ECO$_2$Clouds

Experimental Awareness of CO_2 in Federated Cloud Sourcing

ECO$_2$Clouds team

Barbara Pernici, Politecnico di Milano

This project is partially funded by European Commission under the 7th Framework Programme - Grant agreement no. 318048

Overview ECO2Clouds

Partners: 6
Project type: STREP
Duration: 24 months
Start date: October 1, 2012

Work programme topic addressed:
Objective-ICT-2011.1.6 c) Fire Experimentation

Web site: http://eco2clouds.eu

ECO2Clouds rationale and motivation

Rapid proliferation of **cloud-based IT infrastructures**

Ecological implications form a gap in current state of the art in research and practice

To date little is known about how to incorporate carbon emissions and energy consumption into **application development and deployment decision models**.

Addressing this gap is vital to have an impact on future **sustainable developments**

Optimization of Energy Consumption in the Cloud Infrastructure

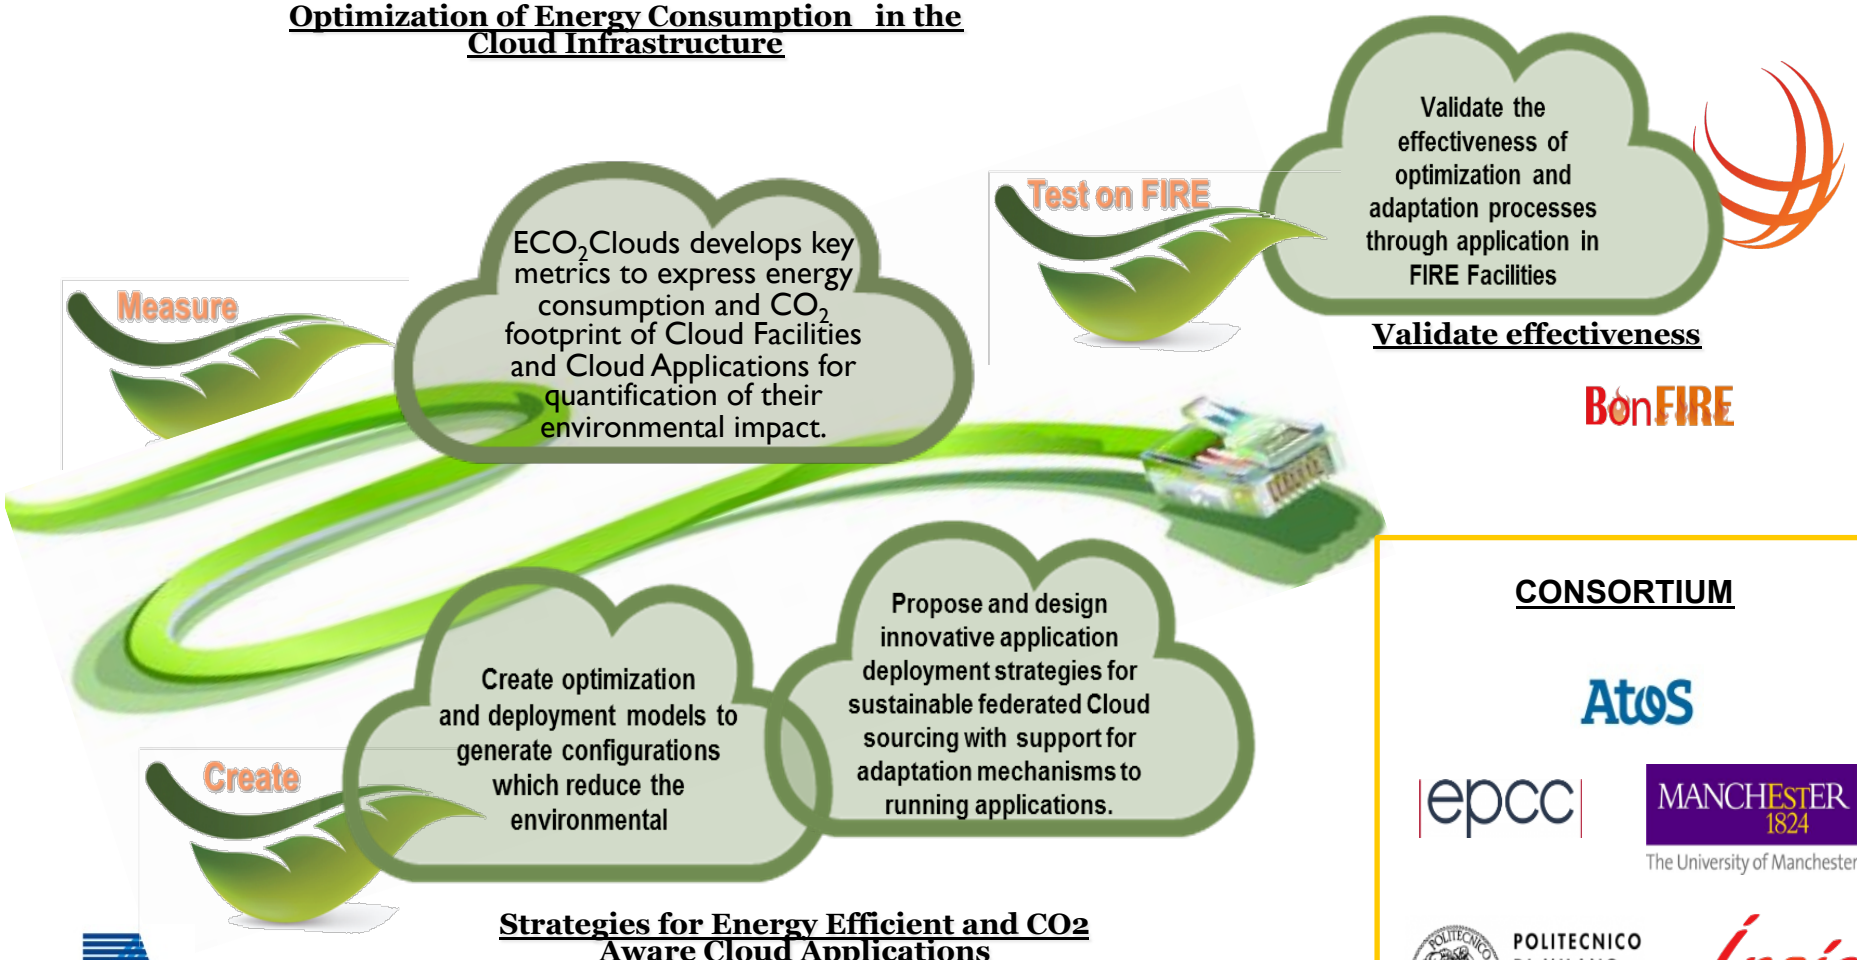

Strategies for Energy Efficient and CO2 Aware Cloud Applications

SEVENTH FRAMEWORK PROGRAMME

CONSORTIUM

Atos

epcc

MANCHESTER 1824
The University of Manchester

POLITECNICO DI MILANO

Inria
INVENTORS FOR THE DIGITAL WORLD

HLRS

Project overview: objectives

- Create **extensions** to the Cloud application programming interface and mechanisms to **expose eco-metrics** at the levels of applications, VM, and infrastructure.
- Complete implementations to **collect key eco-metrics** at VM and infrastructure level by **leveraging consumption probes** of physical nodes and assigning the measured consumption to virtual machines in a Cloud infrastructure.
- Develop software to **implement the optimization and deployment models** while ensuring infrastructure support for the deployment models and adaptation process.
- Validate the effectiveness of the proposed optimization and deployment models and adaptation process through **challenging application case studies**

ECO2Clouds use of BonFIRE

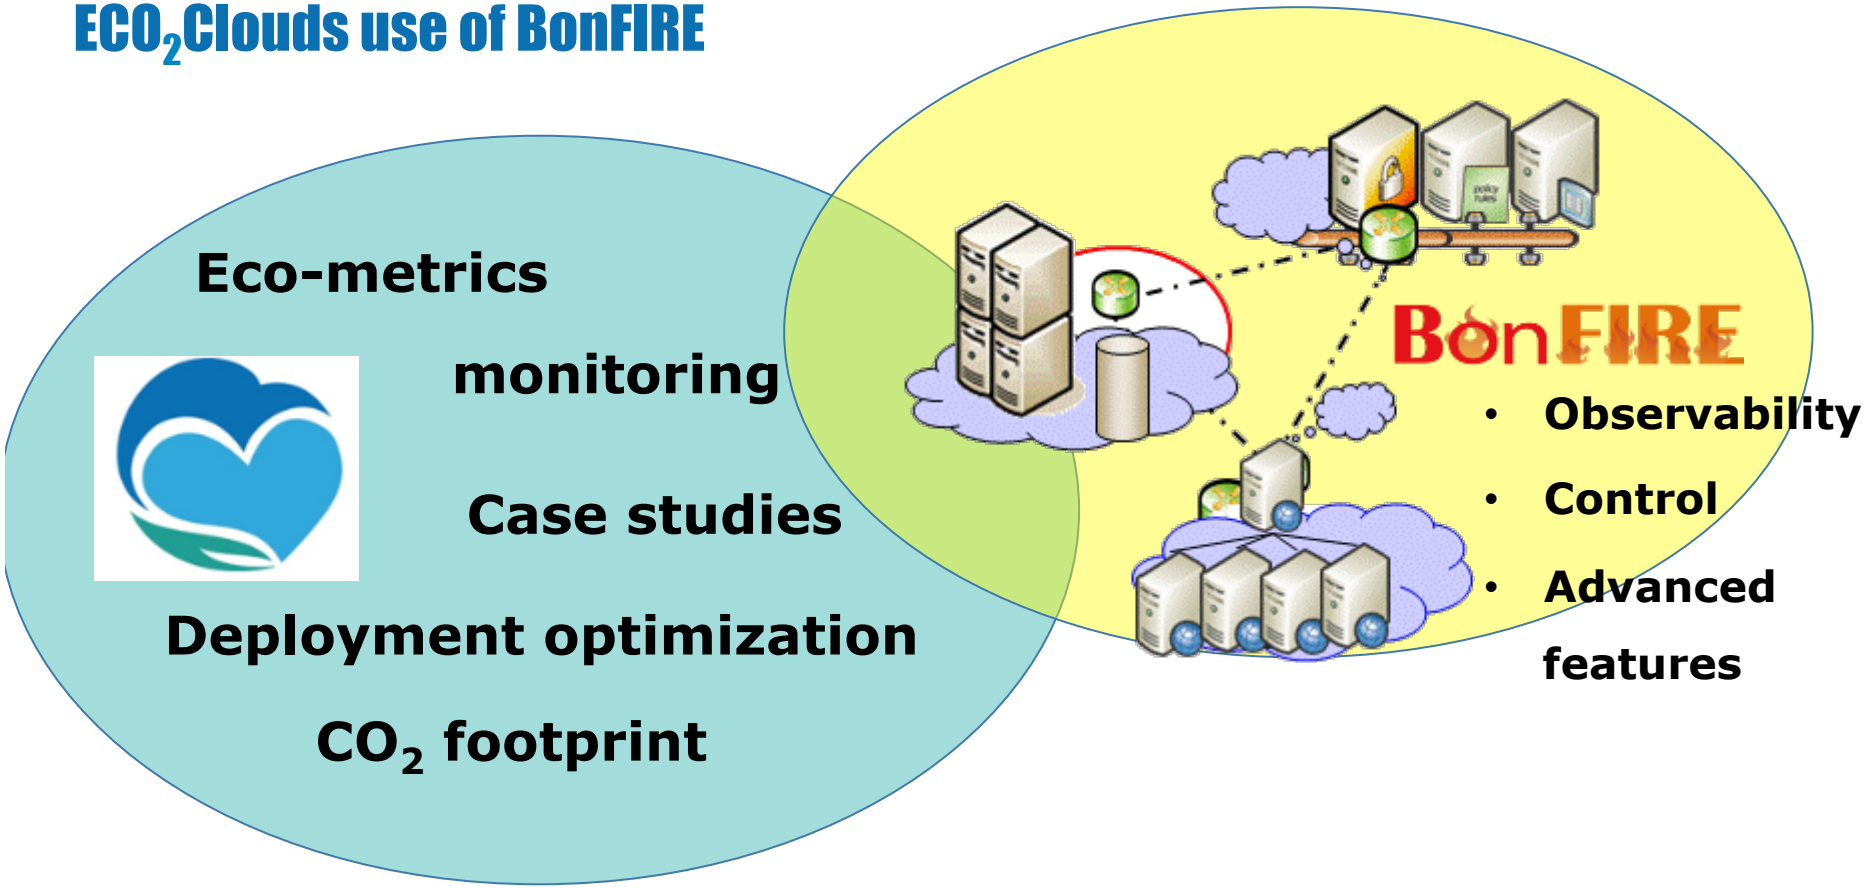

Integrate the carbon-aware mechanisms into BonFIRE so as to test, validate and optimize the eco-metrics, models and algorithms developed

… adding the ECO dimension to FIRE

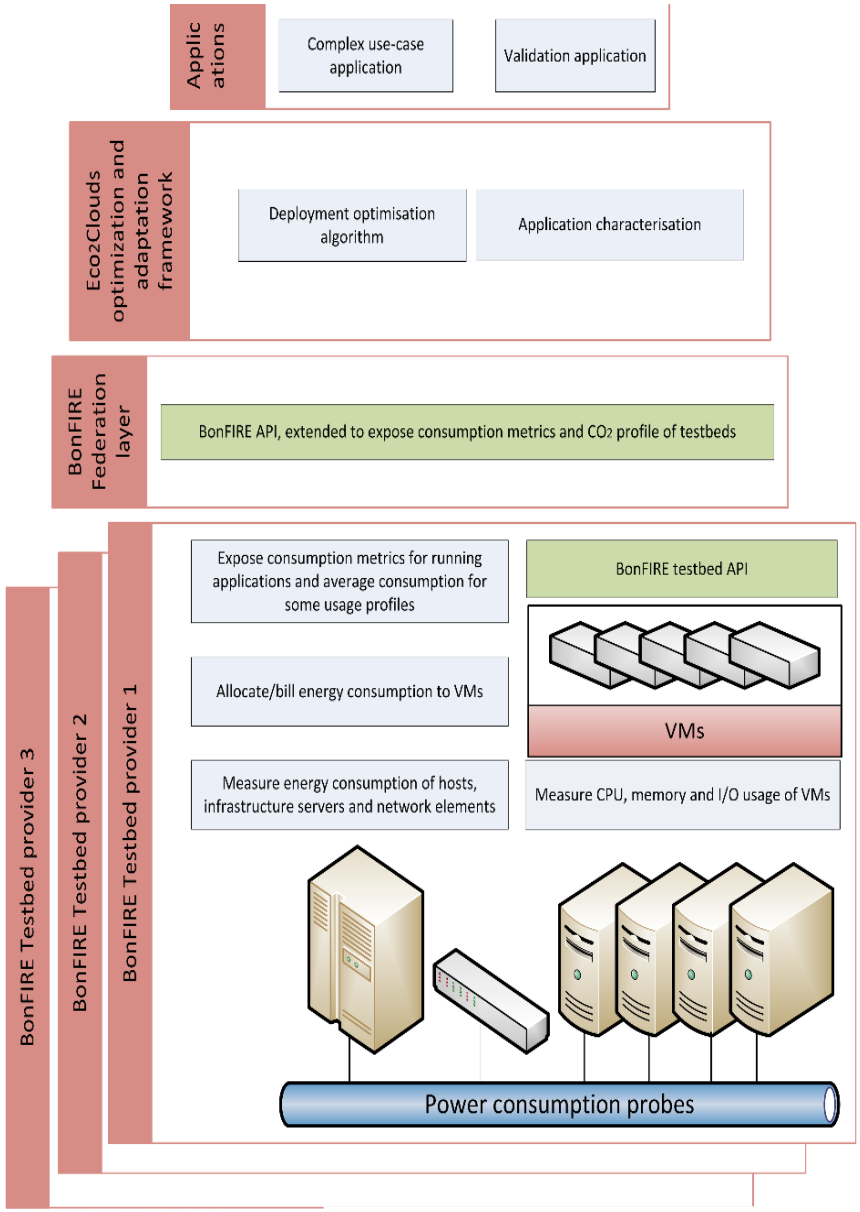

Three gardens

USTUTT-HLRS Data Center

The ECO_2Clouds site USTUTT-HLRS runs OpenNebula 3.6 in a dedicated version derived for BonFIRE.

Resources: 17 dedicated worker nodes and 36 on-request nodes

EPCC Data Center

UK-EPCC runs OpenNebula, in a version derived from OpenNebula 3.2 for BonFIRE.

Resources: EPCC provides 3 dedicated nodes as permanent resources. Two of these nodes offer four, 12-core AMD Opteron 6176 (2.3GHz)

Inria Data Center

FR-Inria runs OpenNebula, in a version derived from OpenNebula 3.6 for BonFIRE.

Resources: 4 dedicated worker nodes (DELL PowerEdge C6220 machines) and can expand over the 160 nodes of Grid'5000 located in Rennes.

Implementations

- All 4 layers are considered
  - Energy
  - Physical
  - Virtual
  - (Service)
    - User needs to implement this type due to application differences!
- Include additional templates
  - Infrastructure aggregator
  - BonFIRE aggregator

ECO2Clouds Architecture

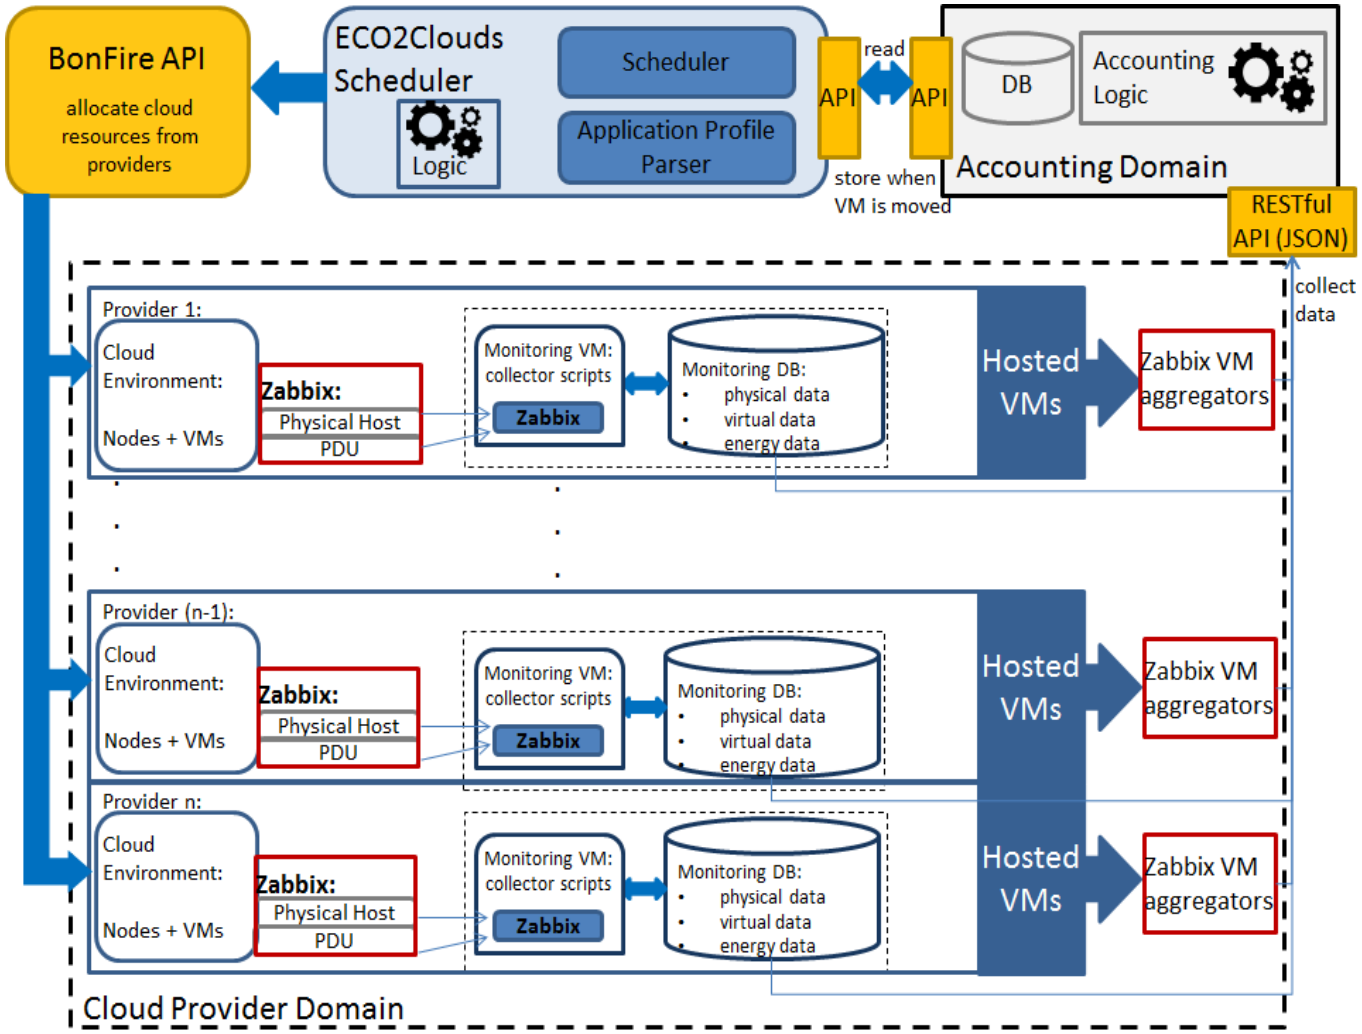

Metrics: infrastructure layer

Host

Metric	*Definition*
CPU utilization	Average utilization of the processors inside a host. For each processor, this metric indicates how much of the processor capacity is in average in use by the system
IOPS	Total number of I/O operations per second (when performing a mix of read and write tests)
Availability	The probability that a request is correctly served by a specific host within a maximum expected time frame. In order to assess this metric, it is necessary to compare the number of satisfied requests with the total amount of the requests received by the analysed host. Note that, at the host level, the request concerns the deployment of a VM on a given host
Energy Consumption	The energy consumed by the analysed host in a specific time period

Site

Metric	*Definition*
PUE	Measure that compares the power used by the entire infrastructure with the power used for computation.
Site utilization	Average utilization of the power drawn by the IT equipment with respect to the power capacity of the site
Storage utilization	Percentage of storage used with respect to the overall storage capacity within the site.
Availability	The probability that a request is correctly served by a site within a maximum expected time frame
Green Efficiency Coefficient (GEC)	Percentage of energy consumed by the site that is produced by green energy sources. This metric is calculated as the ratio between the green energy consumed by the site and the total energy consumed by the site.
CO_2 emissions	Quantity of CO_2 emitted by the site.

Metrics: virtualization layer

Metric	*Definition*
CPU Usage	Processor utilization percentage for a running application over a run time interval. It is calculated by using the ratio between the amount of used CPU and the amount of allocated CPU.
Storage Usage	Storage utilization percentage for data-read and -write operations on the corresponding storage device, computed as the ratio between the used disk space and the allocated disk space.
I/O Usage	Percentage of process execution time in which the disk is busy with read/write activity.
Memory Usage	Ratio of the average size of the portion of memory used by the process to the total amount of memory available for the application.
Energy Consumption	The energy consumed by the analysed VM in a specific time period.
VM-PUE	Measure of how efficiently a VM uses the provided power
VM-EP (VM Energy Productivity)	Ratio between the output of the VM in a certain time interval and the enrgy consumed
VM-GE (VM Green Efficiency)	Information about the portion of energy consumed by the VM that is produced by green energy sources

Metrics: application layer

Metric	*Definition*
Task Execution Time	The time taken to execute the specific task.
Application Execution Time	The time taken to execute the whole application.
Energy Consumption	The energy consumed from the analysed application in a specific time period. This metric is calculated by aggregating the energy consumed by the VMs through which the application is deployed.
Response Time	The average time taken to handle user requests. This metric is particularly relevant for interactive application. Note that for batch application the response time will coincide with the application execution time.
Throughput	Number of executions of an application within a specific time frame.
A-PUE	Measure of how efficiently an application uses the provided power
A-EP (Application Energy Productivity)	Ratio between the number of execution of an application in a certain time interval and the energy consumed
A-GE (Application Green Efficiency)	Provide information about the portion of energy consumed to execute a specific application that is produced by green energy sources.

Monitored metrics – example

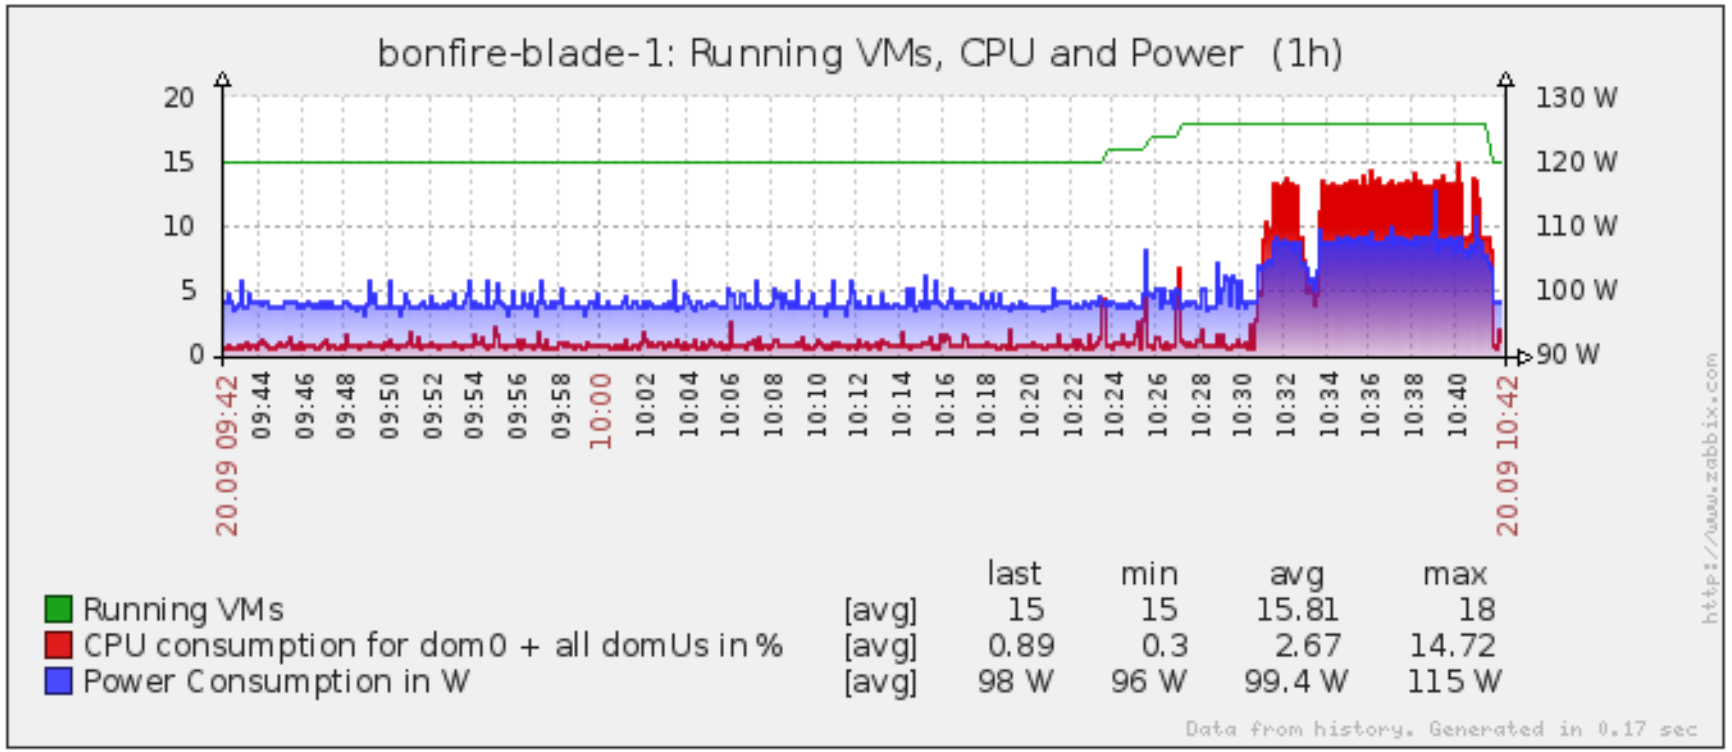

Source: Inria

Metrics: Energy Consumption

- Energy consumption is measured by
  - PDUs
  - Energy sensors (blade servers at HLRS)
- Available at INRIA, HLRS and EPCC
  - PDU scripts, usable at provider sites
- Energy Metrics:
  - calculated by use of PDU/sensor data
  - calculated by use of energy mix statistics

Greenhouse Gas metrics - Energy mix

- Live data at INRIA and EPCC
- Static values at HLRS
  - Fixed by contract

Energy mix at HLRS

fixed values at HLRS

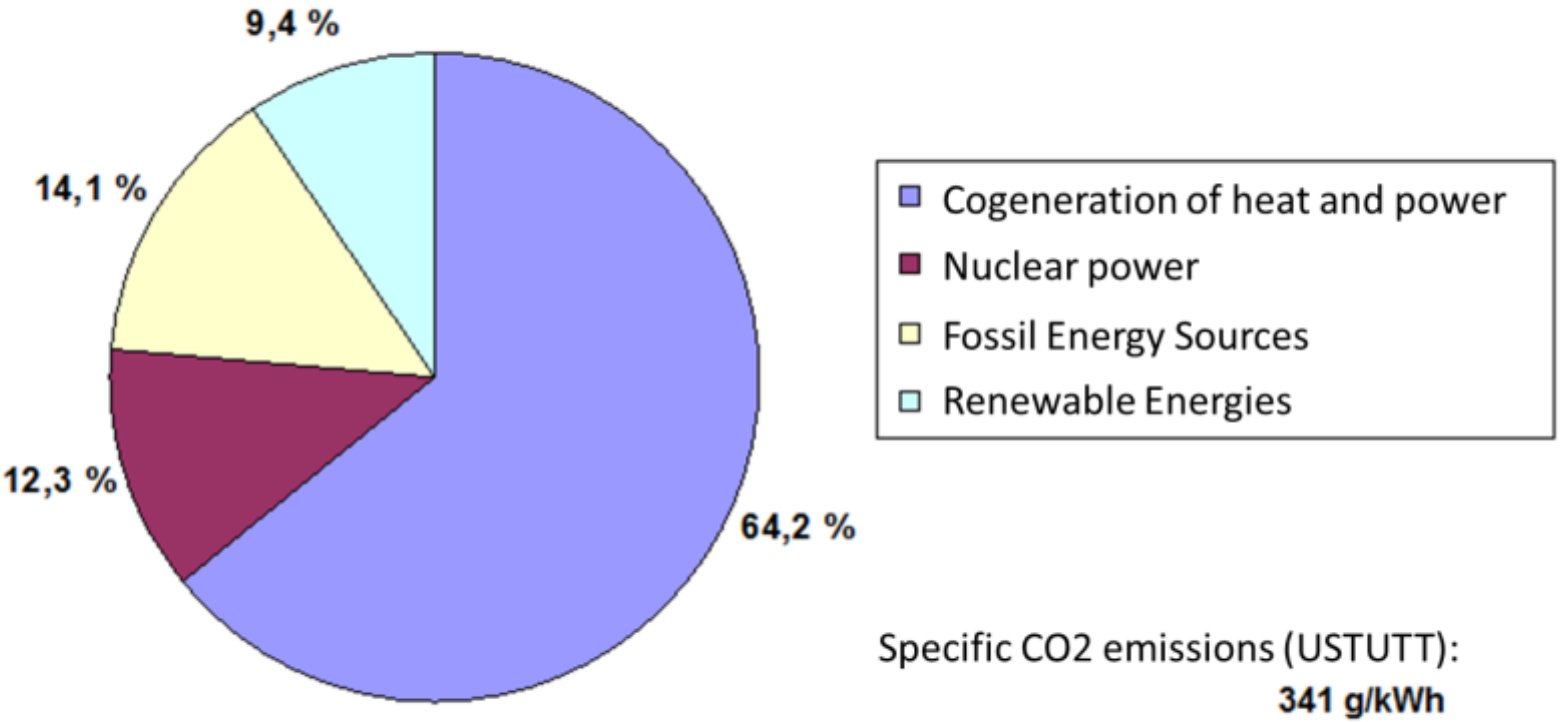

Specific CO2 emissions (USTUTT):
341 g/kWh

German CO2 emissions in general:
503 g/kWh

Energy mix at Inria – live feed from France's electricity transport company (RTE)

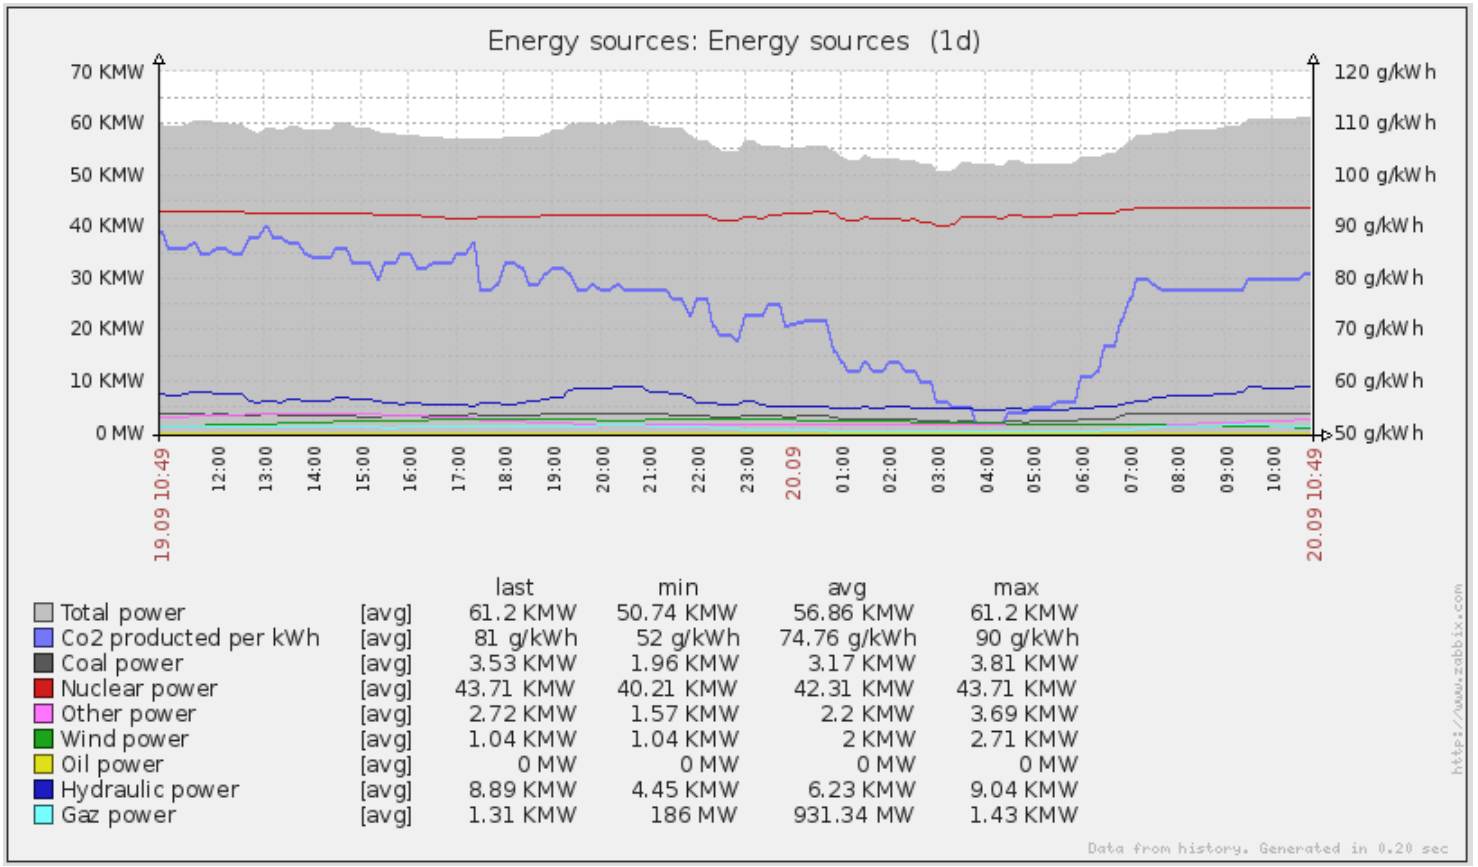

Case studies

- **Data analysis in clinical do**
  HPC
- **Eels case study**
  HPC
- **e-business with services**

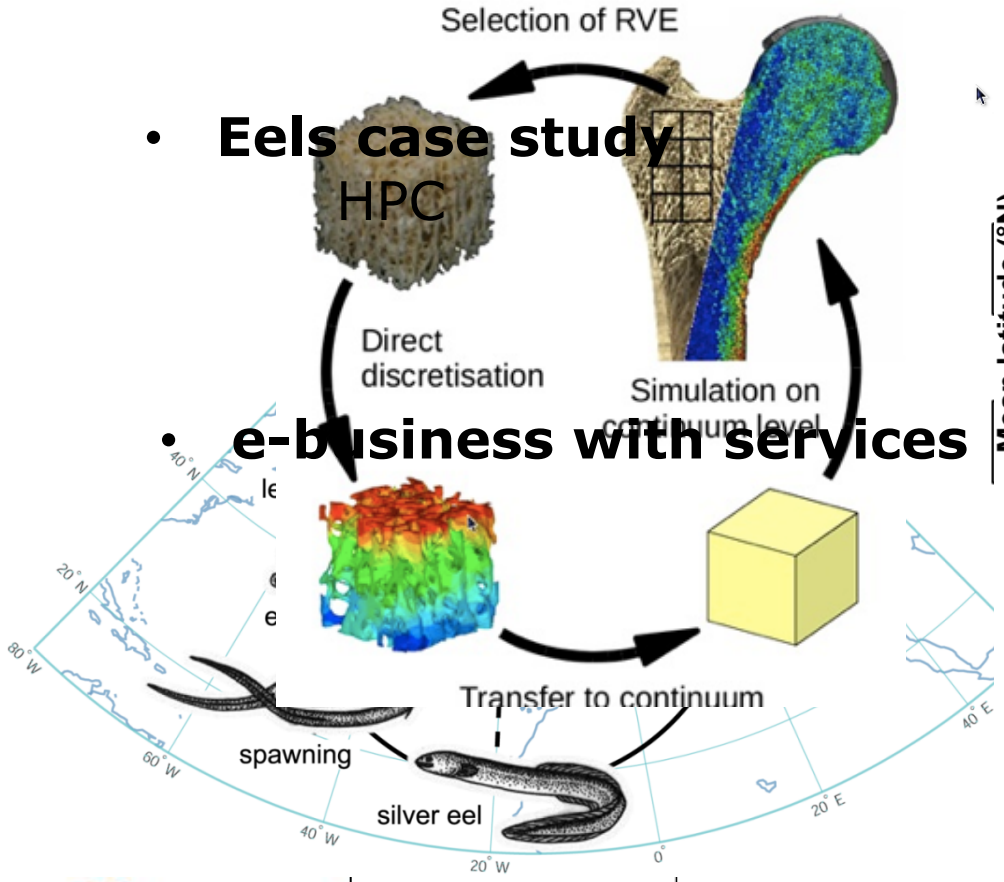

Latitudinal trend of arrivals
(40-yr simulations)

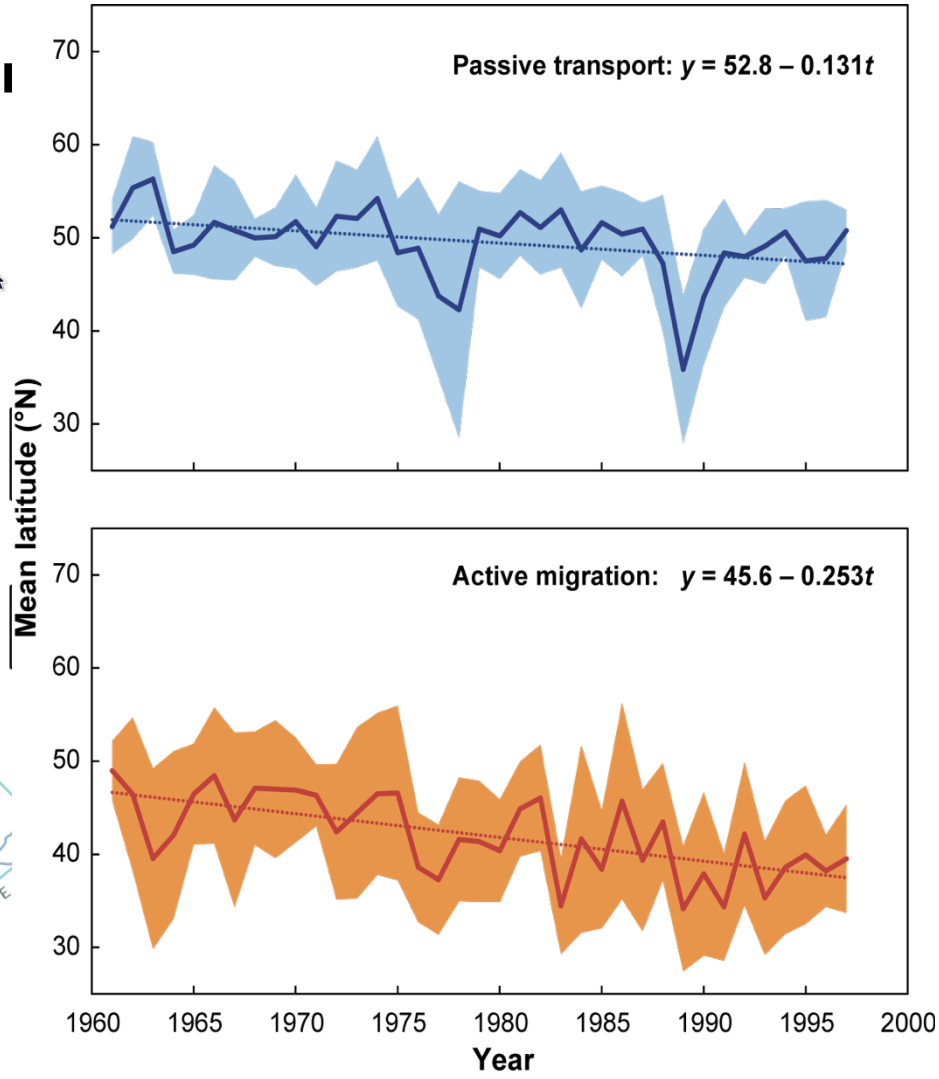

[prob = 10%

[prob = 5%

Future Work

- Derive new requirements from ongoing experimentation
- Selection of more suitable eco-metrics at different levels
- Data mining solution
- Optimization:
  - Application deployment strategies (configurations of requested resources)
  - Design-time advanced scheduling
  - Runtime adaptation

QUESTIONS?

- Further information:
  http://www.eco2clouds.eu

- Contact – project coordinator:
  Julia Wells, Atos Spain