

# The Green Computing Observatory: a data curation approach for green IT

Cécile Germain-Renaud<sup>\*</sup>, Frederic Fürst<sup>†</sup>, Michel Jouvin<sup>‡</sup>, Gilles Kassel<sup>†</sup>, Julien Nauroy<sup>§</sup>, Guillaume Philippon<sup>‡</sup>

<sup>\*</sup>Université Paris Sud, Cecile.Germain@lri.fr

<sup>†</sup>Université Picardie Jules Verne, {Frederic.Furst, Gilles.Kassel}@u-picardie.fr

<sup>‡</sup>IN2P3-CNRS, {Jouvin, Philippon}@lal.in2p3.fr

<sup>§</sup>INRIA-Saclay, Julien.Nauroy@lri.fr

## I. MOTIVATION AND GOALS

The importance of energy saving in IT systems is so widely acknowledged that there is no need to detail it here, nor the explosion of related research. [1] pointed some fundamental limitations on the path of energy-efficiency improvements.

- Energy consumption is a complex system. Manufacturers have created sophisticated HW/SW adaptive control dedicated to energy saving in processors, motherboards, and operating systems, e.g. Advanced Configuration and Power Interface, or the Intel technology for dynamically over-clocking single active cores. Administrators define management policies, such as scheduling computations and data localization with various optimization goals in mind. Finally, usage exhibits complex patterns too [2].
- The metrics remain to define. Energy efficiency should be the a ratio of energy to service delivery, but for data centers and Clouds service output is difficult to measure and varies among applications.
- Almost no public data are available. Benchmarking requires empirical data and ideally behavioral models.

The Green Computing Observatory (GCO) addresses the previous issues within the framework of a production infrastructure dedicated to e-science, providing a unique facility for the Computer Science and Engineering community.

The first barrier to improved energy efficiency is the difficulty of collecting data on the energy consumption of individual components of data centers, and the lack of overall data collection. GCO collects monitoring data on energy consumption of a large computing center, and publish them through the Grid Observatory [3]. These data include the detailed monitoring of the processors and motherboards, as well as the global site information, such as overall consumption and overall cooling, as optimizing at the global level is a promising way of research [4]. A second barrier is making the collected data usable. The difficulty is to make the data readily consistent and complete, as well as understandable for further exploitation. For this purpose, GCO opts for an ontological approach in order to rigorously define the semantics of the data (what is measured) and the context of their production (how are they acquired and/or calculated).

The overall goal is to create a full-fledged *data curation* process, with its four components: establishing and developing

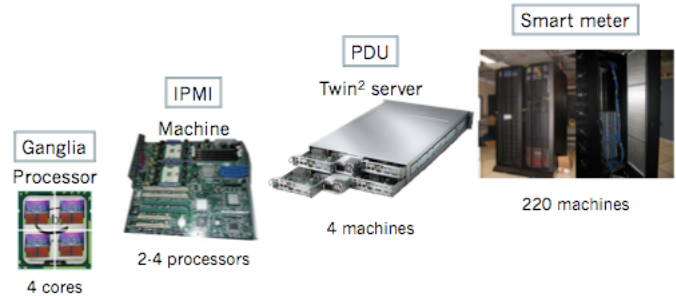


Fig. 1: Monitoring instruments and scale

a long-term repositories of digital assets for current and future references, providing digital asset search and retrieval facilities to scientific communities through a gateway, tackling the good data creation and management issues, and prominently interoperability, through formal ontology building, and finally adding value to data by generating new sources of information and knowledge through both semantic and Machine Learning based inference. This paper reports on the first achievements, specifically acquisition and ontology.

## II. THE ACQUISITION APPARATUS

We define the data curation process by developing the complete hardware and software infrastructure for monitoring the computing center of Laboratoire de l'Accélérateur Linéaire (CC-LAL), and making the resulting data publicly available. The CC-LAL is mainly a Tier 2 in the EGI grid, but also includes local services and a Cloud infrastructure, with a common characteristics of production quality and 24x7 availability. The computing site currently includes 13 racks hosting 1U systems, 4 lower-density racks (network, storage), resulting in  $\approx 240$  machines and 2200+ cores, and 500TB of storage. The following list shows why the CC-LAL is a good testbed.

- Heterogeneity, as the site features DELL and IBM systems, with classical and Twin<sup>2</sup> technology. Thus, designing a general information model is mandatory.
- Classical, cold-water based central cooling (traditional racks plus cooling) as well as the more advanced water-cooled racks are present and monitored.
- The traffic is dominated by High Energy Physics experiments, which are high throughput or loosely coupled,

data-intensive, as opposed to High Performance, compute intensive, strongly parallel workloads. We thus get some approximation of the behaviour of a data center.

- It hosts the experimental infrastructure of the StratusLab FP7 project, which builds and operates a Cloud on top of the EGI grid resources, thus creating an opportunity for Cloud-oriented monitoring.

Most work on energy consumption is based on the measurement of the inlets associated with the blades, through smart PDUs (Power Distribution Units). With the advent of Twin<sup>2</sup> servers, the granularity of energy consumption measurements becomes limited to the 8-16 processors of the server, which is clearly too coarse. We exploit the wealth of information provided by the IPMI (Intelligent Platform Management Interface) technology. IPMI is an industry agreement defining a standardized, abstracted, message-based interface to intelligent platform management hardware, and standardized records for describing platform management devices and their characteristics. Besides power consumption, extremely detailed information about the motherboard, such as the operating voltage or the fan speed is available. Nonetheless, some servers are equipped with PDUs, in order to exploit the opportunity to calibrate one instrument with another. To be useful, the energy data have to be related to computational usage. We use Ganglia to capture CPU, memory and network usage. Finally, a smart meter reports on the overall energy consumption of the site, including central cooling, and its ambient temperature. Fig. 1 shows a typical configuration of the sources. With a 5 minutes sampling period, the volume is in the order of 1GByte per day.

### III. THE INFORMATION MODEL

To rigorously define the semantics of the data while addressing the problem of their heterogeneity, an ontological approach is implemented to define concerned entities and correlate them. This approach relies on an ontology of measurement providing a general framework [5] - *magnitudes* are *assigned* to *qualities* hosted by *objects* - which is refined into a model of energy consumption in computing centers; the general framework is itself an extension to the foundational ontology DOLCE [6]. *Objects* are physical ones (e.g. computers and associated components) or temporal objects such as *processes* (e.g. rotating movement of fan) or *events* (e.g. motherboard failure). *Qualities* are description dimensions of these objects, different according to whether they are hosted by physical objects (e.g. temperature of a component), processes (e.g. speed of rotation of a fan) or events (e.g. duration of a power failure). Measurement instruments have their own qualities (e.g. resolution, calibration). *Magnitudes* can be boolean, numerical, scalars or vectors [7]. The magnitude (or measurement) assigned to a quality is obtained either by data acquisition from a sensor or by calculation from other qualities as in the case of derived qualities (e.g. power depends on voltage, intensity and power factor) or in the case of missing values (e.g. extrapolation).

This ontology translates into an XML schema. Implementation further requires supporting scalable exploration (selection,

projection, and more complex requests) of the datasets. Given the performance limitations of XML querying, we decided to offer maximal flexibility to the user. Firstly, a common schema is used for the three acquisition sources (IPMI, Ganglia and PDU), and files are structured at a fine grain (one per day, machine and source) ; flexible aggregation is made possible through the standard XInclude tool. Second, we publish both the native bulky data, and selected ones, with a typical 70% volume reduction.

The next step is to integrate this approach into a higher level view, including both acquisition from other sites, and the data dissemination process. It should be oriented towards users and usage, statistical analysis of time series. SDMX (Statistical Data and Metadata Exchange) [8] is a de-facto standard (and ISO norm for SDMX1.0) within the sphere of economic data and the extension of the ontology will be done in line with the SDMX information model.

Finally, the issue of Linked Data [9] should be considered: multiple repositories of IT energy data may appear. Examples such as [10] indicate that building on the SDMX experience makes the transition to Linked Data access manageable.

### IV. CONCLUSION

The final motivation of monitoring and semantics is optimization. Due to the lack of previous data at this scale, and the complexity of the datasets, a necessary intermediate step is to build parsimonious representations and descriptive models from the large dimension space available from the detailed monitoring. Our ontology-oriented approach offers the perspective of combining semantic and logical inference for achieving this goal.

#### Acknowledgments

This work has been partially supported by the EU FP7 project EGI-InsPIRE INFISO-RI-261323, the France Grilles initiative, and the CNRS program PEPS 2011-2012.

### REFERENCES

- [1] U.S. Environmental Protection Agency, "Report to Congress on Server and Data Center Energy Efficiency," Tech. Rep., 2007.
- [2] I. Rodero and al., "Energy-efficient application-aware online provisioning for virtualized clouds and data centers," in *Int. Conf. on Green Computing*, 2010, pp. 31–45.
- [3] C. Germain-Renaud and al., "The Grid Observatory," in *11th IEEE/ACM Int. Symp. on Cluster Computing and the Grid*, 2011, pp. 114–123.
- [4] R. Das, J. O. Kephart, J. Lenchner, and H. Hamann, "Utility-function-driven energy-efficient cooling in data centers," in *7th Int. Conf. on Autonomic computing*, 2010, pp. 61–70.
- [5] W. Kuhn, "A functional ontology of observation and measurement," in *3rd Int. Conf. on GeoSpatial Semantics*, 2009, pp. 26–43.
- [6] S. Borgo and C. Masolo, "Foundational choices in dolce," in *Handbook on Ontologies*, 2nd ed. Springer, 2009.
- [7] T. Gruber and G. Olsen, "An ontology for engineering mathematics," in *4th Int. Conf. on Principles of Knowledge Representation and Reasoning*, 1994, pp. 166–181.
- [8] *SDMX2.1 Technical Specification - Section 2: The Information model*, The SDMX initiative Std., 2011.
- [9] C. Bizer, T. Heath, and T. Berners-Lee, "Linked Data - The Story So Far," *Int. J. on Semantic Web and Information Systems*, no. 3, pp. 1–22, 2009.
- [10] R. Cyganiak and al., "Semantic Statistics: Bringing Together SDMX and SCOVO," in *Linked Data on the Web Workshop at the 19th Int. World Wide Web Conference*, 2010.