



Contribution ID: 10

Type: **not specified**

The Green Computing Observatory

Friday, 30 March 2012 11:00 (30 minutes)

Description of the Work

The LAL site includes 13 racks hosting 1U systems, 4 lower-density racks (network, storage), resulting in 240 machines, 2200+ cores, and 500TB of storage.

- The site features DELL and IBM systems, with classical and Twin2 technology; thus, designing an interoperable information model is mandatory.
- Classical, cold-water based central cooling as well as advanced water-cooled racks are present.
- LAL hosts the experimental infrastructure of the StratusLab FP7 project, creating an opportunity for Cloud-oriented monitoring.

The de facto standard IPMI (Intelligent Platform Management Interface) technology provides detailed motherboard information: power consumption, operating voltage or the fans speed, and events reporting, are available. Some servers are equipped with network-enabled PDUs (Power Distribution Units) too, 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, smart meters report on the overall energy consumption of the site, and its ambient temperature. With a 5 minutes sampling period, the volume is in the order of 1GByte per day.

To rigorously define the semantics of the heterogeneous data, an ontological approach defines concerned entities and correlates them. It refines the general ontology of measurement of [Kuh09] - magnitudes are assigned to qualities hosted by objects, itself an extension to the foundational ontology DOLCE [Bor09]. This ontology translates into an XML schema 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. 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.

Conclusions

The overall goal of GCO is to create a full-fledged data curation process, with its four components: establishing 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.

[Kuh09] W. Kuhn, "A functional ontology of observation and measurement," in 3rd Int. Conf. on GeoSpatial Semantics, 2009

[Bor09] S. Borgo and C. Masolo, "Foundational choices in dolce," in Handbook on Ontologies, Springer, 2009.

[EPA] U.S. Environmental Protection Agency, "Report to Congress on Server and Data Center Energy Efficiency", 2007.

Impact

The 2007 U.S. Environmental Protection Agency report [EPA] 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
- The metrics remain to define. “Energy efficiency” should be the 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.

Its combination of large-scale acquisition and fine-grain monitoring technology have no equivalent. Most related work on energy consumption is based on the measurement of the inlets associated with the blades, through smart PDUs. With the advent of Twin2 servers, the granularity of measurements becomes limited to the 8-16 processors of the server, which is clearly too coarse.

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) is a de-facto standard (and ISO norm for SDMX1.0) within the sphere of economic data; the extension of the ontology will be done in line with the SDMX model.

Overview (For the conference guide)

The first barrier to improved energy efficiency of IT systems is the lack of large-scale collections of experimental data. The Green Computing Observatory (GCO) monitors a large computing center (Laboratoire de l'Accélérateur Linéaire - LAL) within the EGI grid, and publishes the data through the Grid Observatory. 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. 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).

Primary authors: GERMAIN-RENAUD, Cecile (CNRS); FURST, Frederic (Universite Picardie Jules Verne); KASSEL, Gilles (Universite Picardie Jules Verne); PHILIPPON, Guillaume (IN2P3-CNRS); NAUROY, Julien (INRIA-Saclay); JOUVIN, Michel (IN2P3-CNRS)

Presenter: NAUROY, Julien (INRIA-Saclay)

Session Classification: Perspectives

Track Classification: Operational services and infrastructure