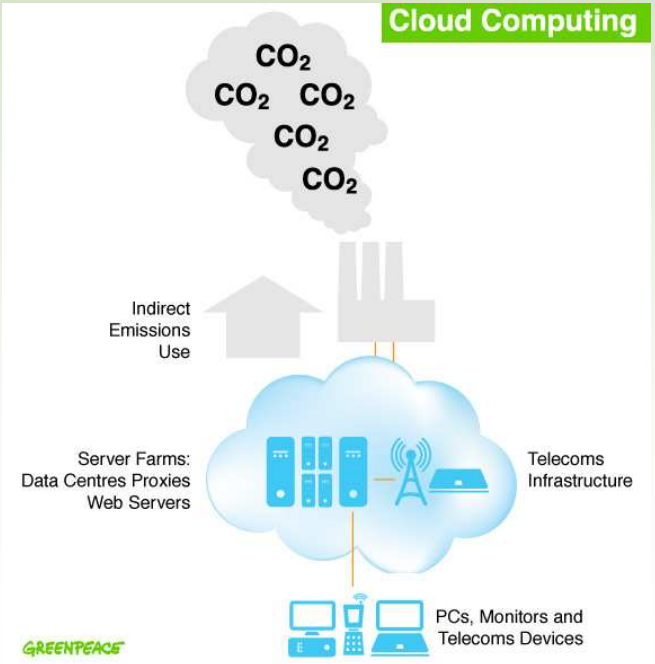

Green leverages on Clouds

- Designing the Green Open Cloud architecture based on the ERIDIS model
- Evaluatind and supporting features like live migration, tuning capping for agregating

ERIDIS support for Green Open Cloud

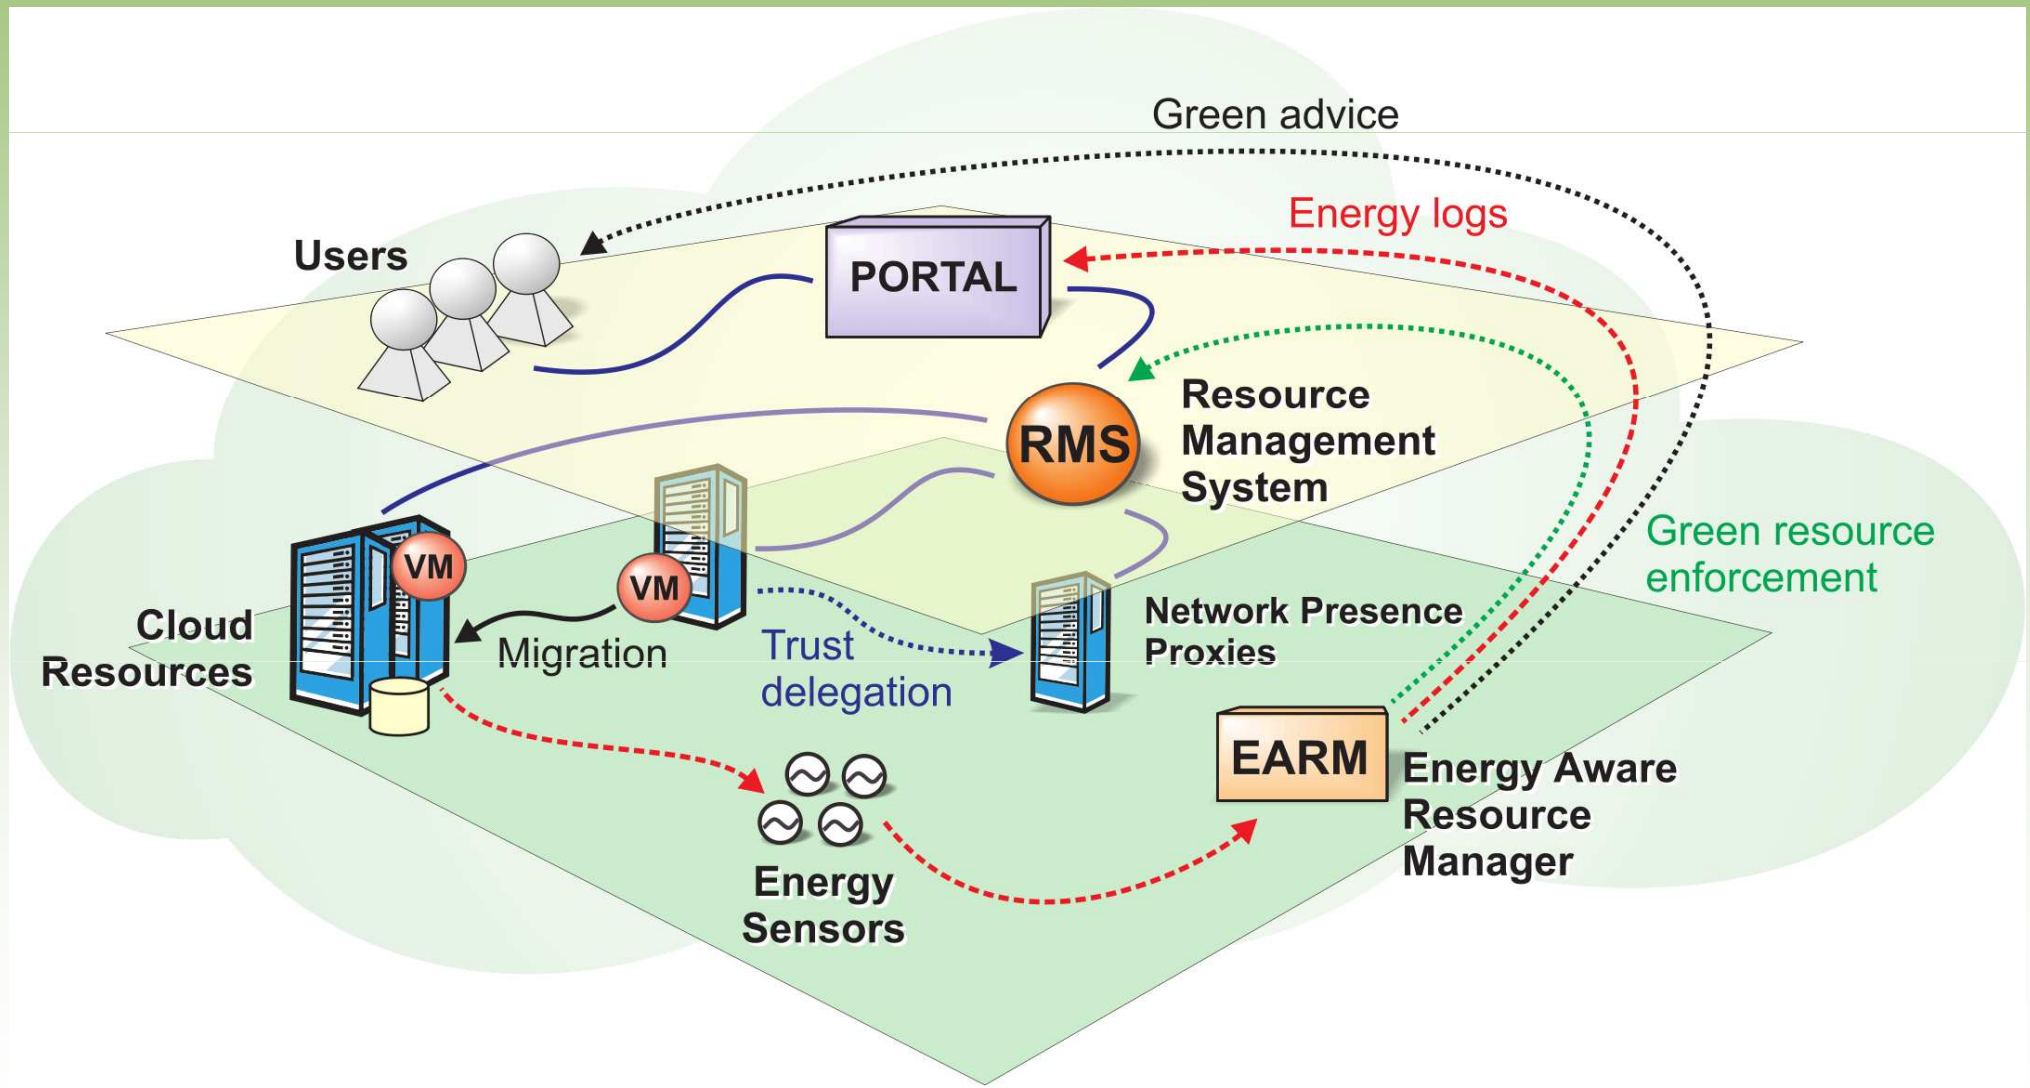

GOC Resource Manager

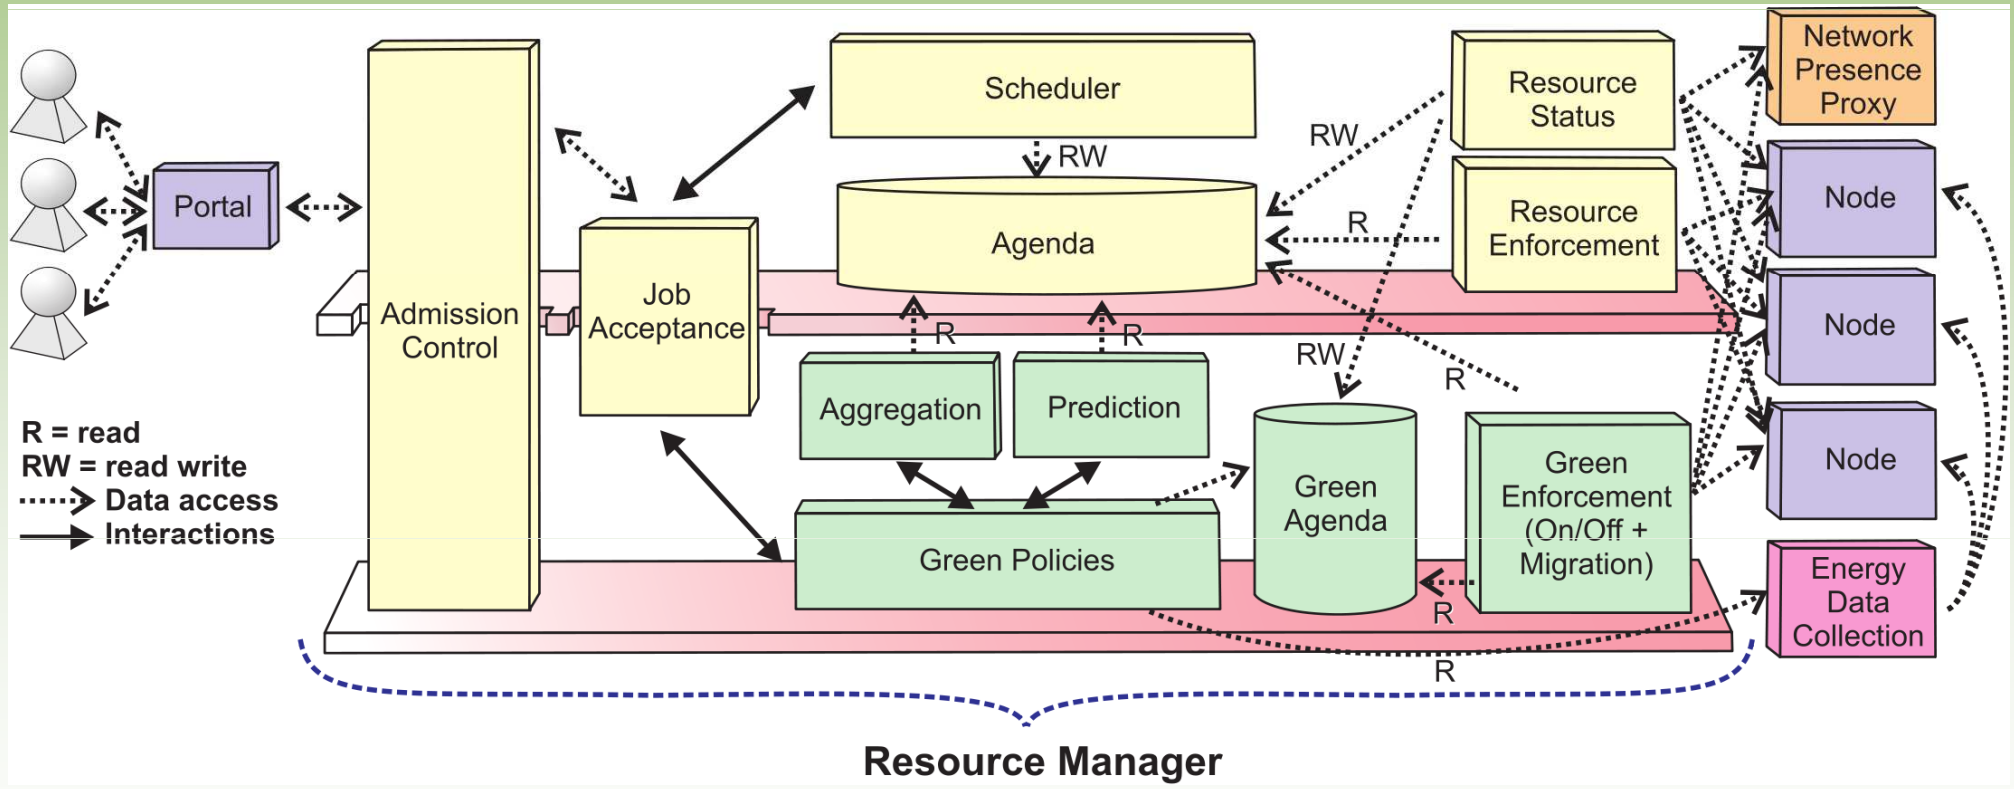

Evaluating leverages : Boot, Run and Halt

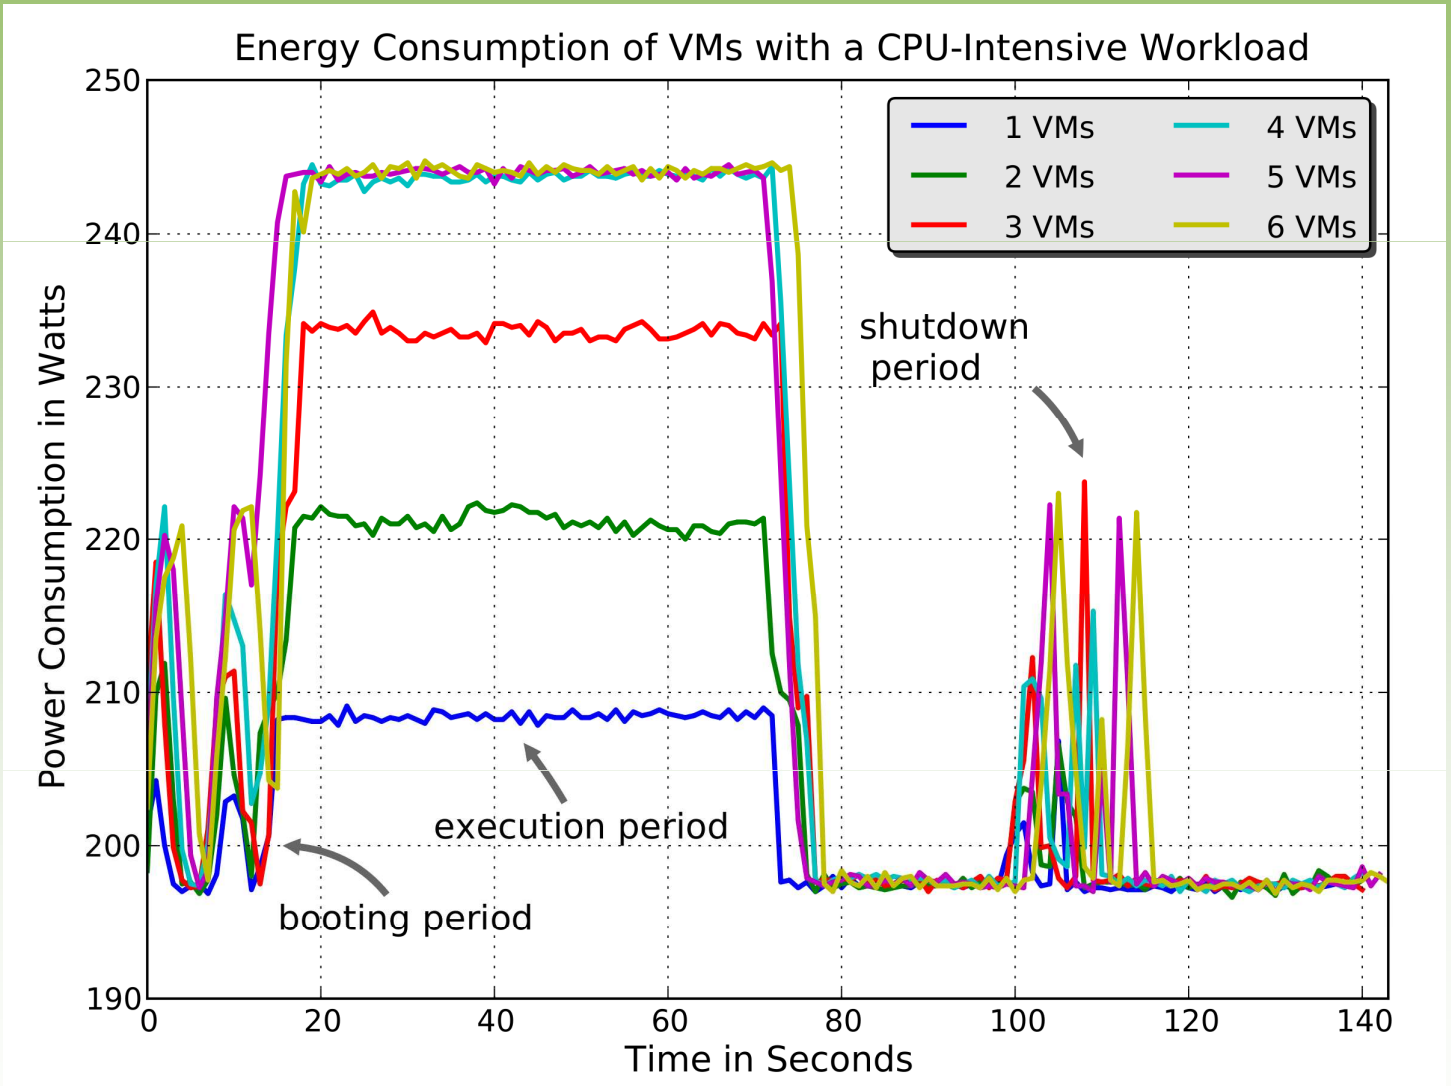

HP Proliant 85 G2 Servers (2.2 GHz, 2 dual core CPUs per node) XenServer 5.0

- 6% increase of energy with 1 VM running

Evaluating cloud leverages

- Capping

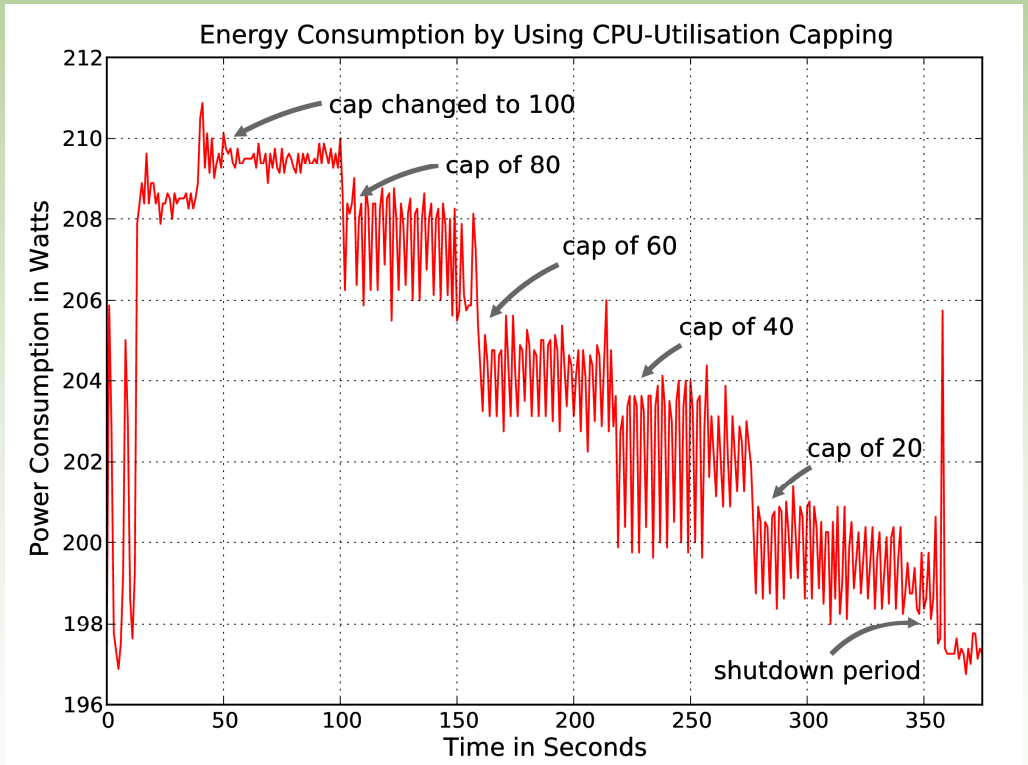

- Tuning

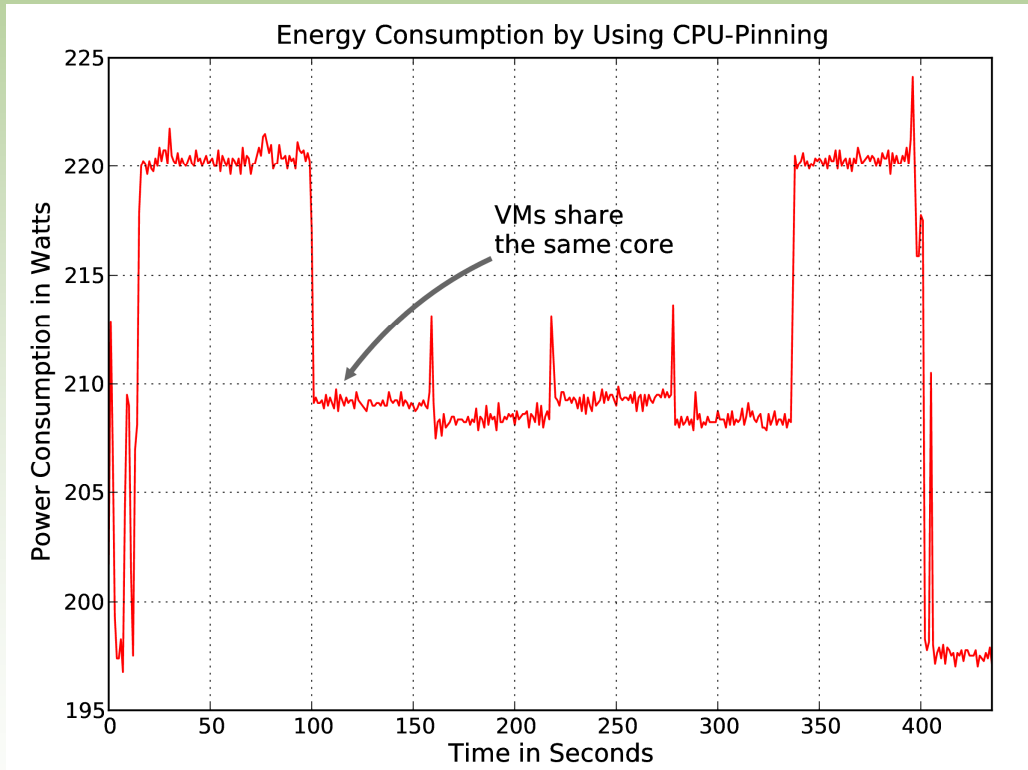

Evaluating Cloud leverages : Migration

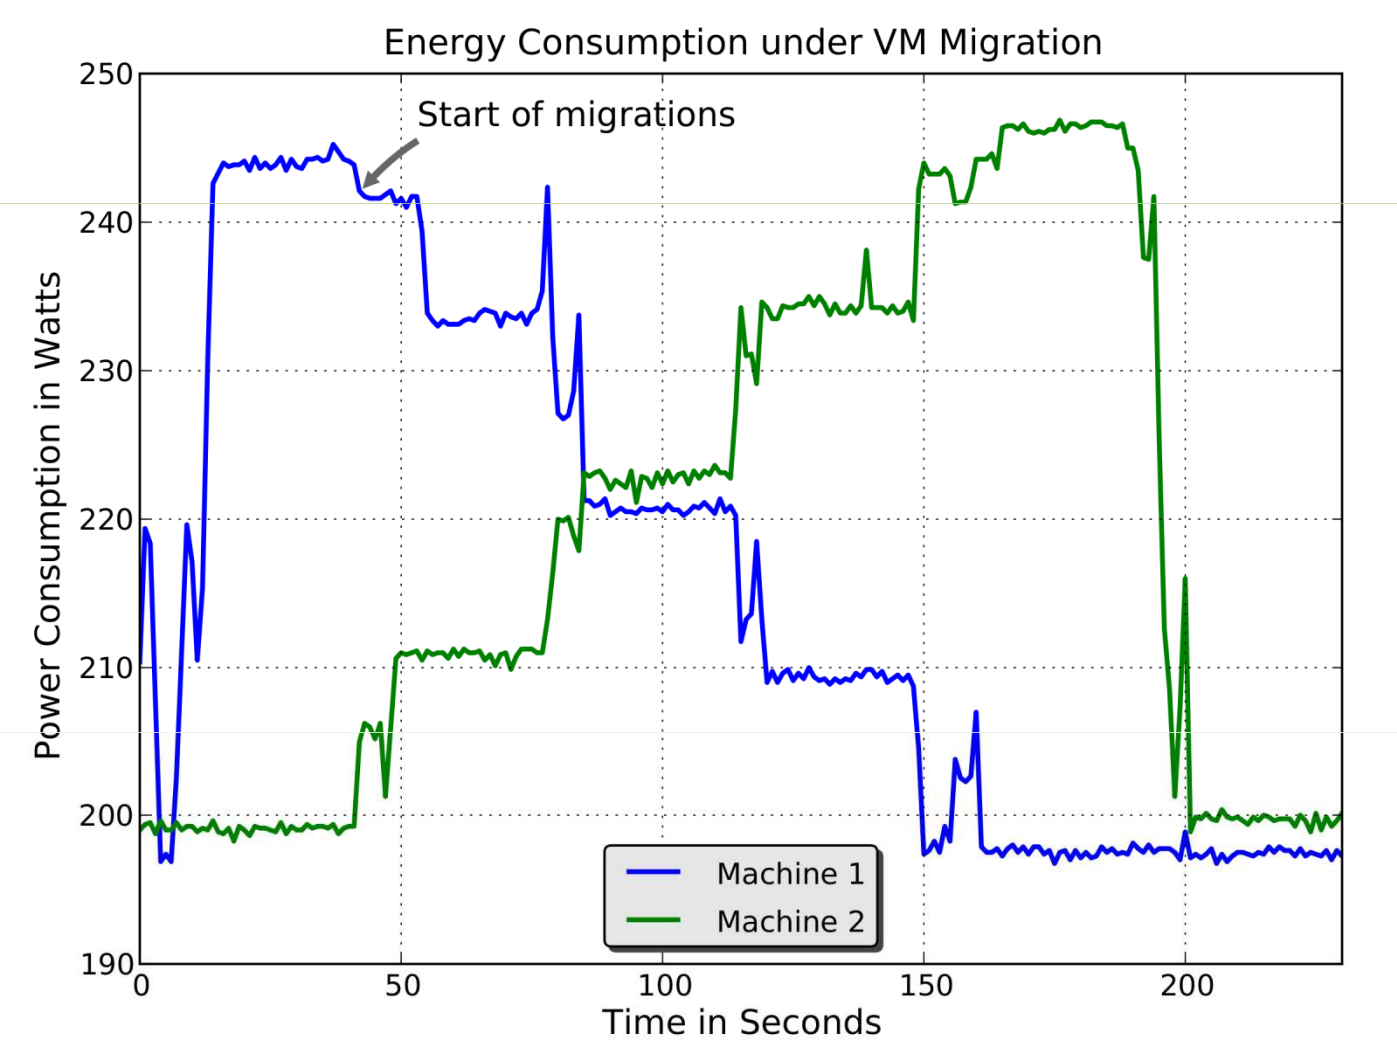

- Bad moment in energy during the migration

Applying some leverages : From « perfect » load balance (Perf oriented) to unbalance (Green)

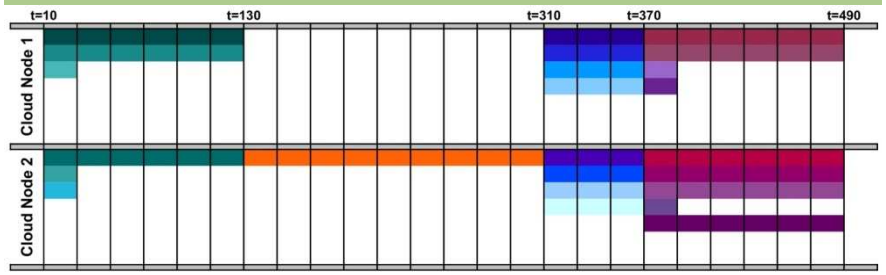

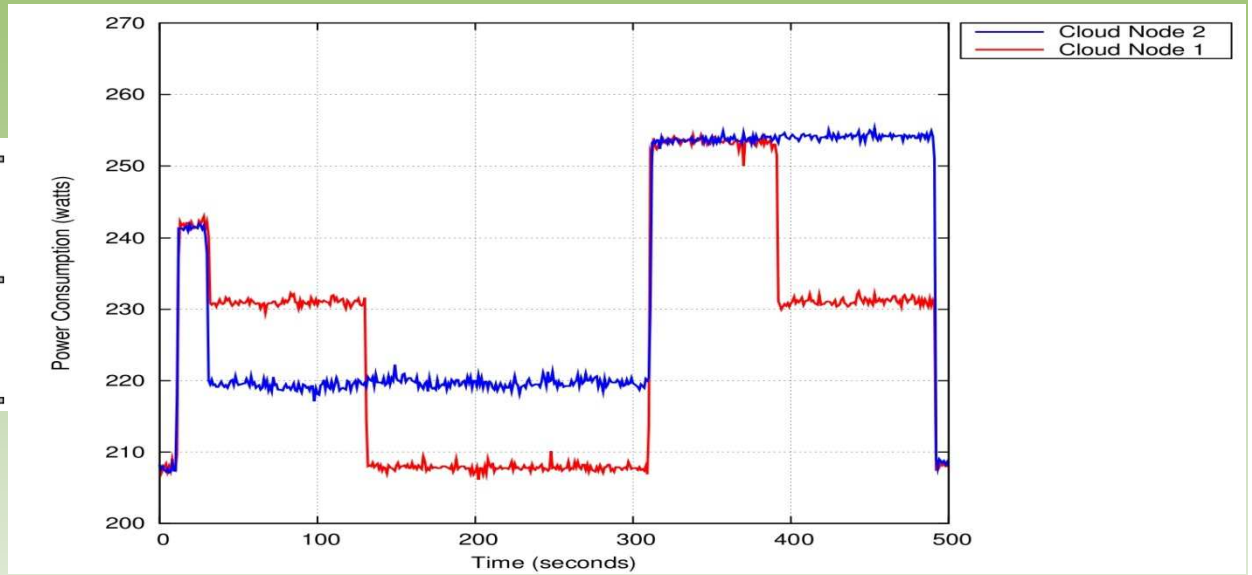

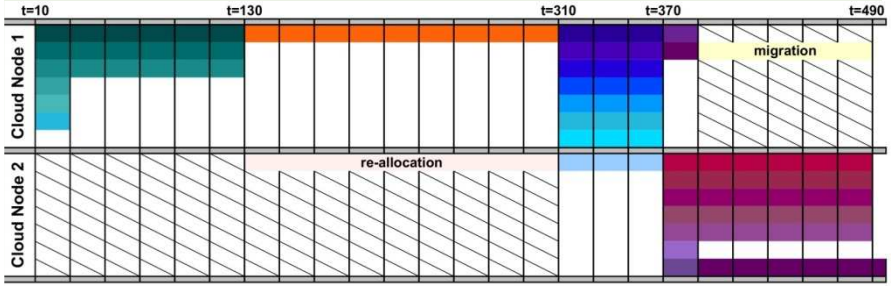

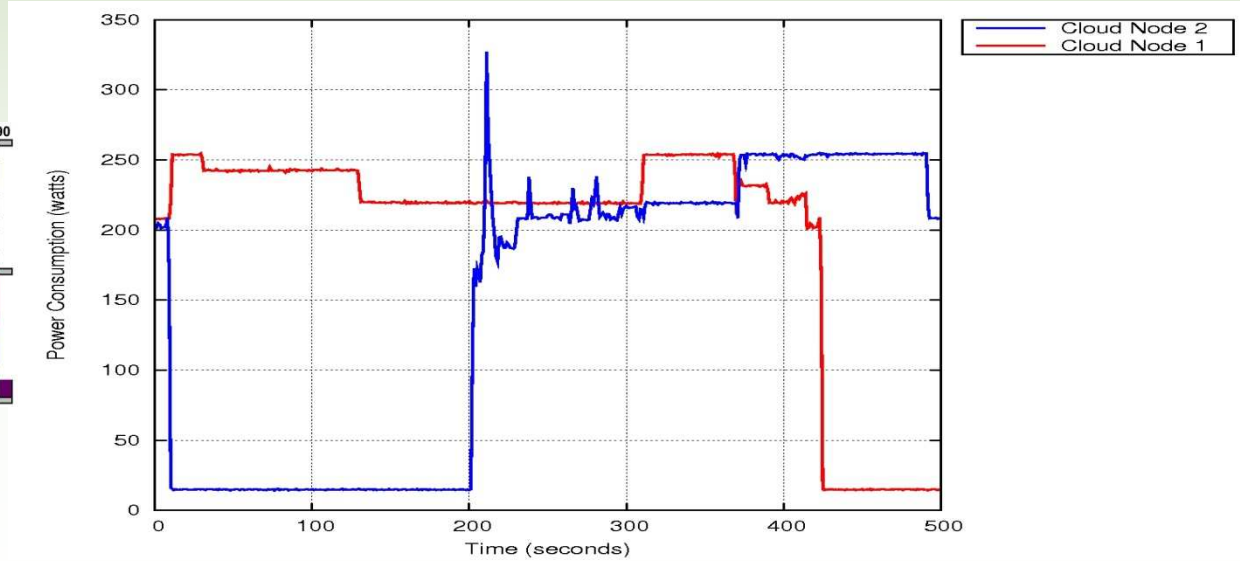

French Cloudware

- CompatibleOne : BULL (Leader), ActiveEon, CityPassenger, eNovance, Eureva, INRIA, Institut Télécom, Lost Oasis Mandriva,Nexedi SANuxeo SA, OW2, Prologue SA, XWiki
- FUI Project (2010-2012)
- Proposing an open source cloud software stack (with energy aware software components)

Conclusion and Perspectives

- At the beginning, the Grid was “green” (cycle stealing, old machines usage…) like the Cloud (aggregation)
- Energy aspects change the way we design Grids&Clouds applications, protocols, services and policies (i.e. load balancing is not always the best solution)
- Challenge : design energy proportional equipments and frameworks (computing, memory or network usage)
- Focus on usage : current step Action interfaces IdG/Aladdin INRIA – live measurement of production centers

Questions?

Laurent Lefèvre
Laurent.lefevre@inria.fr