Energy efficient approach for grids, clouds and networks

Laurent Lefèvre
INRIA – Université de Lyon

laurent.lefevre@inria.fr

CCGSC2010, Flat Rock, 7th September 2010

Energy : 1st challenge for large scale systems ?

- Yes ! as important as latency, bandwidth, fault tolerance…
- Exascale systems from 20 to 100MW (F.Cappello)
- 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

Public Concerns !

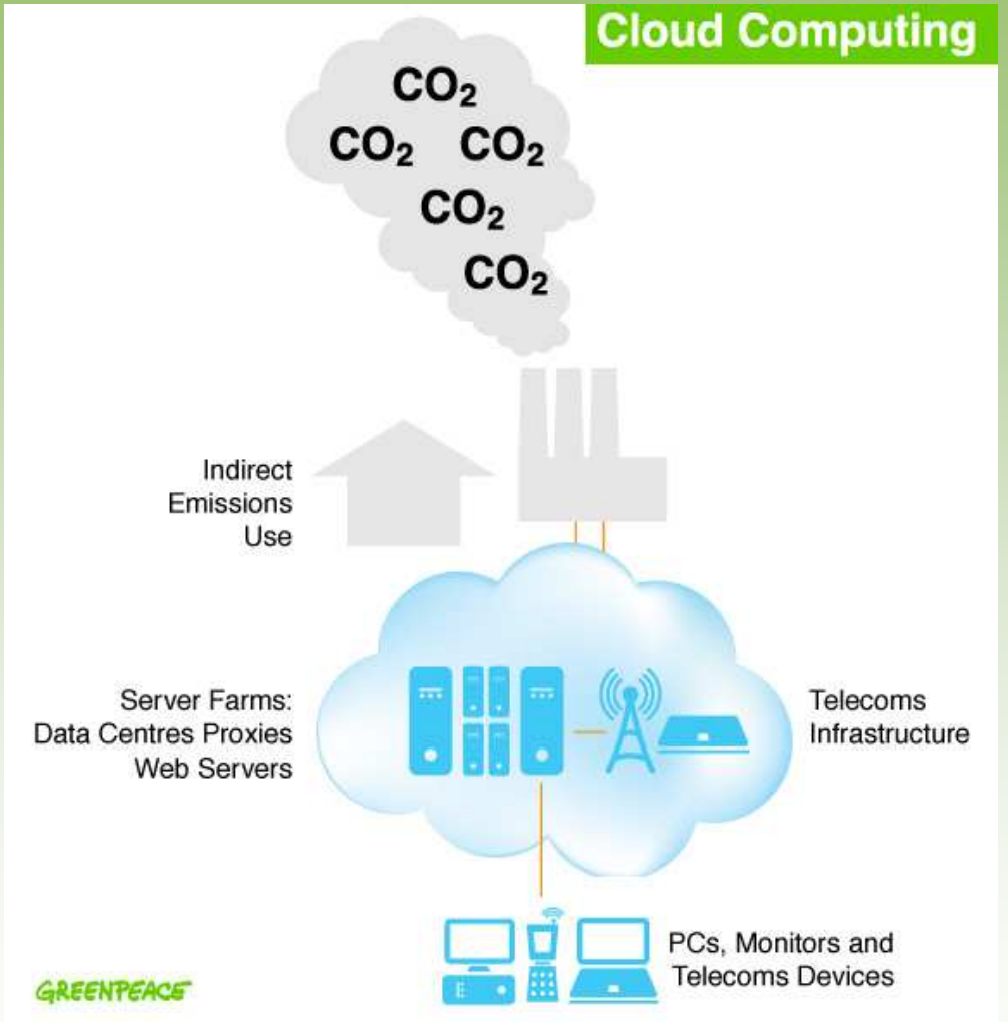

Scientific Concerns !

Today

Tomorrow ?

- First approach :evaluate / analyze
- Push simulations vs large experiments
- Reduce amount of co2 consuming research ?

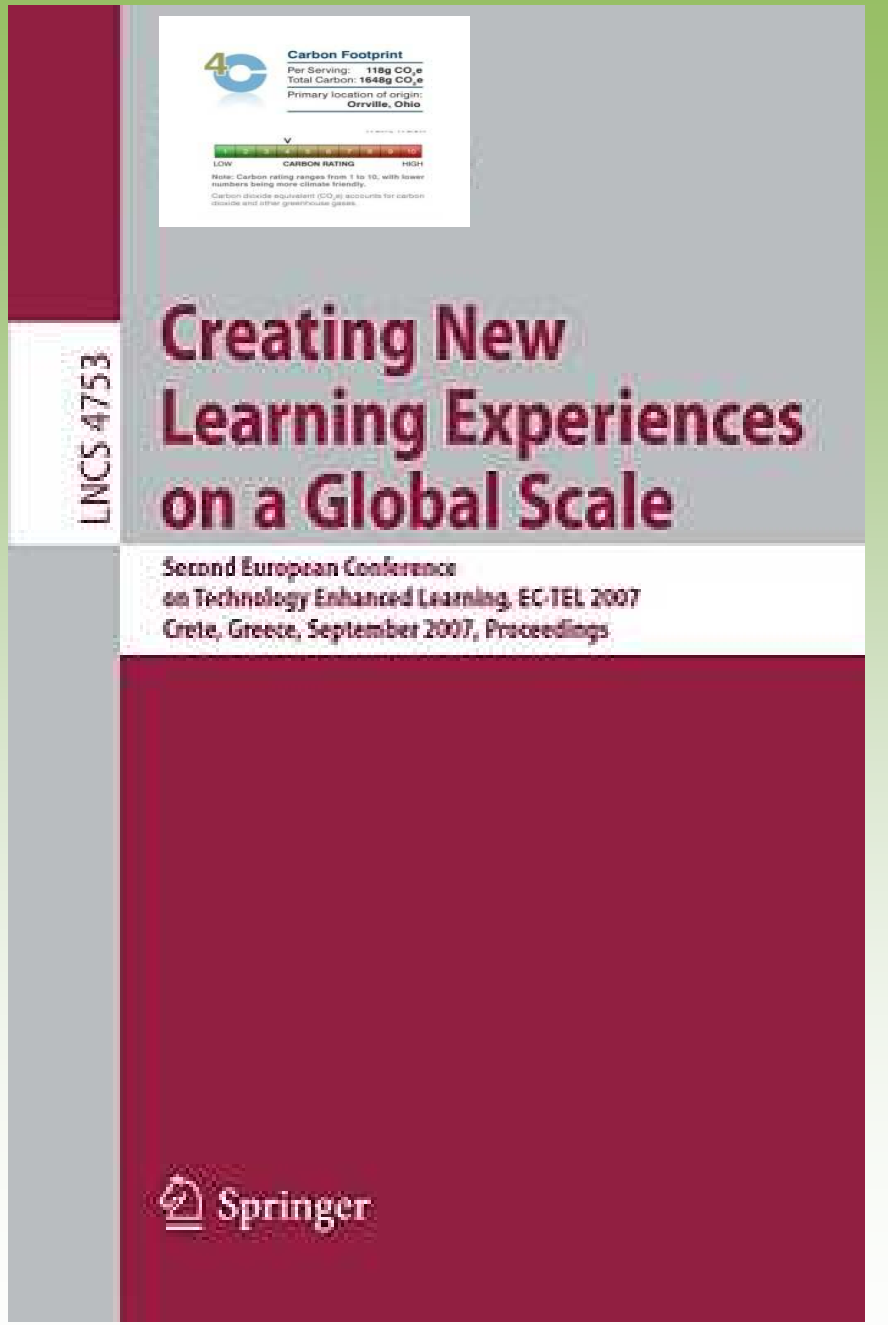

Towards Energy Aware Large Scale Systems : open questions

How to decrease the energy consumption 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 (i.e. energy aware libraries)
- **Coordinating :** using large scale approaches to enhance green leverages

Explosion of initiatives

For each domain

- **Data centers/HPC : Green500, EU CoC**
- **Grids : The Green Grid / Open Grid Forum**
- **Storage : SNIA**
- **Networks : Green Touch / EEE**

Methodology

- Proposing a generic 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 : switching on/off VMs
    - Networks : lighting or switching off paths, nics, links, routers
  - Adapting « speed » (and consumption) to the need of appliciations/users
    - HPC, Grids : dvfs
    - Clouds : tuning, capping
    - Networks : adaptive link rate

The CEPAP model

- 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 : switch on/off or adapt performancs
- Predicting usage of infrastructures

CEPAP

Multi-View Understanding of Large Scale Systems Usage

Monitoring and Analyzing Energy Information

Designing Energy Efficient Frameworks

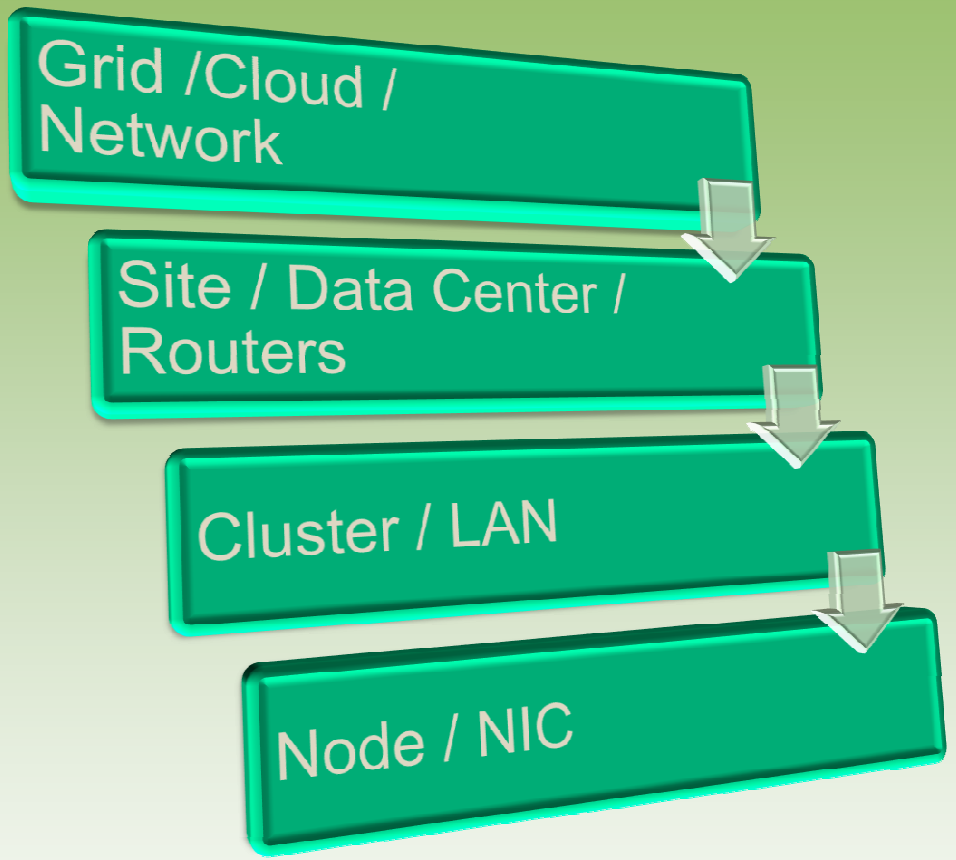

CEPAP

Multi-View Understanding of Large Scale Systems Usage

Monitoring and Analyzing Energy Information

Designing Energy Efficient Frameworks

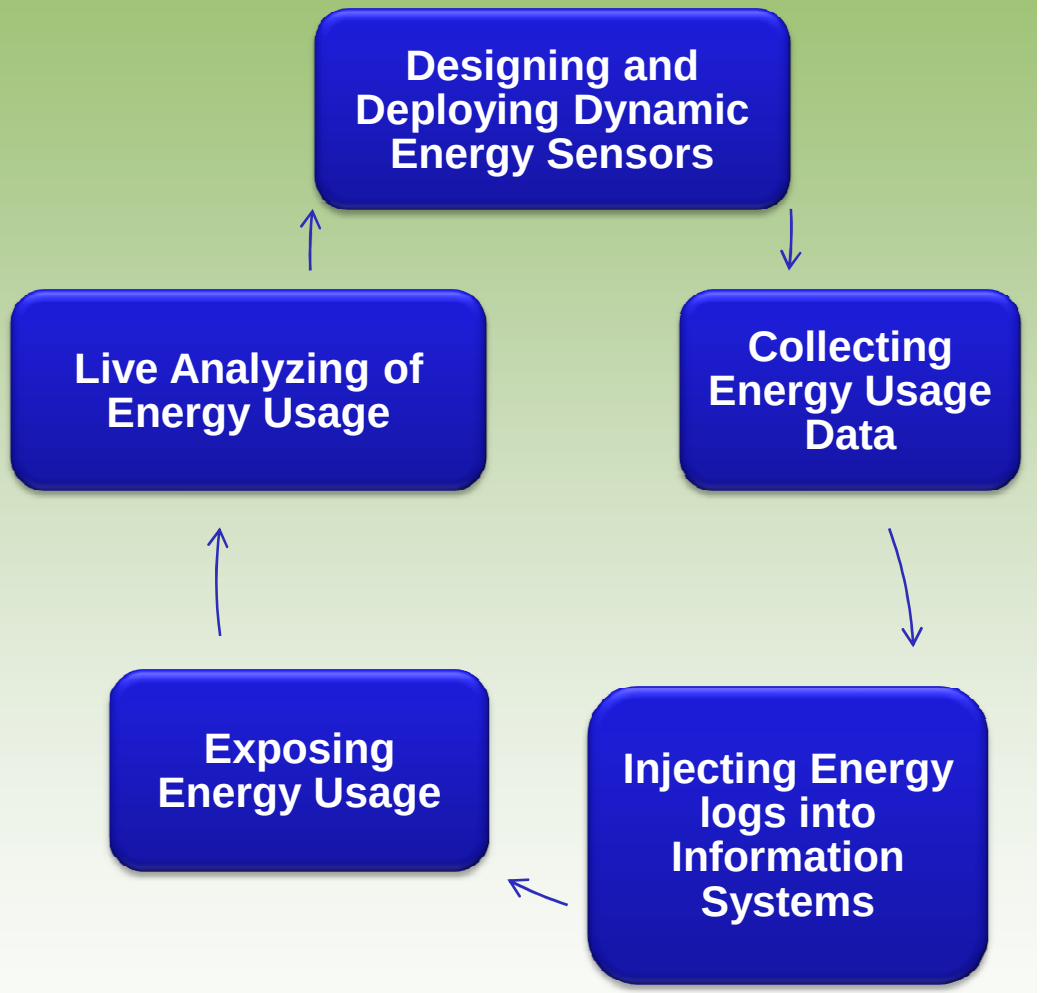

CEPAP

Multi-View Understanding of Large Scale Systems Usage

Monitoring and Analyzing Energy Information

Designing Energy Efficient Frameworks

Energy Aware Reservation Infrastructure

Prediction Systems

Adapted Schedulers and Resource Management

Node and Services Virtualization

Delegated Trust and Network Presence

Green Policies Support

Node Energy Controlers

The CEPAP Framework

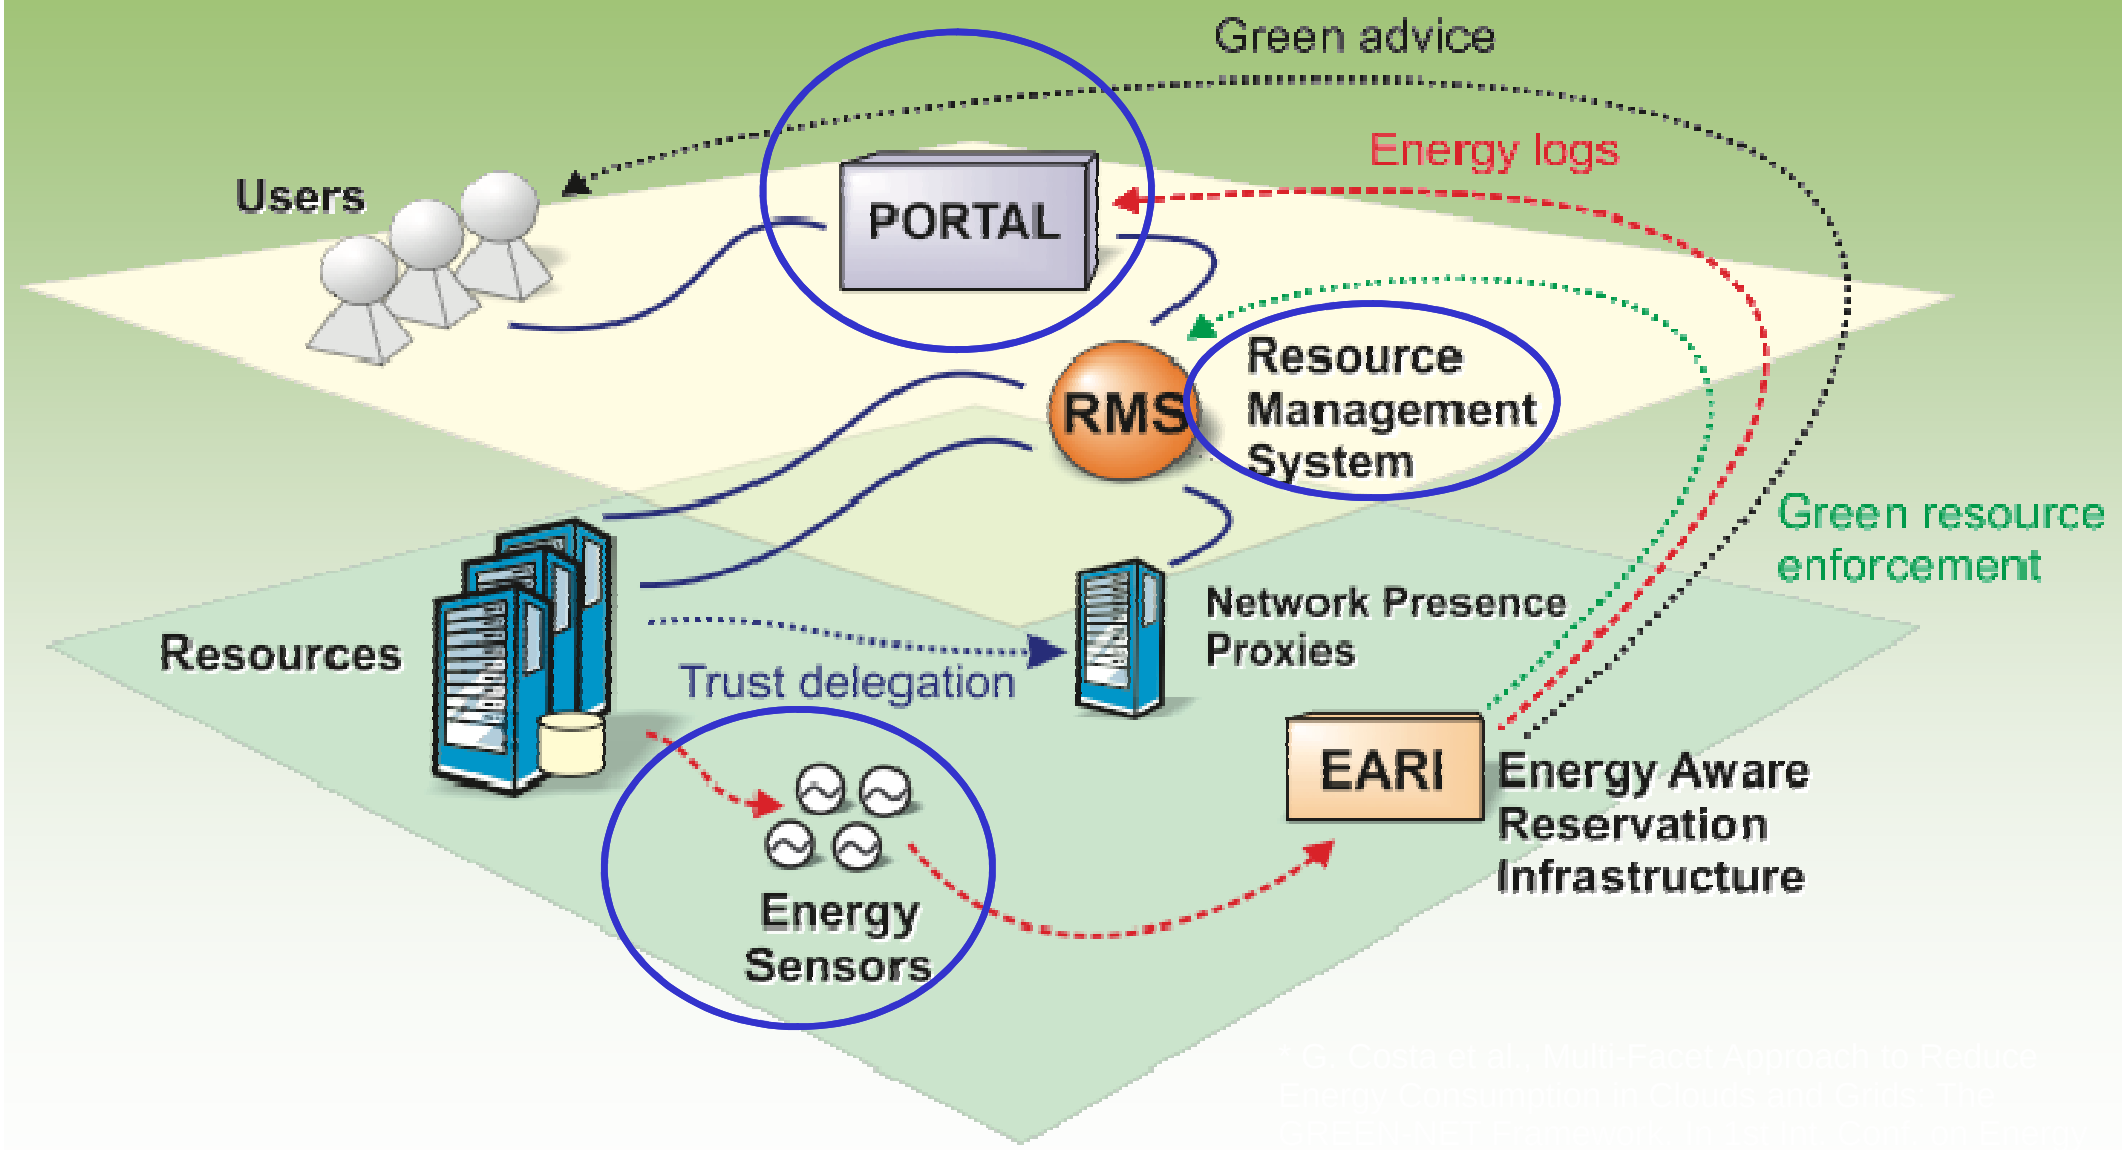

* G. Costa et al., Multi-Facet Approach to Reduce Energy Consumption in Clouds and Grids: The GreenNet Framework, the 1st Int. Conf. on Energy-Efficient Computing and Networking

1st focus : Collecting and exposing

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

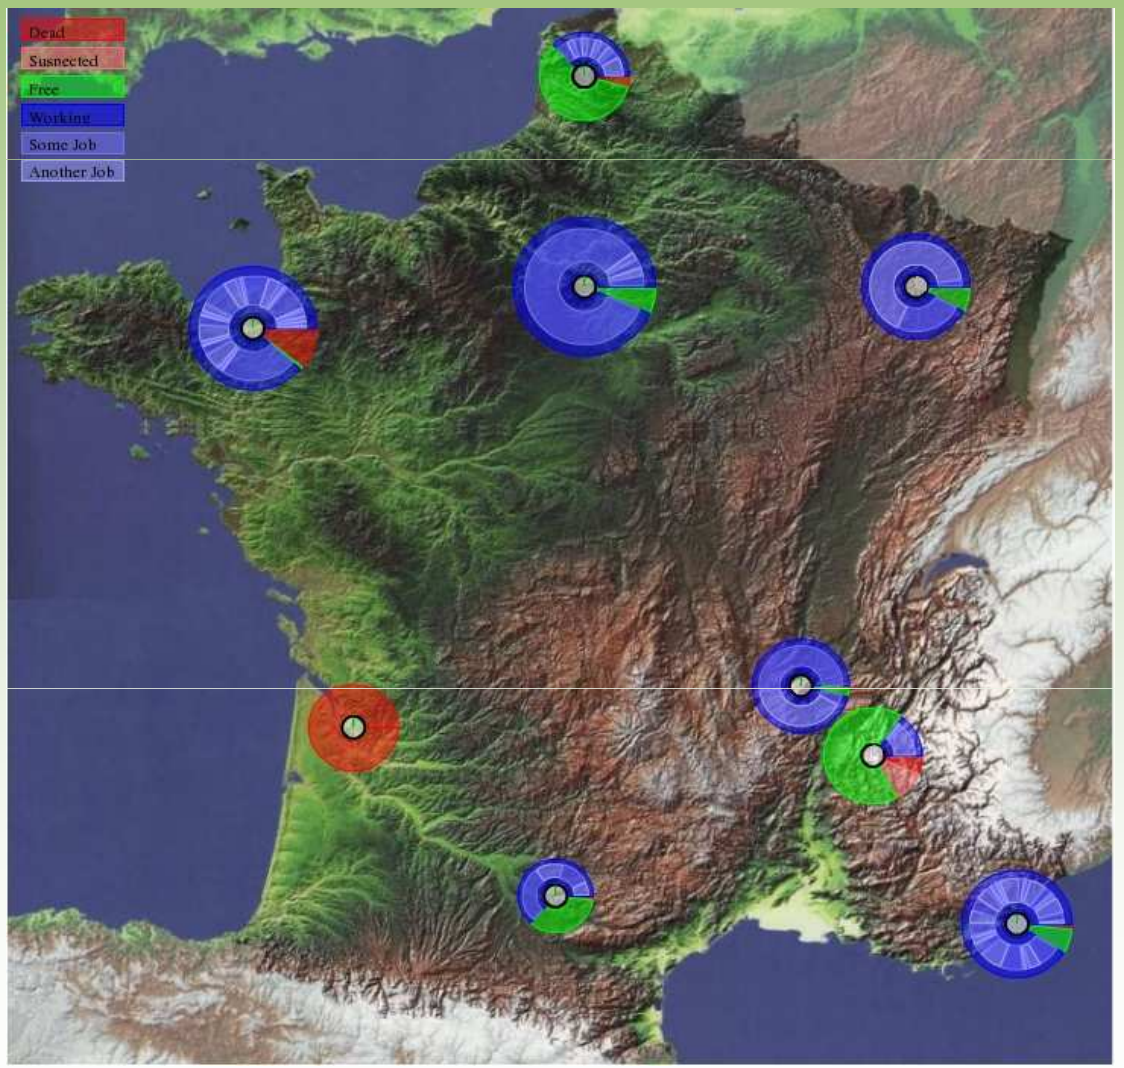

The Green Grid5000

Energy sensors

- 6 or 48 ports, communication via serial port
- Deployed on three sites of Grid'5000
- Library for **interfacing** with **energy sensors**
- **Client-side** applications to **obtain** and **store** the energy consumption data

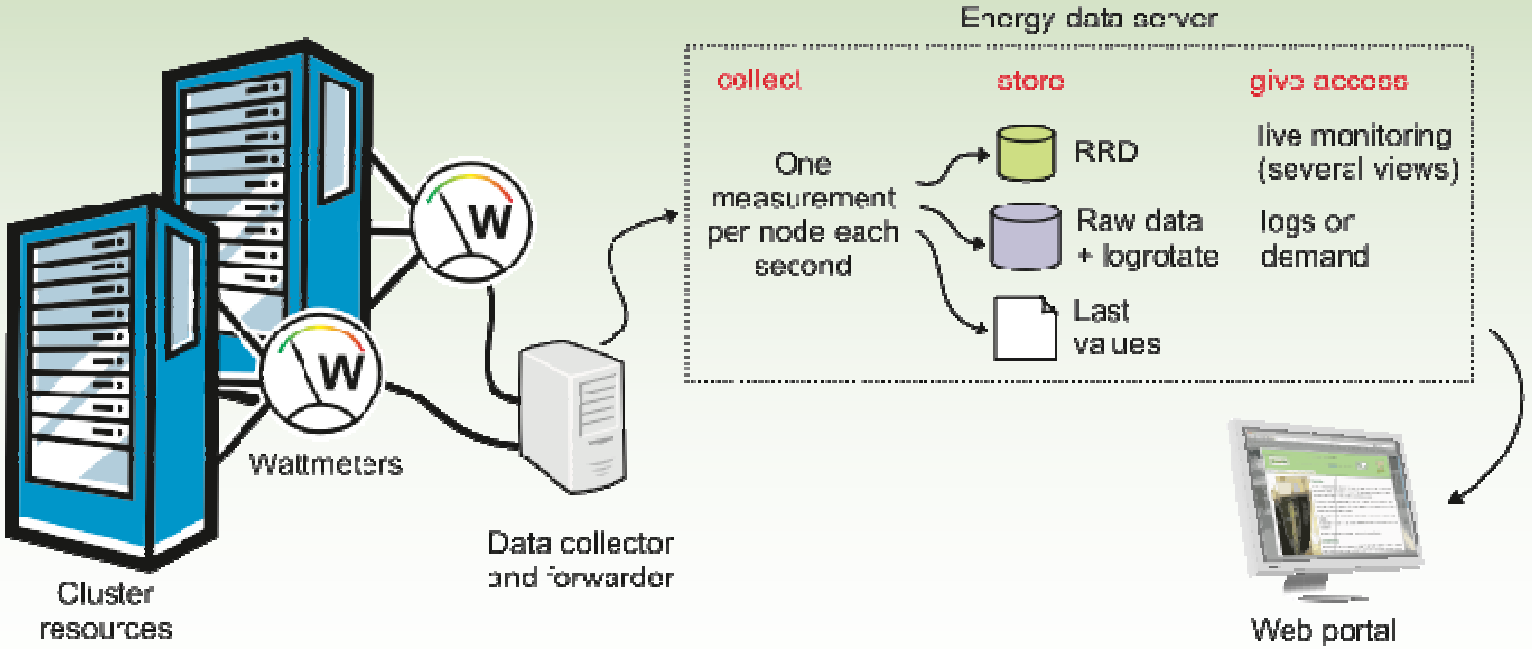

Electrical consumption / Usage

Periodicity of **energy measurements**:

One measurement per **second** for each equipment

*

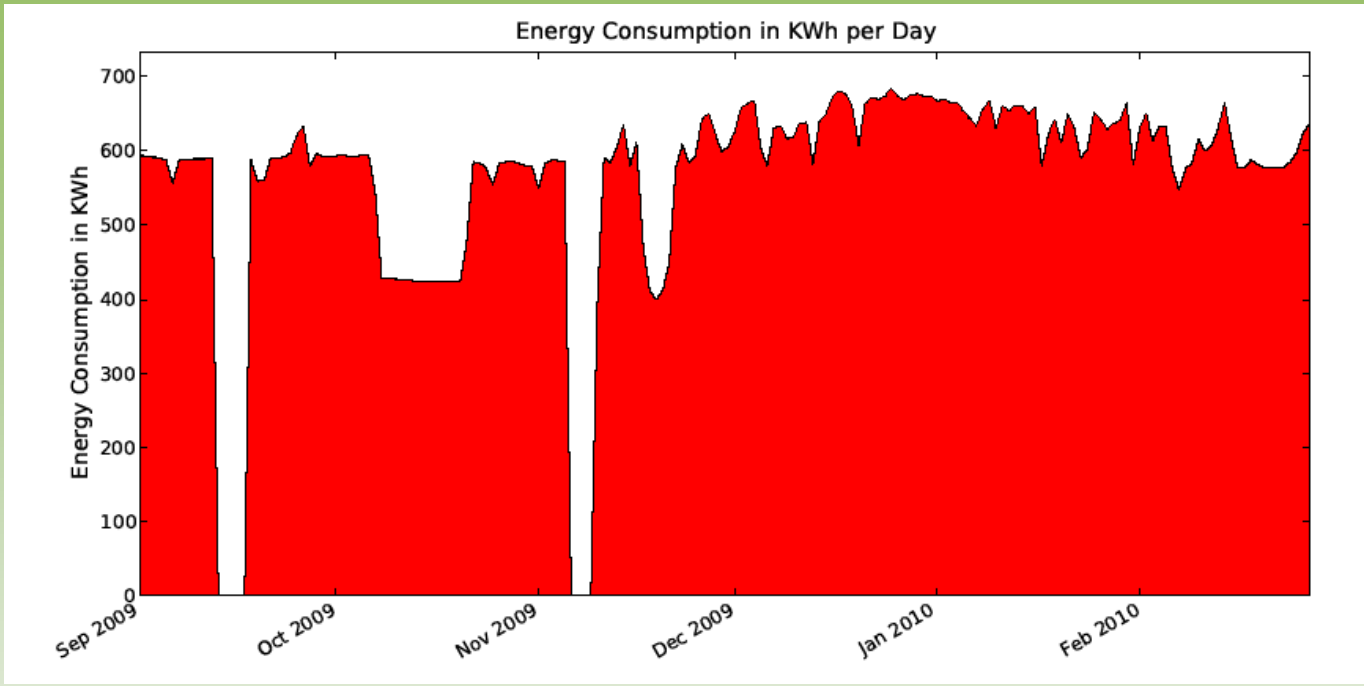

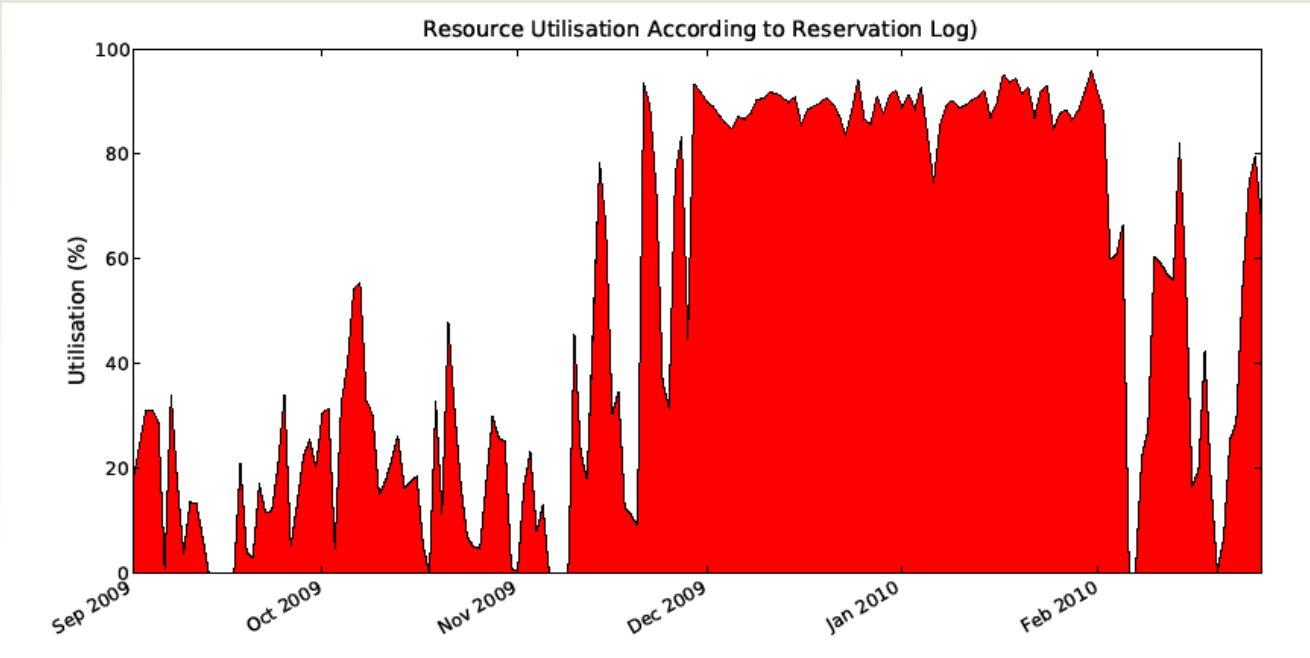

Example I : Profiling applications

Profiling the energy consumption of applications

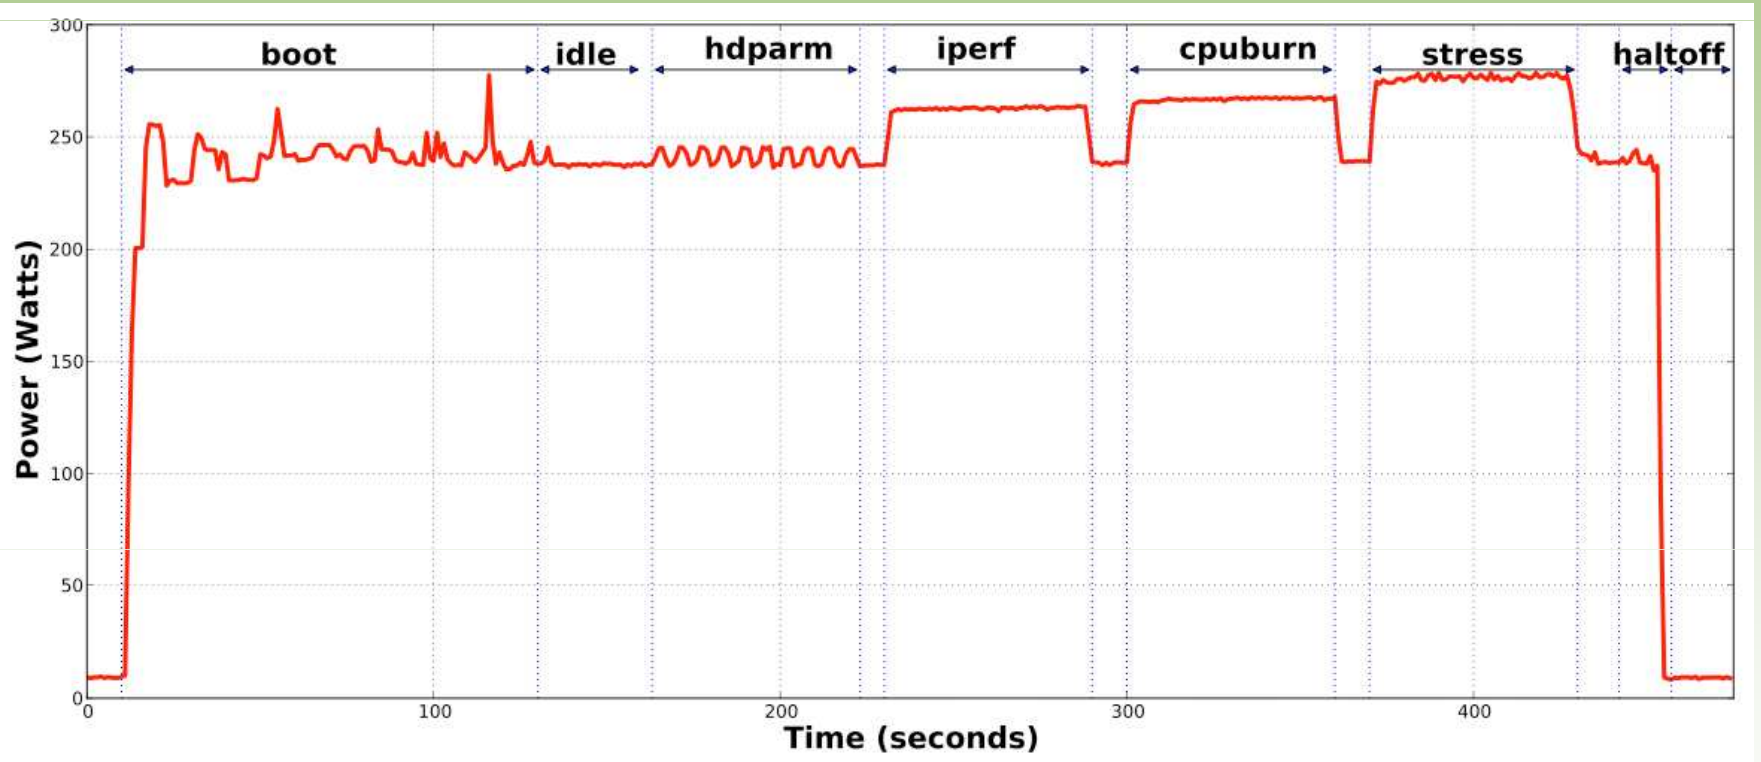

Example II : detecting anomalies

Improving frameworks/middleware and policies

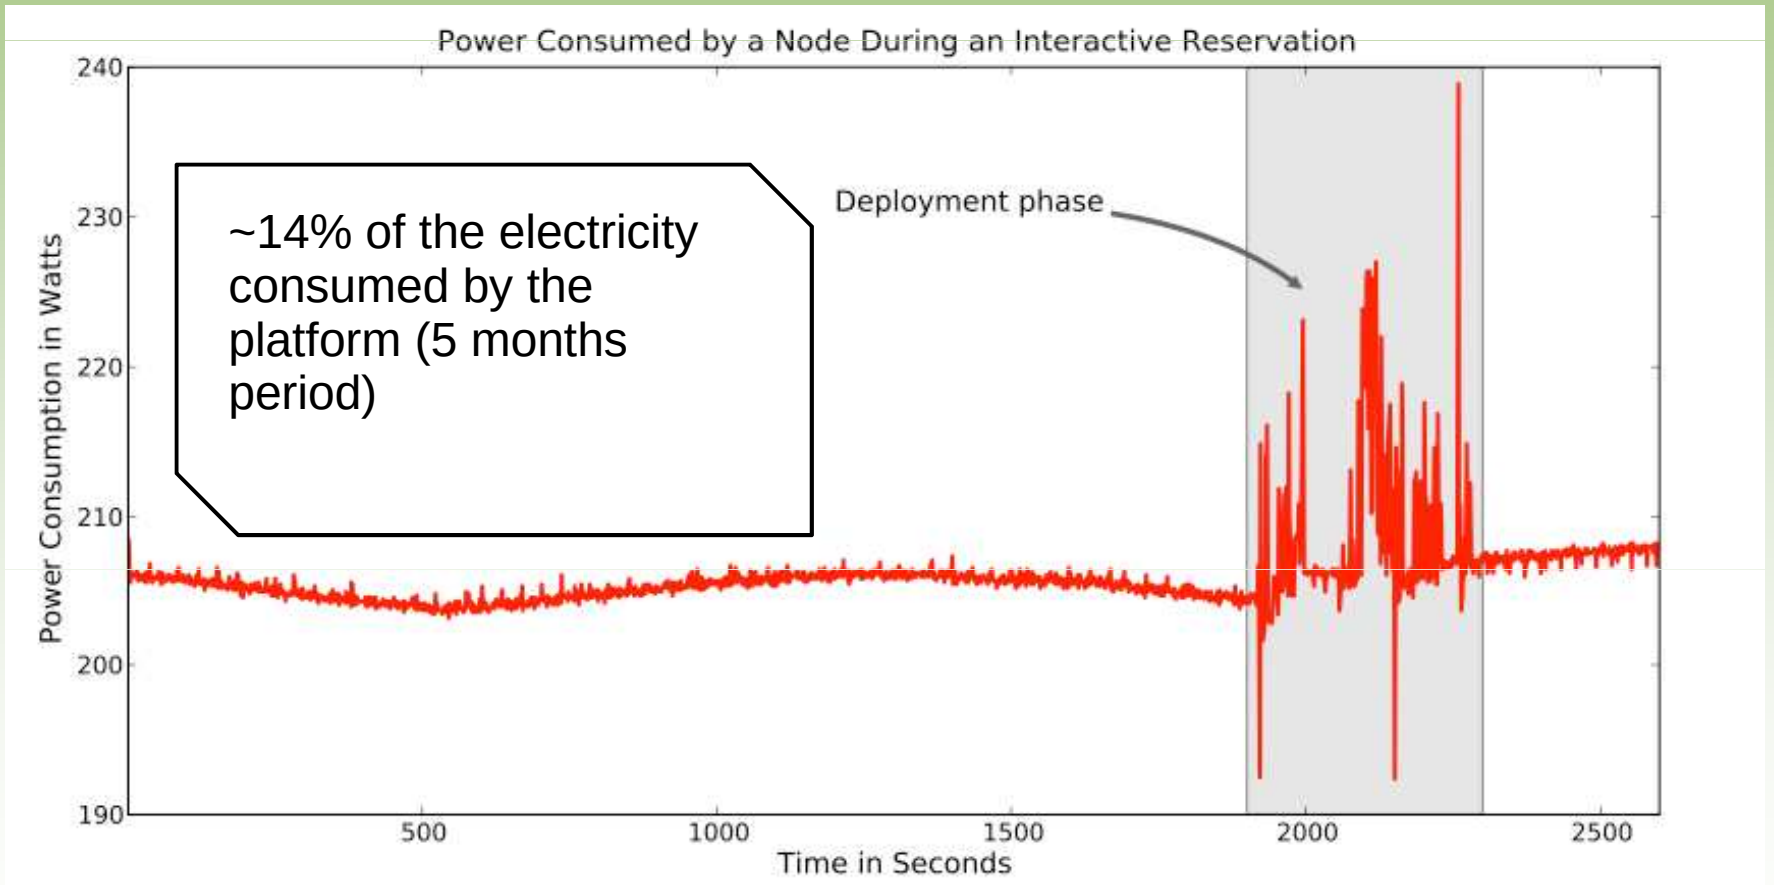

Example III : providing global views

Understanding the overall infrastructure

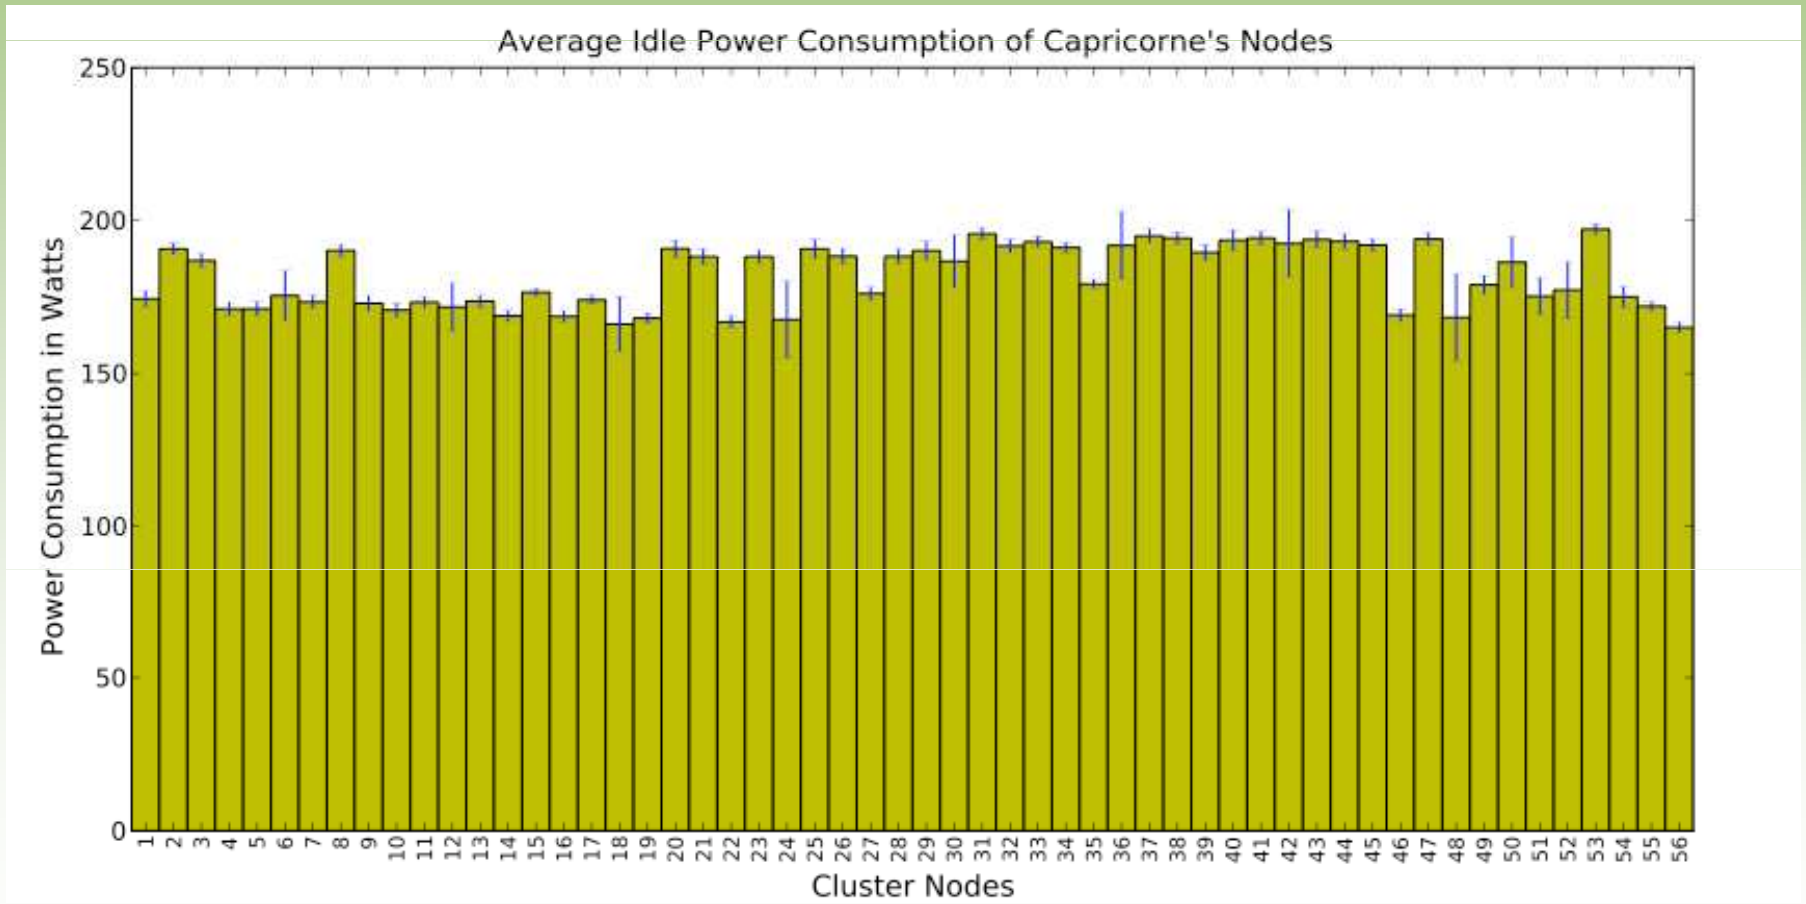

Large scale energy exposing

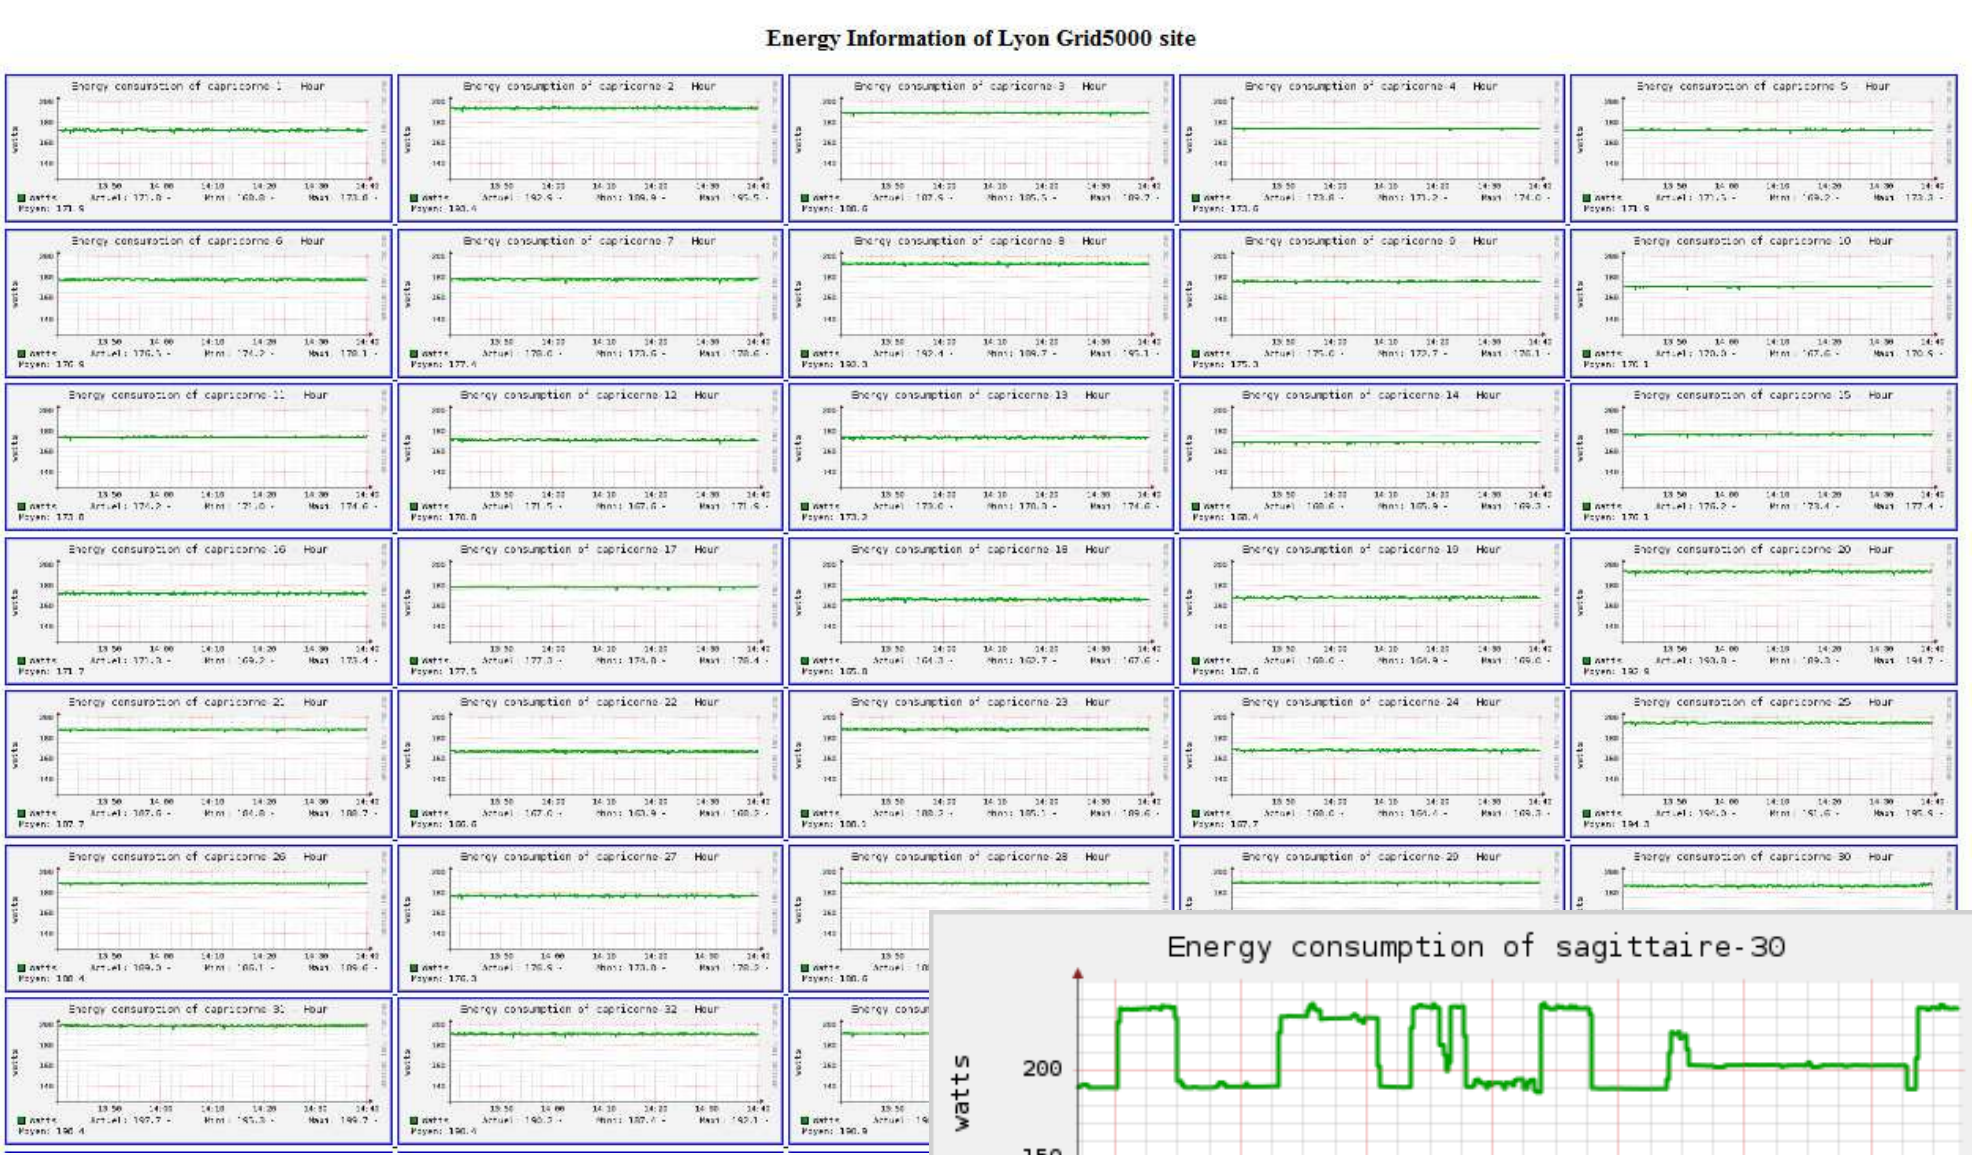

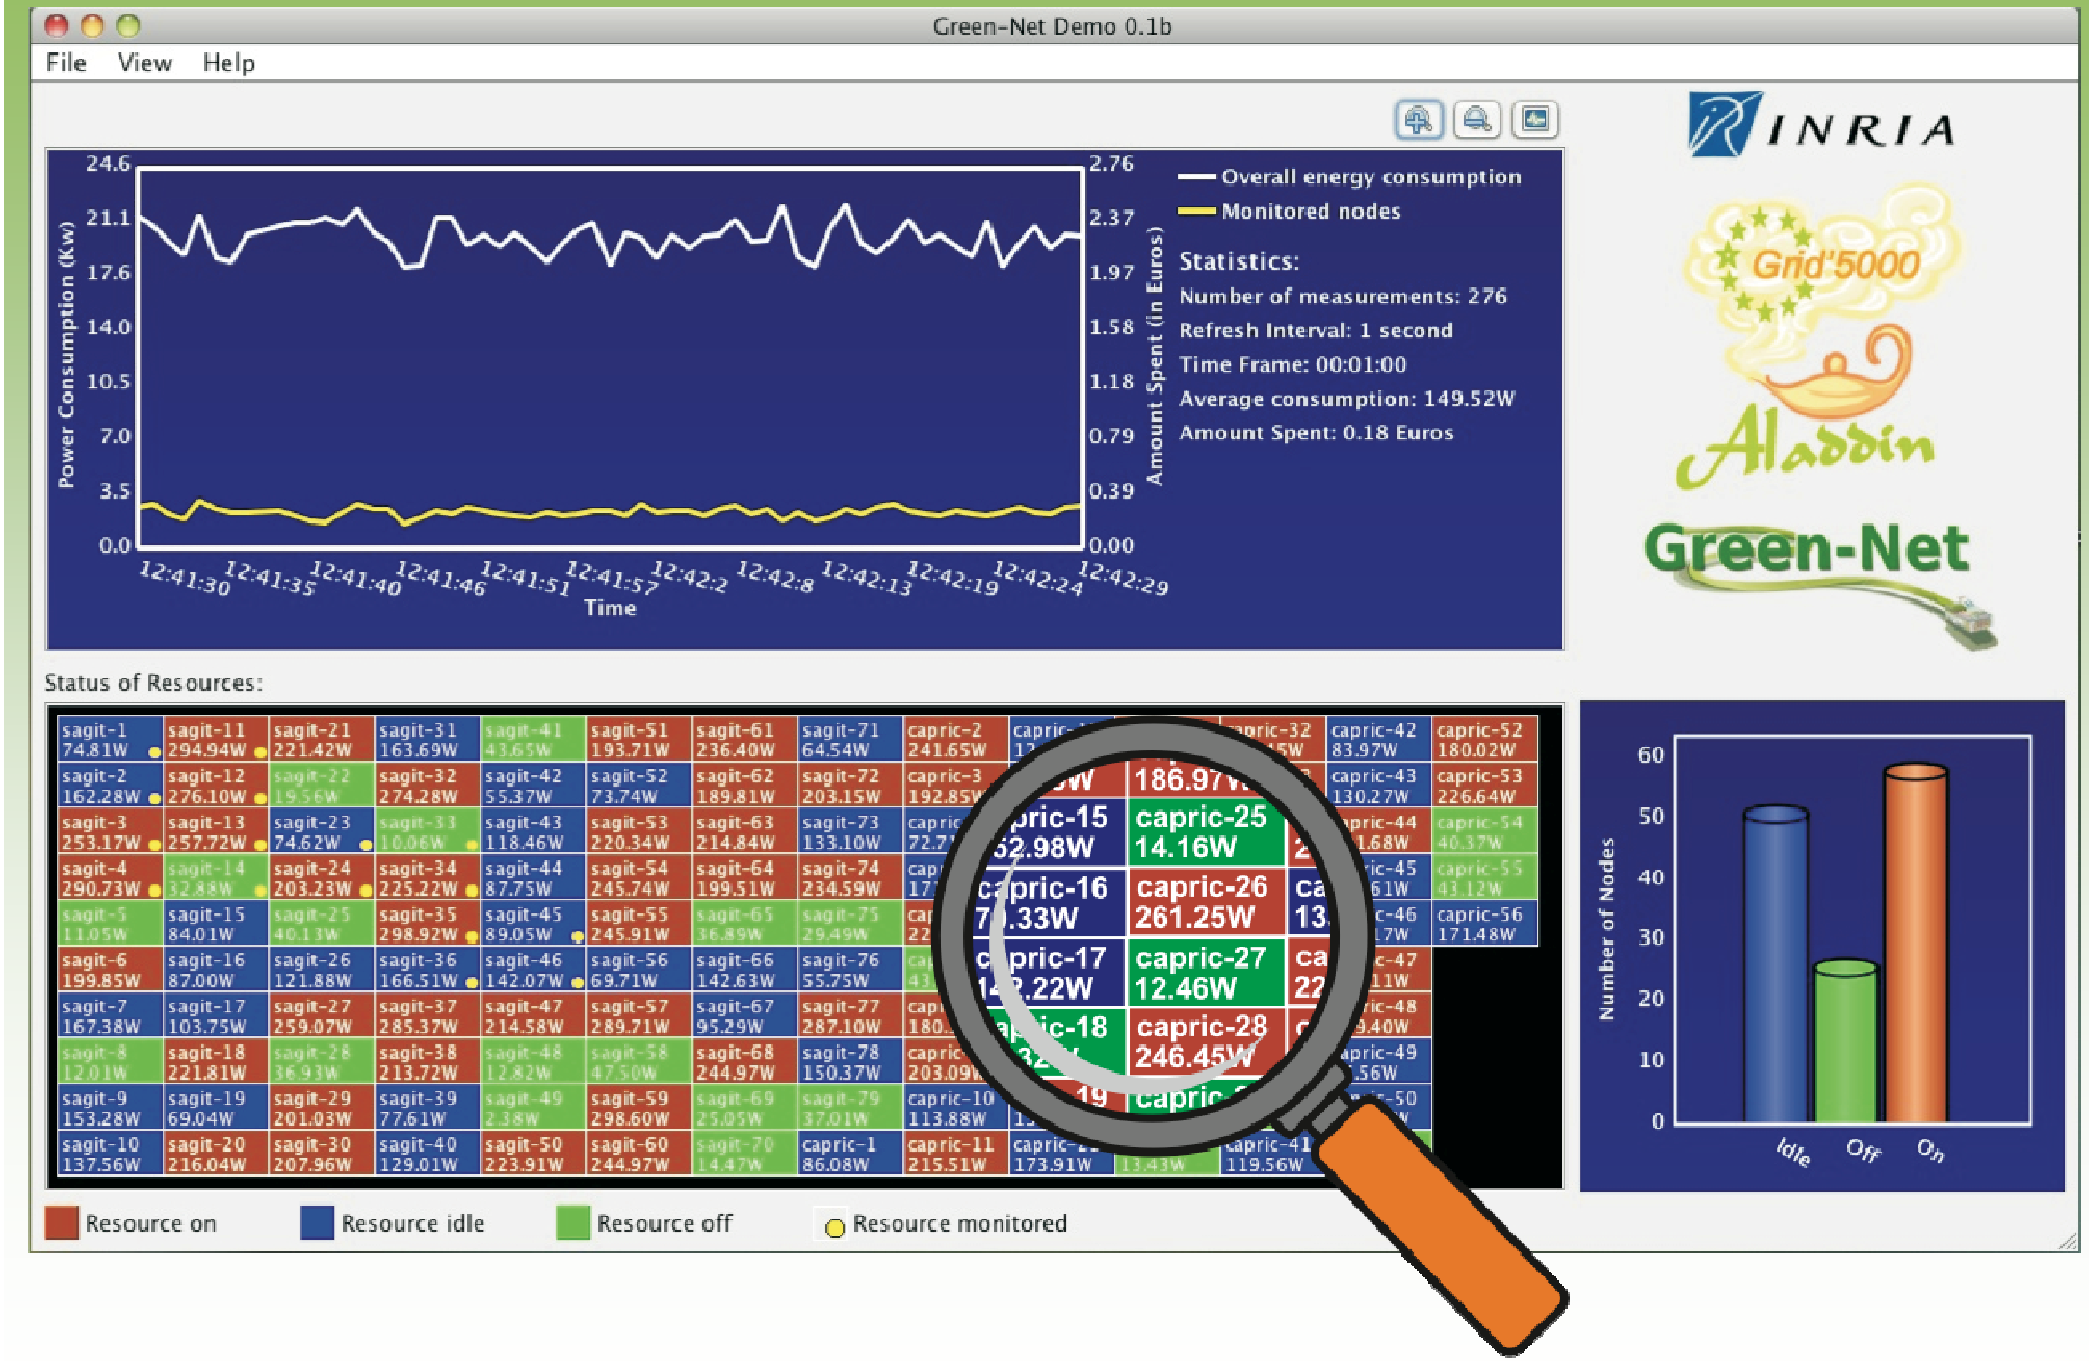

ShowWatts : M. Dias, J.P. Gelas, L. Lefèvre, A.C. Orgerie

Focus 2 : GOC

- Designing the Green Open Cloud architecture (GOC) based on the CEPAP model
- Supporting advanced features like live migration, tuning capping for agregating

Boot, Run and Halt

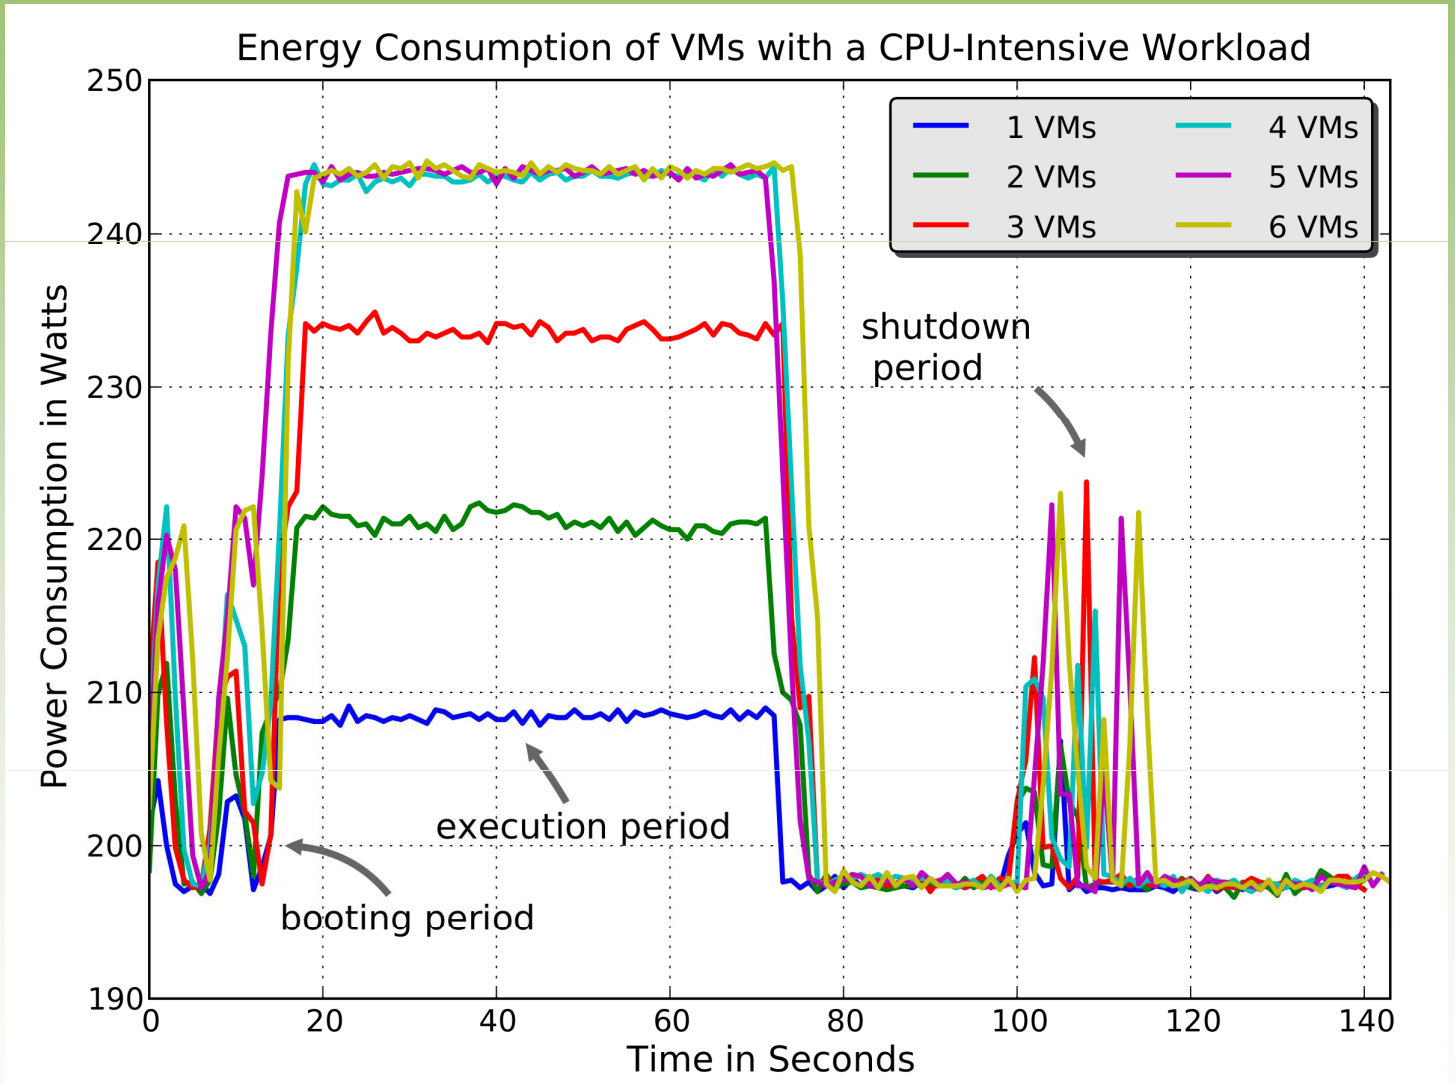

- 6% increase of energy with 1 VM running

Migration

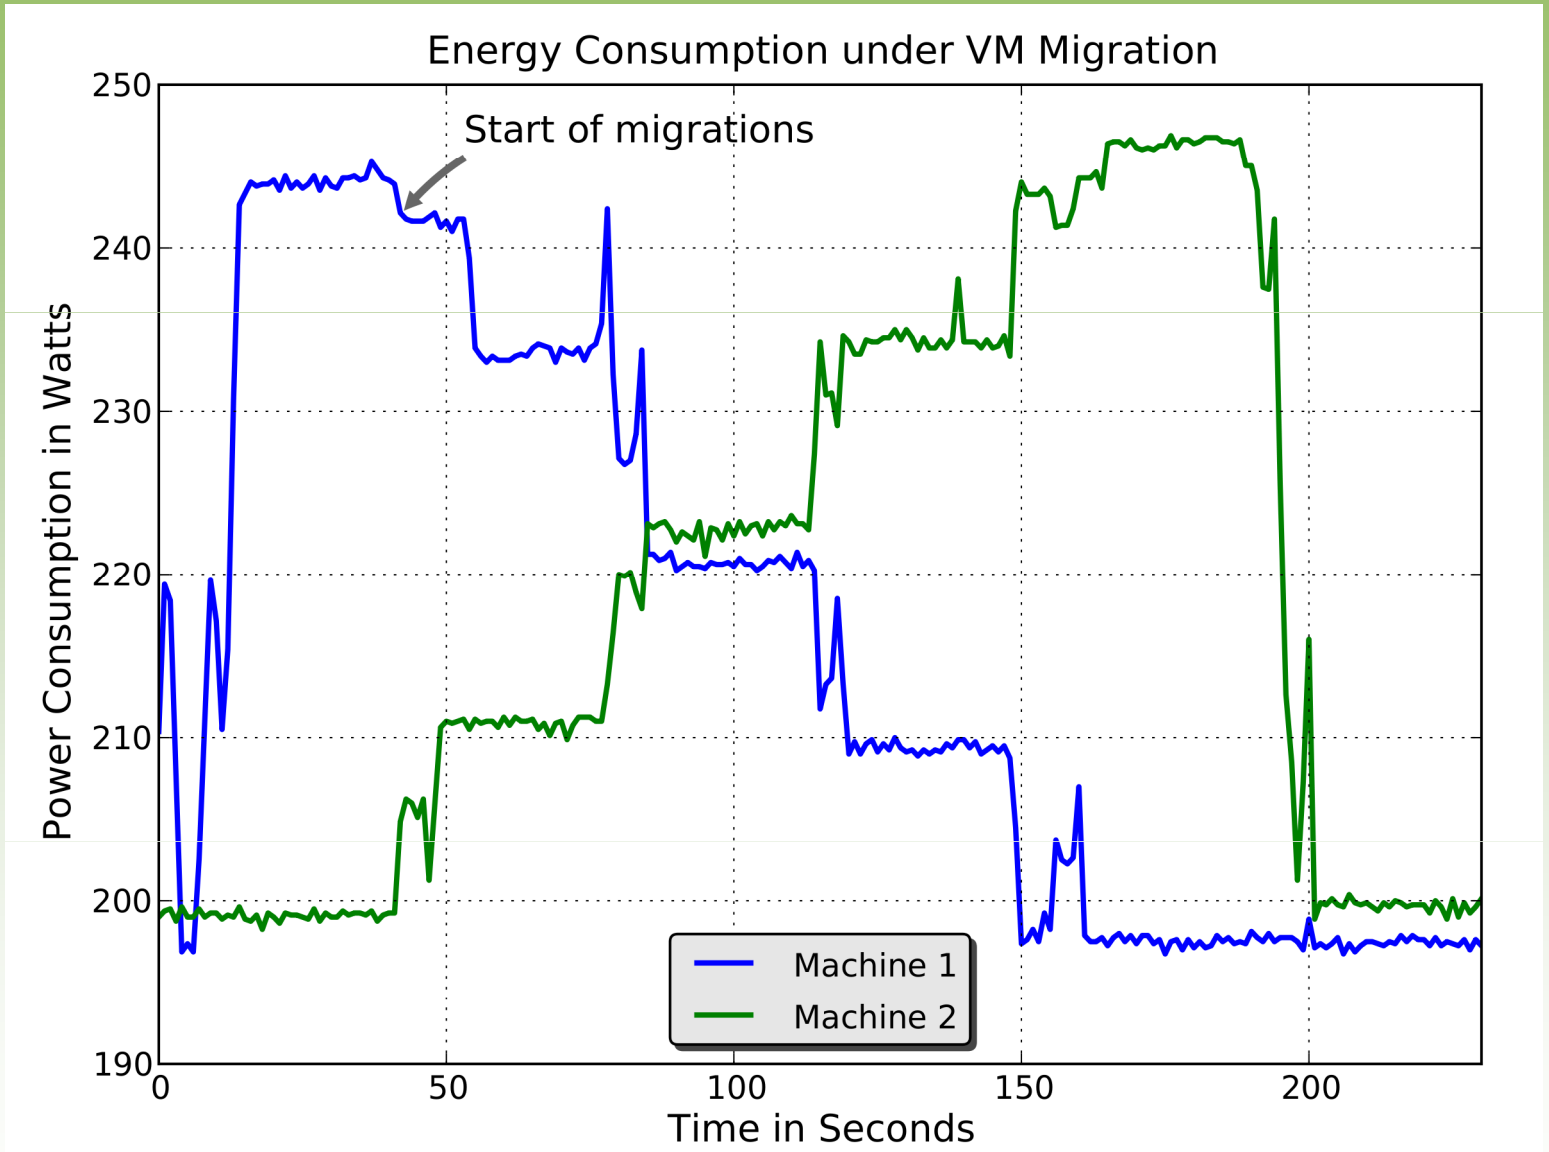

- Bad moment in energy during the migration

Load balancing tasks in VMs

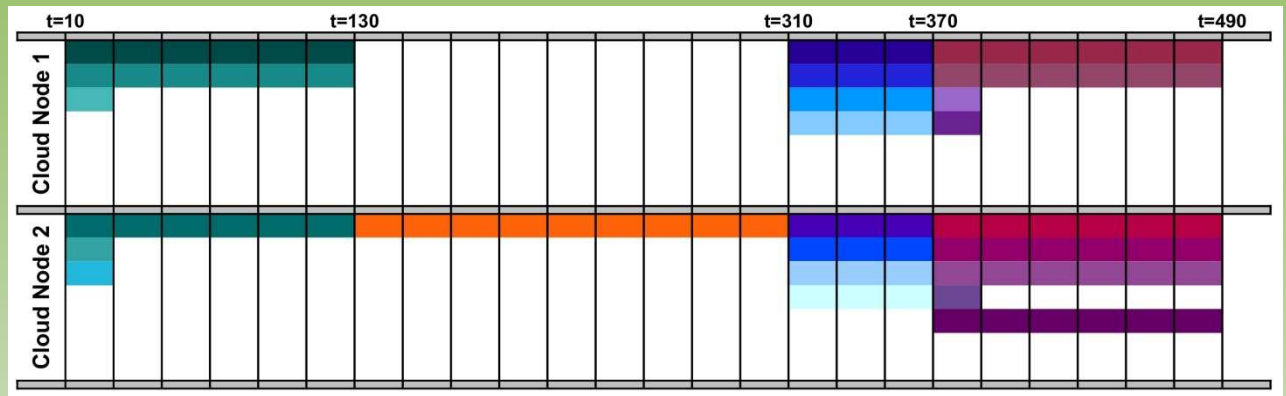

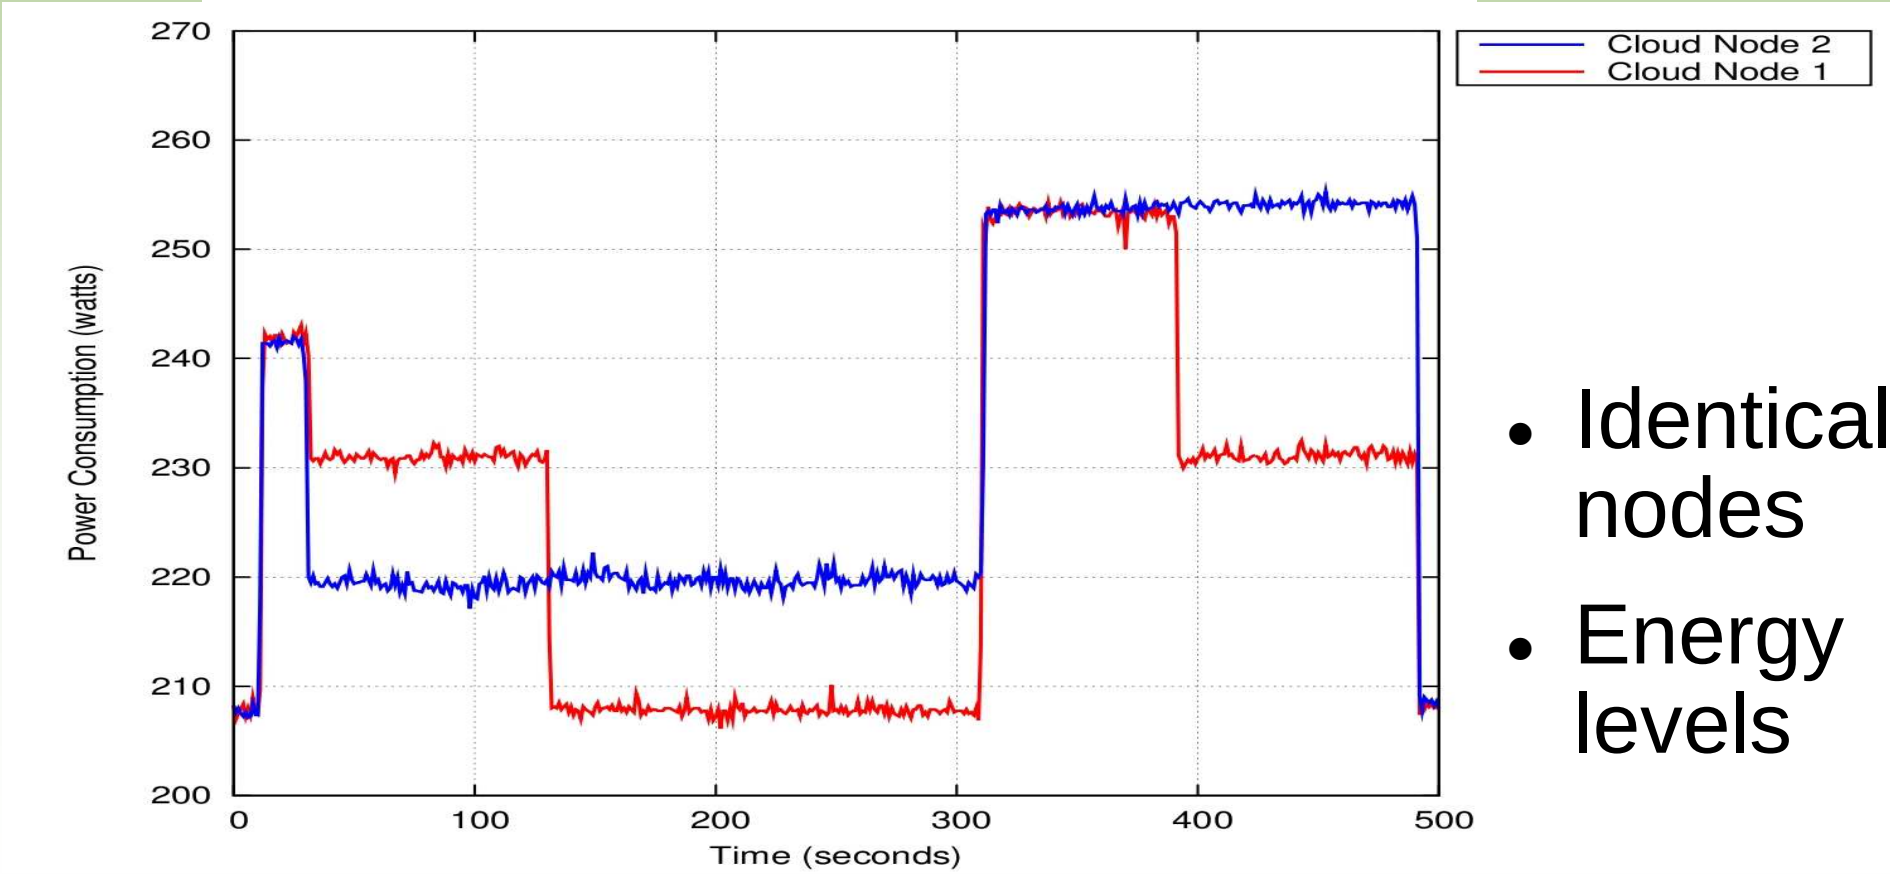

- Identical nodes
- Energy levels

Unbalanced with Green Scenario

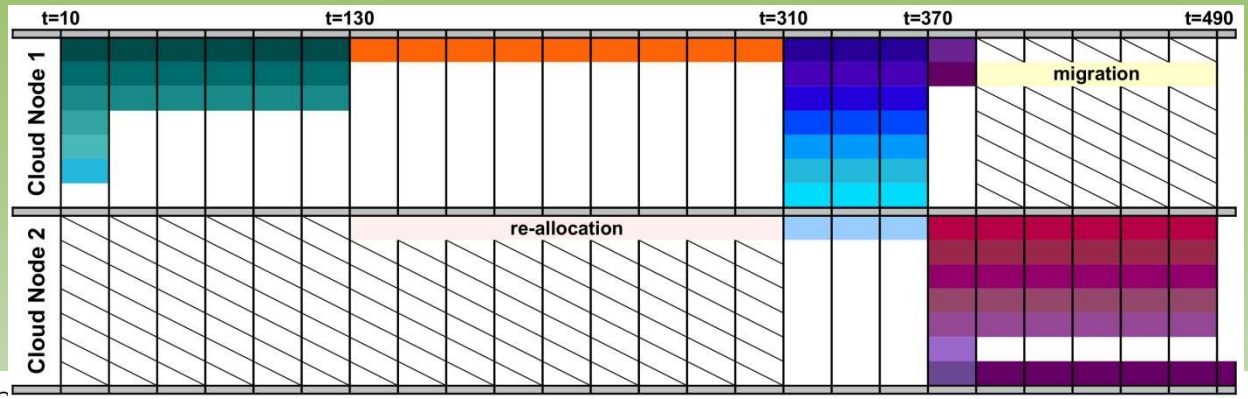

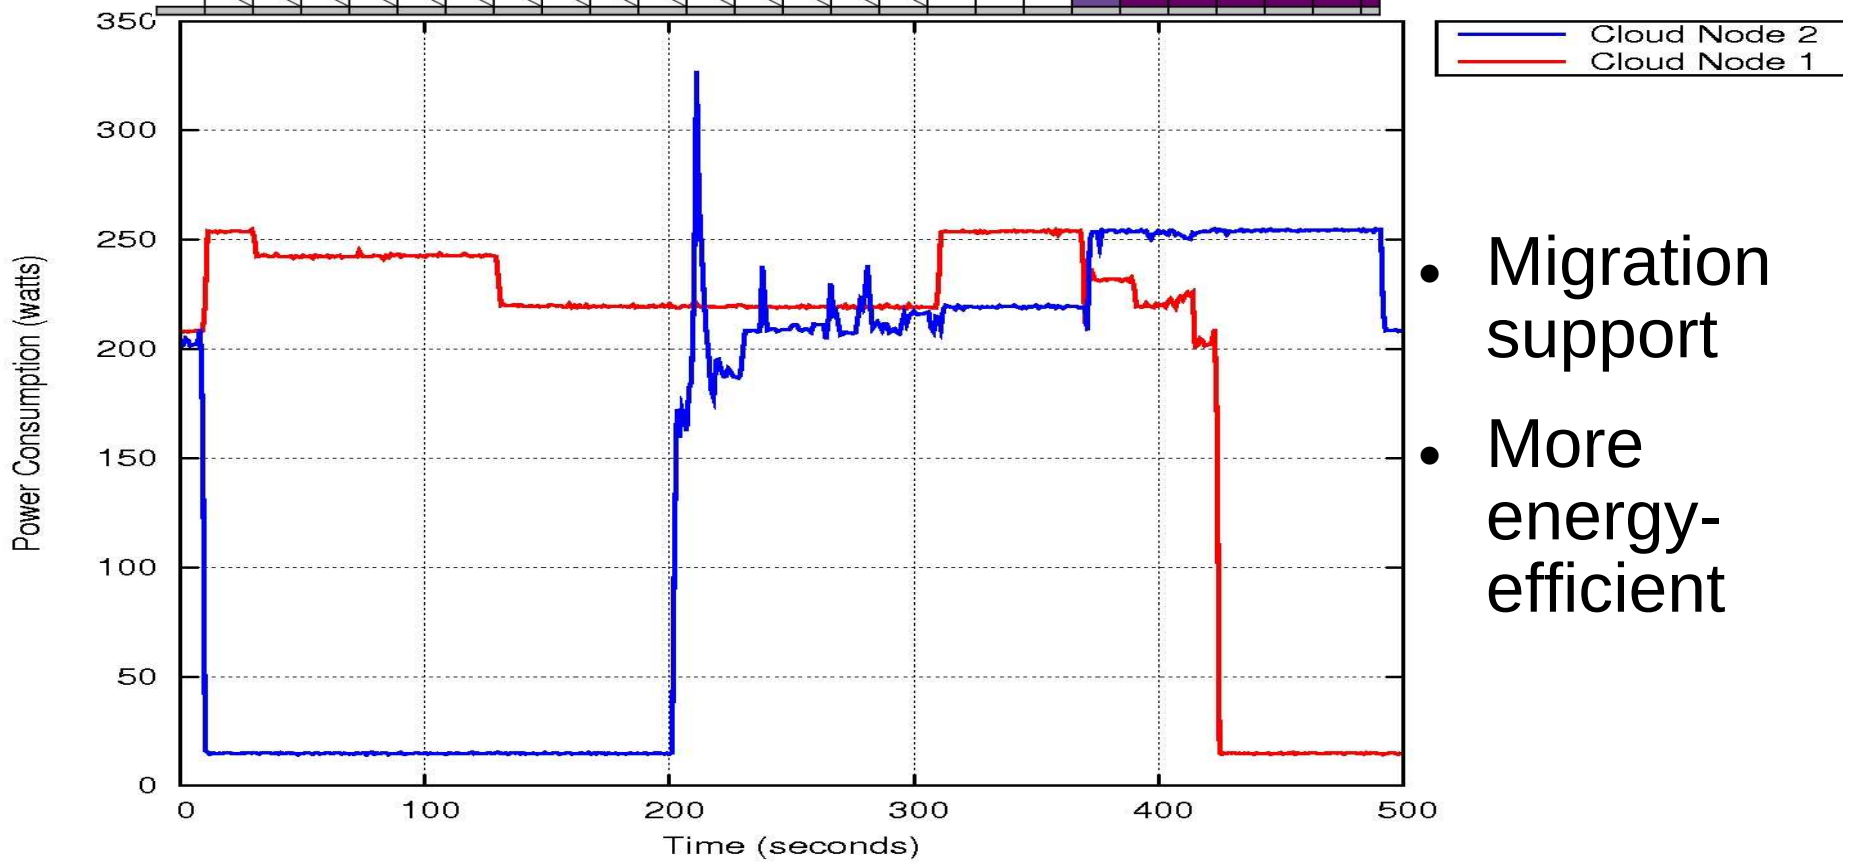

- Migration support
- More energy-efficient

Contributions and Perspectives

- Energy aspects change the way we design 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)
- Current works : energy efficiency in Green Wired Networks

Questions?

Laurent Lefèvre
laurent.lefevre@inria.fr

Thanks to A.C. Orgerie, M. Dias de Assuncao, J.P. Gelas