


## Description of the Energy Consumption Data

This document describes log files that contain information about the electricity consumed by the computer servers of one of the sites of an experimental Grid platform (i.e. Grid'5000). In addition, the files describe the resource reservations made by users of the infrastructure, which servers were used, and resource failures.

### **The Grid'5000 Infrastructure**

The Grid'5000<sup>1</sup> platform depicted in Figure 1 is an experimental testbed for research in distributed computing which offers around 5000 CPU cores geographically distributed across 9 sites in France; all sites are linked by a dedicated high-speed network. The utilisation of Grid'5000 is specific: a user can reserve a number of nodes in advance to deploy a system image customised to her application, so that a node is entirely dedicated to the user during her reservation.



**Figure 1 - The Grid'5000 Infrastructure in France.**

The data described here was obtained by sensors (i.e. wattmetres) that measure the energy consumption of equipments located at one of Grid'5000 sites, in Lyon. The energy sensing infrastructure deployed on this site monitors 135 servers and the network equipments. The wattmetres have been manufactured by OMEGAWATT; each wattmetre is connected to a data collector via a serial port. A server collects and logs the data from the wattmetres and generates graphs that help understand the behaviour of equipments and user reservations in terms of energy consumption. One power measurement is carried out every second.

This site comprises two clusters namely **capricorne** and **sagittaire**, which are composed of 56 and 79 nodes respectively. Each of the capricorne machines has 2 AMD Opteron CPUs, and 2GB of RAM. The sagittaire machines are equipped with 2 AMD Opteron CPUs, and 16GB of RAM. All nodes have disks with SCSI interfaces.

---

<sup>1</sup> <http://www.grid5000.fr>

## ***Information on Reservation Requests and Resource Failures***

OAR is the Resource Management System (RMS) responsible for scheduling user reservation requests on sites of Grid'5000. The files provided here contain information on resource reservations made by users from 1 September 2009 to 27 February 2010. This interval corresponds to the same period when the energy consumption data - described later - was available. OAR provides a detailed log with information about the jobs and reservations it schedules, but we select only the fields that we considered relevant to explore correlations between resource usage and energy consumption of computer servers.

The file **reservations.txt** contains information about the resource reservation requests submitted to the Lyon site and abides by the following principles:

- Each row represents a resource reservation request;
- Lines contain a predefined number of fields separated by whitespace; and
- It is an ASCII file.

Each line of this file contains the following data fields:

1. A request ID, which is the same given by OAR scheduler.
2. The time at which the request was submitted as a UNIX timestamp.
3. When the request has started its execution as a UNIX timestamp.
4. A UNIX timestamp indicating the finish time of the request.
5. The final status of the request. The possible values are **Error** and **Terminated**. A reservation that goes until the end of the initial finish time stipulated by the user has a status Terminated. Failed requests receive the Error status. However, reservations that are cancelled by users before they approach the initial finish time receive the status Error. That is, an Error status does not necessarily mean that the request has failed.
6. An ID identifying the user who made the request.
7. The type of request. The possible values are **PASSIVE** and **INTERACTIVE**.
8. A description of the queue to which the reservations were

submitted. Here, **besteffort** means that it is a job that uses a default system image (i.e. environment) and executes when enough servers are available; **default** corresponds to resource reservation requests in which the user chooses the environment she wants to deploy; and **admin** are resource reservations made by the system administrator.

9. Indicated whether the user has made a late environment deployment (i.e. the user deployed the system disk image by herself after the start of the reservation). The possible values are **Y** for yes and **N** for no.

The file **reservations equipments.txt** details the equipments (i.e. computer servers) utilised by each resource reservation request. This file can be interpreted as a relational-database table whose first field is the key containing the reservation ID. The file contains one line for each server reserved by a given request. Hence, each line contains the following fields:

1. A reservation ID.
2. A string containing the name of the (or one of the) resource(s) utilised.

The file **equipment\_failures.txt** contains information about the failures of computer servers. This file presents time intervals during which the computer servers were not available to user requests. Each line of this file contains the following fields:

1. An incremental number.
2. The name of the server that failed.
3. The time at which the resource failed, as a UNIX timestamp.
4. A UNIX timestamp indicating when the resource resumed its operation.
5. The type of failure.

### ***The Energy Consumption Data***

The energy consumption data provided here spans six months, more precisely from 1 September 2009 to 27 February 2010. There are intervals over which the power consumption of the platform could not be determined. These intervals represent periods when the electricity supply was cut off, or the software and hardware responsible for collecting the energy consumption data failed. Although most of these issues have been identified and addressed, they have led to small gaps in the energy consumption logs during the months considered here. The file **energy\_failures.txt**

summarises the periods during which the energy consumption data is not available. Each line represents a failure in the energy measurement infrastructure and contains the following fields:

1. An incremental number.
2. The time at which the infrastructure failed, as a UNIX timestamp.
3. A UNIX timestamp indicating when the infrastructure resumed its operation.

The information about the power consumed by the computer servers is split in 135 files. A file contains information about the power consumed by an individual computer server, where each line has the following fields:

1. A UNIX timestamp indicating when the measurement was carried out.
2. The power consumption of the server in Watts.