# ECO$_2$Clouds

Experimental Awareness of $CO_2$ in Federated Cloud Sourcing

## Newsletter September 2013

### New to ECO$_2$Clouds? Here is a a short summary...

The ECO$_2$Clouds project investigates strategies that can both ensure effective application deployment on the Cloud infrastructure and reduce energy consumption and $CO_2$ emissions. The need for novel deployment strategies becomes more evident when an application spans multiple Clouds.

ECO$_2$Clouds will provide a challenging and innovative approach to Cloud computing service delivery by:

- developing extensions and mechanisms for Cloud application programming interfaces to quantify their environmental impact;
- developing energy-efficient Cloud sourcing and application deployment strategies.

The carbon-aware mechanisms will be integrated into the FIRE facility BonFIRE.

During the past semester, the ECO$_2$Clouds partners have working on the following topics:

- Tuning of metrics for achieving energy efficiency and $CO_2$ awareness.
- Consolidation of green indicators.
- Implementation of the architecture proposed for energy efficient and $CO_2$ aware ECO2Clouds solution.
- Simulation/Implementation of monitoring, optimization and run-time adaptation techniques.
- Implementation of scheduling and data mining models

## *Research work focus of this issue:*

- **Models and Metrics**
- **Monitoring with Monitoring Collector**

### *Models and Metrics*

A key research area of ECO2Clouds concerns models and metrics of Cloud Computing. For this we look at a Federated Cloud Model and the associated Energy Metrics, taking into account infrastructure, virtualization and application levels.
We will now look at each of these items in more detail.

***1) Federated Cloud Model.***
The characteristics of the sites and resources available in each site represent the infrastructural layer. The virtualization layer contains the hypervisor and the VMs that it manages. VMs host (parts of) the applications used to manage the cloud infrastructure (e.g., the monitoring system) or offered to the final users (e.g., business application). The set of applications running on the VMs compose the application layer, where we consider different kinds of applications: siloed applications run on a single VM where all the resources required by the application are the resources that are reserved by the VM. On the other hand we consider federated applications modeled as business process (seen as composition of different tasks), where not only the composing tasks could run on different VMs, but a given task can also be reused in different business processes.

Website: eco2clouds.eu

Follow us on:

@eco2clouds

E-mail: eco2clouds@elet.polimi.it

SEVENTH FRAMEWORK PROGRAMME

Grant agreement no: 318048

ECO2Clouds

Experimental Awareness of $CO_2$ in Federated Cloud Sourcing

# Models and metrics (cont)

2) ***Energy Metrics***. On the basis of the defined model, we adopt a layered metric approach: application, infrastructural and virtualization layers, organized in a hierarchical structure.

A. Infrastructural layer metrics. These include the host layer and the site layer. The former focuses on the analysis of the behavior of the single host inside a site. The latter considers the whole cloud site providing an overall picture of its greenness.

- ***Host Layer Metrics*** include *CPU utilization (*load of a specific host): a relevant factor since a good application deployment strategy should avoid low host utilization. Energy efficiency of a host can be analyzed by observing CPU utilization together with information about the host energy consumption. Data about *I/O operations per second* are useful to understand the type of application running on a specific host. *Availability* is an index of the host reputation (the greater the availability the less the probability to have unsatisfied requests): an effective application deployment strategy should prefer highly available hosts.
- ***Site Layer Metrics***: These characterize the site activity and are needed to evaluate the most suitable cloud site for application deployment. Site selection depends on the site utilization and on storage (site congestions should be avoided). *Site availability* provides indications about the reliability of the site. To obtain greener application deployment, sites with high *Green Efficiency Coefficient (GEC)* are to be preferred. The *carbon footprint* measures the total amount of greenhouse gas (GHG) emitted during the full lifecycle of a product, service or system. It is calculated in terms of carbon dioxide equivalent (CO2e) using the relevant 100-year global warming potential The overall carbon footprint of a data centre includes:
  - ✓ the carbon embedded in the building hosting the centre, facilities and IT equipment;
  - ✓GHG emissions due to electricity and/or fuel consumption during the operational phase;
  - ✓ the final decommissioning of the centre and recycling or final disposal of all the materials.

  The relative contribution of the embedded carbon depends on the characteristics of the infrastructure, but for most data centres emissions associated with the operational phase are 2 to 3 orders of magnitude larger than the contribution of the infrastructure. Emissions during operation are due to the energy consumption of the IT equipment and of the cooling systems.

The carbon footprint of electricity consumption depends on the energy production mix (the proportion of different power generation technologies) and the efficiency of the distribution grid. The energy mix is the most significant factor affecting GHG emissions.

A general agreement is still lacking on the emission factors of nuclear power. Anyway, strategies for reducing $CO_2$ in a federated cloud system include the increase of energy efficiency and fuel switching in our approach, the evaluation of the $CO_2$ emissions is based on the emission factors (gCO2e/kWh) provided by the national grids. Assuming that we know the average power consumption (AP) for a specific site the energy consumed in one year of operating site:
Energy = APx365daysx24hours. Multiplying the energy consumed by the emission factor, it is possible to quantify $CO_2$ emissions.

B. Virtualization layer metrics. These characterize VMs on which the applications are running. Information about the usage of VM resource has to be analysed in order to evaluate if the current deployment can be further improved and thus optimized. The analysis of VM energy consumption aims to understand how the energy consumed by the host is distributed among the VMs deployed on it.

C. Application layer metrics. We have defined new metrics inspired by the data center metrics proposed recently, especially by the *Green Grid Consortium*, redefining e.g. PUE and DCeP at application level to measure the impact of an application in terms of energy consumptions and carbon emissions. The original PUE (comparing power used by the entire infrastructure divided by power used by the IT infrastructure) has been modified defining *A-PUE (Application PUE)* which compares the total amount of power required by the VM with the power used to execute the application tasks. The *Data Centre Energy Productivity (DCeP)* - ratio between the work output of the data centre and the total energy for the data center - is redefined at the application level as the *Application Energy Productivity (AEP)*.

Website: eco2clouds.eu
Follow us on:

@eco2clouds

E-mail: eco2clouds@elet.polimi.it

SEVENTH FRAMEWORK PROGRAMME

Grant agreement no: 318048

ECO2Clouds

Experimental Awareness of $CO_2$ in Federated Cloud Sourcing

## Metrics (cont)

we are also interested in the greenness of the application in terms how much green energy is used to run the application Ai. To do this, we consider the Green Efficiency Coefficient (GEC) factor and we multiply it to the energy consumed by the virtual machine

3) ***Metrics Assessment Infrastructure***. The monitoring infrastructure allows retrieval of most of the data needed to assess our metrics. Additional data provided by the cloud scheduler are needed to compute the availability at the infrastructure level and to collect details about the start/ completion of the execution of tasks for assessment of some of the application layer metrics. Energy contracts have to be analyzed to define the *Green Energy Coefficient*. The figure below presents a simplified view of the architecture needed for metric assessment (bidirectional arrows denote that the metrics at the various levels are used to calculate other metrics). The figure presents a simplified view of the architecture needed for metric assessment (bidirectional arrows denote that the metrics at the various levels are used to calculate other metrics).

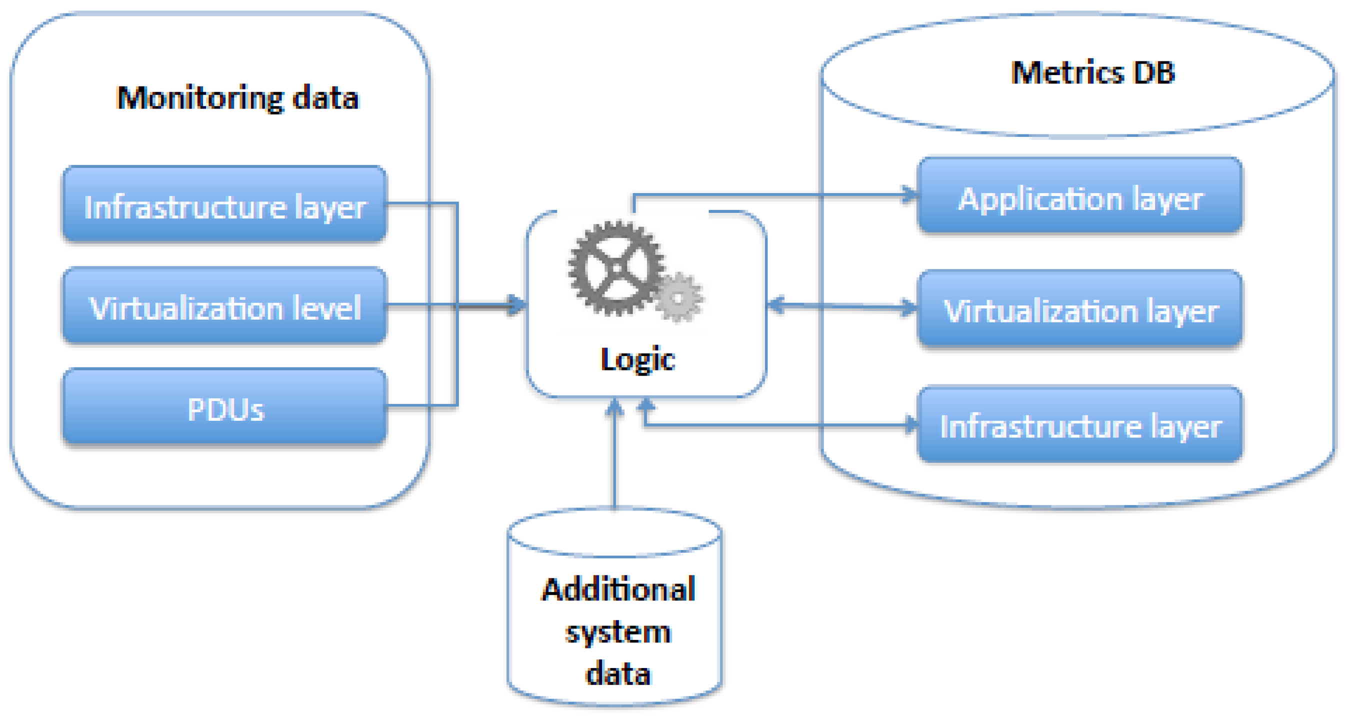


**Monitoring Data and Metrics DB**

## Monitoring

The latest choices undertaken in ECO2Clouds for the monitoring infrastructure are addressed by using the Zabbix monitoring framework, providing in detail the monitoring and data analysis environment and presenting how the approach is being implemented (see Figure on the right).

The focus is on being able through Monitoring to derive metrics referring to the interrelation between different components of the IT cloud infrastructure. Monitoring deals with three-layers metrics as follows:

- Tuning the metric definition.
- Implementing the metrics considered as relevant for the application, VM and infrastructure layers, including existing metrics and refinements of metric definitions (to address the novel challenges of ECO2Clouds).
- Defining how metrics will be used in the application profile (introduced in D2.1 "ECO2Clouds Requirements") and describing the developments done on the "application profile" concept.
- Defining how the monitoring infrastructure is used to gather the data required for metric assessment.
- Implementing energy metrics to compute the energy consumption of VMs based on physical host parameters and the energy consumption measured by the PDUs.

The monitoring architecture considers the three layers and, additionally, the Power Distribution Units (PDU) used for gathering energy related metrics.

For the monitoring system, Zabbix is deployed on a VM running on a server within the given site in order to monitor physical and energy related data. For the virtual hosts, the framework foresees an additional virtual aggregator which connects to the infrastructure aggregators to retrieve all the related parameters. Both types of aggregators rely on Zabbix and are integrated in the monitoring infrastructure.

The *Accounting Service (AS)* contacts the Zabbix servers at the distributed locations and retrieves the current monitoring information for the infrastructure and energy related metrics. Furthermore, the monitoring information (virtual and application) captured by the virtualized Zabbix aggregator is retrieved and stored in a MySQL database backend. This process allows gathering all available metrics together and guarantees that no parameters will be lost after the execution of an experiment. In order to keep the amount of stored data as small as possible, a data mining approach is implemented for the AS.

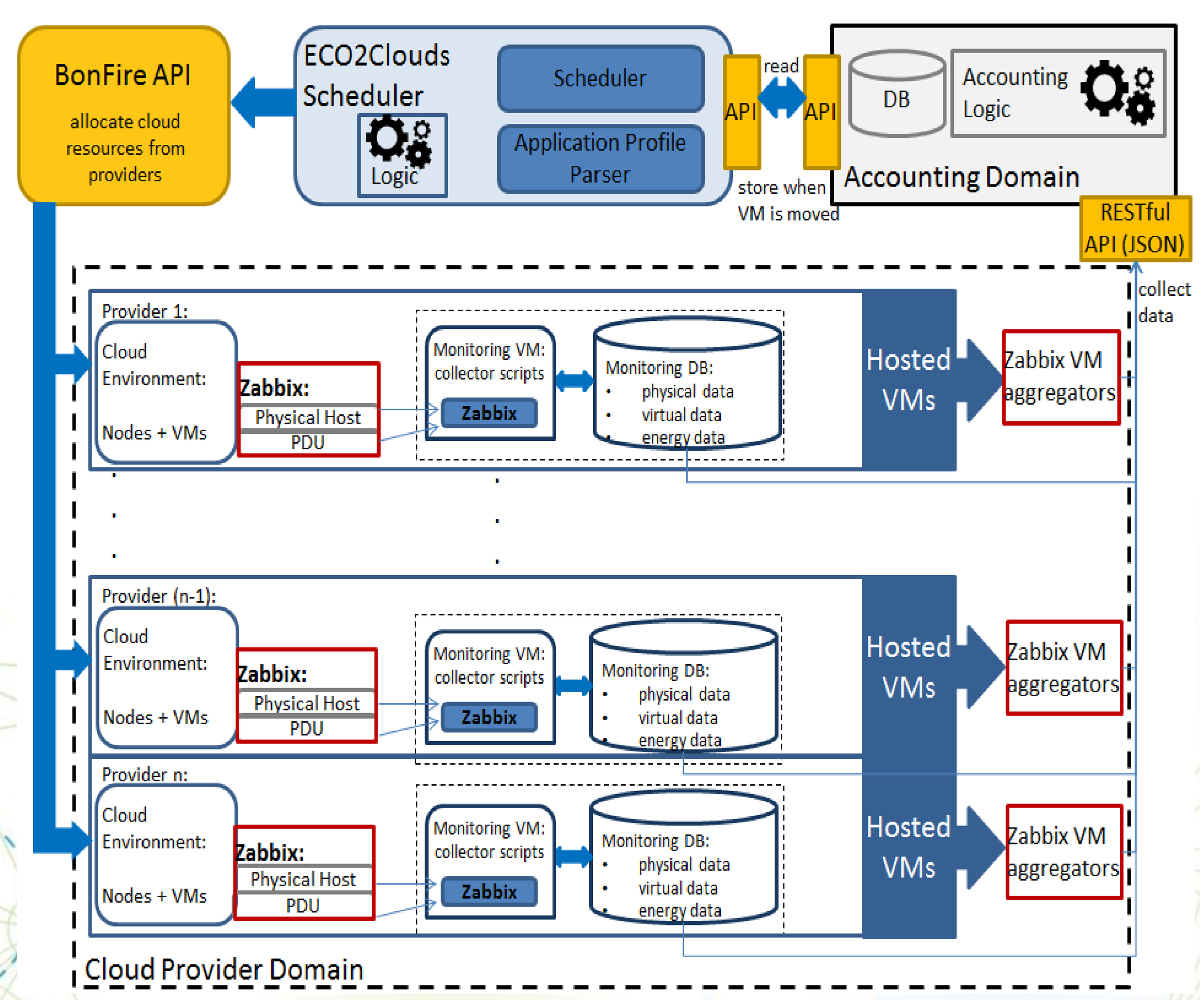


**$ECO_2$Clouds Monitoring Infrastructure**


Website: eco2clouds.eu
Follow us on:

@eco2clouds

E-mail: eco2clouds@elet.polimi.it



SEVENTH FRAMEWORK PROGRAMME

Grant agreement no: 318048

ECO2Clouds

Experimental Awareness of $CO_2$ in Federated Cloud Sourcing

# A practical example

## Some results from the testbed

This screenshot from infrastructure monitoring at Inria's testbed illustrates how variable $CO_2$ cost of energy can be. This data, taken from the live feed from France's electricity transport company (RTE) is made available to the higher levels of the ECO2Clouds stack, as well as to users to as to enable optimization strategies.

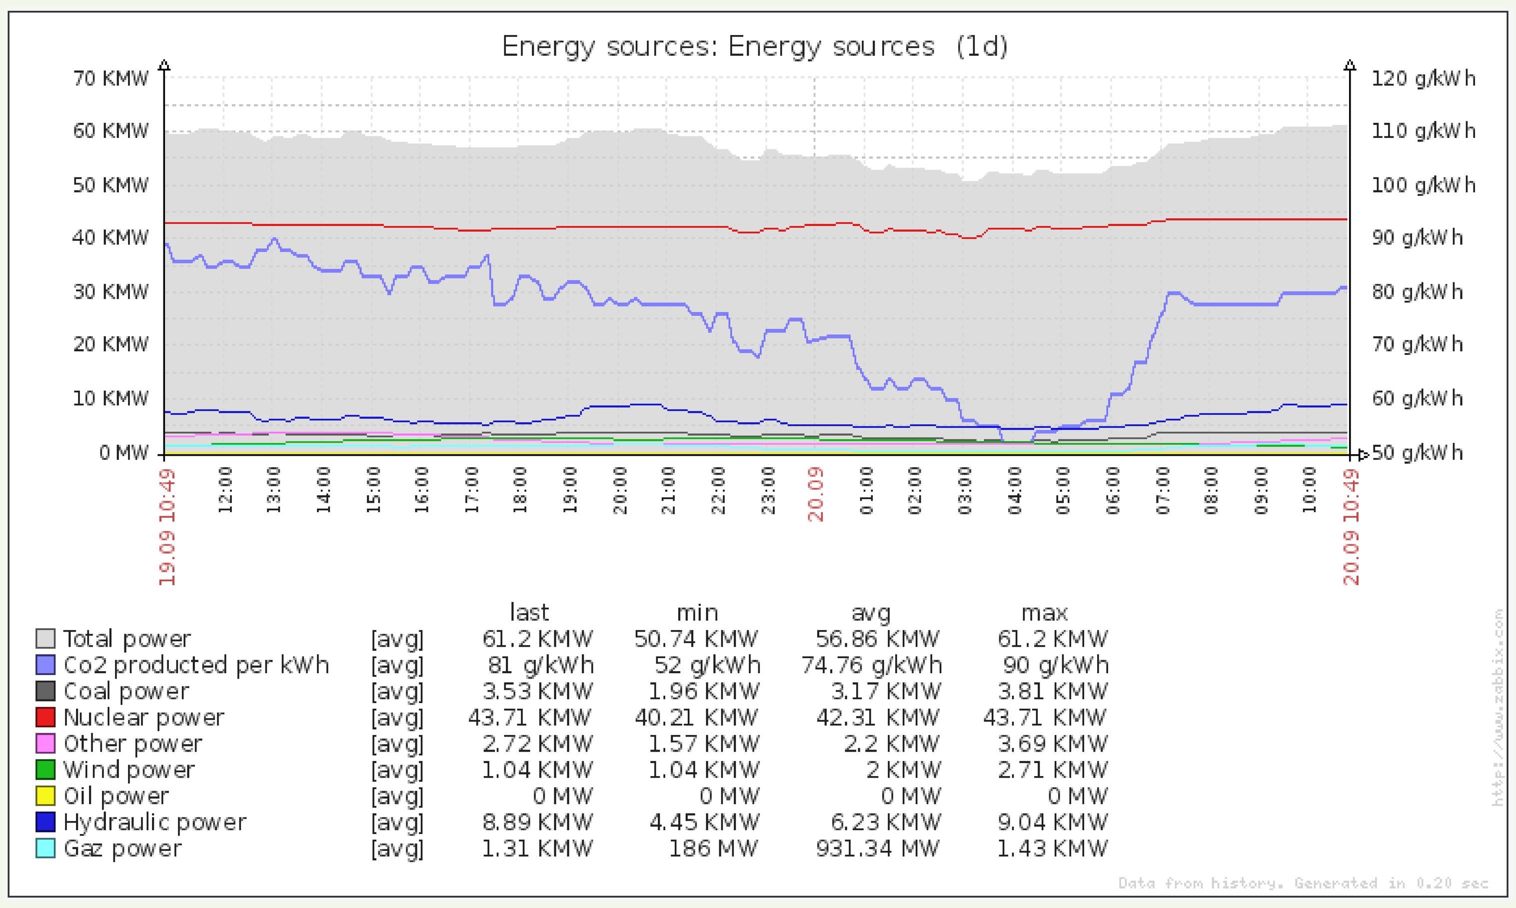


This second screenshot illustrates the variability of power consumption on infrastructure nodes according to the number of VMs running (green line) and their activity (red line). It is by combining the data from these two graphs (and more) that CO2 costs can be attributed to running VMs.

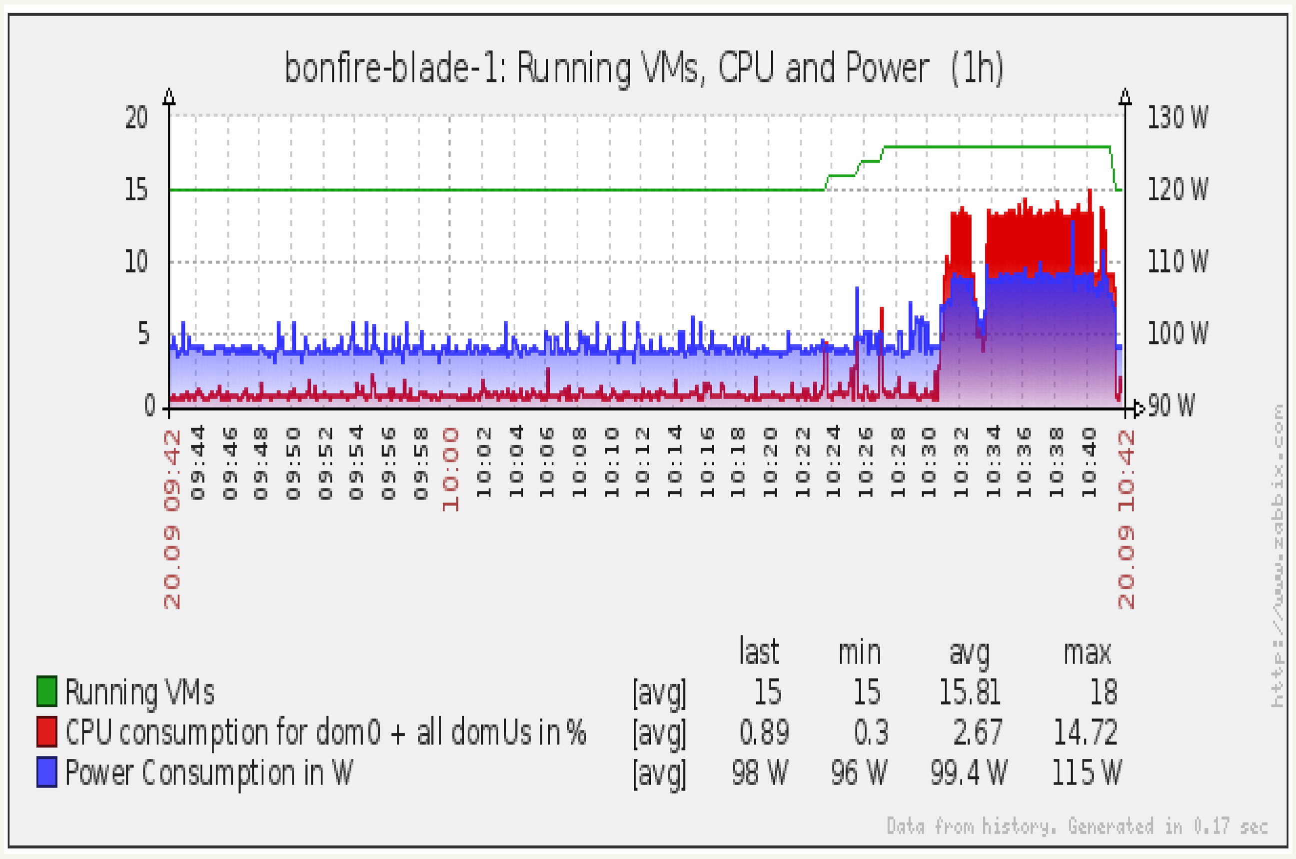



Website: eco2clouds.eu
Follow us on:

@eco2clouds

E-mail: eco2clouds@elet.polimi.it



SEVENTH FRAMEWORK PROGRAMME

Grant agreement no: 318048

ECO$_2$Clouds

Experimental Awareness of $CO_2$

## Publications

1. Cloud Federation Monitoring for an Improved Eco-Efficiency
Axel Tenschert, Michael Gienger, *eChallenges 2013*, October 2013, Dublin, Ireland.
*Description of the monitoring infrastructure for eco-efficient cloud federations and metrics, based on the Zabbix monitoring framework and user-defined metrics.*

2. Energy Efficient and $CO_2$ Aware Cloud Computing: Requirements and Case Study
Usman Wajid, Barbara Pernici, Gareth Francis, Special Session at *IEEE Int. Conf. On Systems, Man and Cybernetics (SMC*), October 13-16, Manchester, UK
*The requirements of the system for monitoring green IT metrics and reducing the environmental impact of applications are illustrated*

3. Monitoring and Assessing Energy Consumption and $CO_2$ Emissions in Cloud-based Systems
Cinzia Cappiello, Sumit Datre, MariaGrazia Fugini, Michael Gienger, Paco Melia', Barbara Pernici *Special Session at IEEE Int. Conf. On Systems, Man and Cybernetics (SMC*), October 13-16, Manchester, UK.
*Green Performance Indicators and environmental impact metrics for ECO2Clouds are defined, as well as the necessary monitoring infrastructure.*

4. Optimizing Energy Efficiency in the Cloud Using Service Composition and Runtime Adaptation Techniques
Usman Wajid, Cesar Marin, Anthony Karageorgos, *Special Session at IEEE Int. Conf. On Systems, Man and Cybernetics (SMC*), October 13-16, Manchester, UK.
*A dynamic approach to optimizing energy efficiency in cloud services is proposed*

5. Energy-Aware Process Design Optimization
Cinzia Cappiello, Pierluigi Plebani, and Monica Vitali, *EuroEcoDC Workshop*, September 2013, Karlsruhe.
*To increase data center energy efficiency, focusing at the application level, this paper proposes a method to support the process design optimizing configuration and deployment.*

6. Improving Energy Efficiency in Data Centers and federated Cloud Environments
Eugen Volk, Axel Tenschert, Michael Gienger, Ariel Oleksiak, Laura Sisó and Jaume Salom, *EuroEcoDC Workshop*, September 2013, Karlsruhe.
*Two complementary energy-efficiency optimization approaches covered in the scope of EU projects are presented: CoolEmAll – with focus on building energy efficient data centers, and ECO$_2$Clouds – with focus on energy-efficient cloud-application deployment in federated cloud-environments.*

7. Using Intelligent Agents to Discover Energy Saving Opportunities within Data Centers
Alexandre Mello Ferreira and Barbara Pernici, *Second International Workshop on Requirements Engineering for Sustainable Systems (RE4SuSy)*, July 15th, 2013 in Rio, Brazil.
*The problem of analyzing events related to violations of thresholds of Green Performance Indicators is discussed.*

## ECO$_2$Clouds social life

***Past Events (2013)***

- April 2013 Participation in COST 804 School on Energy efficiency in large scale distributed systems.
- 22-Jul-2013 Presentation of the project at the Cloud Communities Day organized by Eurocloud Italy.

***Forthcoming Events***

- EuroEcoDC Worksop
  ECO2Clouds is one of the EU projects co-organizing the EuroEcoDC (E2DC) Workshop co-located with the 3rd International Conference on Cloud and Green Computing (CGC), Sept. 30–Oct. 2, 2013,Karlsruhe,Germany.
  Webpage:http://www.all4green-project.eu/EuroEcoDC
- ECO2Clouds organizes the Energy Efficient Systems Special Session at IEEE International Conference on Systems, Man and Cybernetics (SMC), October 13–16, 2013, Manchester, UK
  Webpage: http://eco2clouds.eu/ees/
- ECO2Clouds related workshop, eChallenges 2013 (Organized by HRLS).

Website: eco2clouds.eu
Follow us on: @eco2clouds
E-mail: eco2clouds@elet.polimi.it

SEVENTH FRAMEWORK PROGRAMME

Grant agreement no: 318048

ECO$_2$Clouds

Experimental Awareness of $CO_2$ in Federated Cloud Sourcing

## Ongoing work

### Our current work

- Use case –based testing of the approach on "eels life cycle" application and on the scientific applications in health environments.
- Testing of the adaptation mechanisms at the infrastructure at site, host and VM level.
- Runtime adaptation techniques to ensure efficient resource utilization and minimum CO2 footprint of applications: testing our approach on deployment of use cases on the clouds.
- Energy-aware deployment strategies at design and at run time using BonFIRE as a test environment.

In the next newsletter (March 2014), we plan to explain further about

- Monitoring collector
- Optimization and Runtime adaptation
- Datamining

## Consortium

**ATOS**, Spain

Atos is an international company focused on Services for Information Technologies (IT) and currently with headquarters located in Paris, with presence in 42 countries.

**University of Manchester**, UK

The University of Manchester (UNIMAN) is the largest campus based university in UK with close links to industry.

**HLRS**, Germany

The High Performance Computing Center is a research and service institution affiliated to the University of Stuttgart. It has been the first national supercomputing center in Germany offering services to academic users and industry.

**EPCC**, UK

EPCC is the High-Performance and Novel Computing Centre of University of Edinburgh. EPCC provides a wide variety of services to academia and industry.

**Politecnico di Milano**, Italy

Politecnico di Milano is a State University in Italy, established in 1863, ranked as one of the most outstanding European universities.

**Inria**, France

Inria is the only public research body fully dedicated to computational sciences. Inria collaborates with the main players in public and private research in France and abroad.

**ECO2Clouds**
**Advancing ecological awareness in the Cloud**

Identifying good practices to improve energy efficiency of Cloud data centers

*

Developing techniques and mechanisms for CO2 aware application deployment on Cloud


**Project Coordinator**
**Julia Wells**
Atos Spain
Tel: +34 93 486 1818
julia.wells@atos.net

**Science and Tech Leader**
**Prof. Nikolay Mehandjiev**
The University of Manchester
Tel: +44 161 275 0579
n.mehandjiev@manchester.ac.uk

**Deputy S&T Leader**
**Usman Wajid**
usman.wajid@manchester.ac.uk





SEVENTH FRAMEWORK PROGRAMME

Grant agreement no: 318048