



Contribution ID: 69

Type: Poster

Gustav : a lightweight approach to CPU accounting

Monday 12 April 2010 17:48 (3 minutes)

Gustav is a CPU usage accounting tool developed by INFN, COMETA and the KISTI institute. Gustav collects accounting records from resources and publishes them to a centralised relational database, that can be queried through a web interface. Gustav is characterised by a lightweight architecture, that makes its usage ideal for small sized infrastructures. However, interoperability with more powerful tools, like DGAS or APEL, is accomplished by the adoption of standard format for accounting records.

Detailed analysis

Gustav publishers, implemented in Python, run on resources and analyse periodically batch systems logs to produce statistics about CPU usage by grid users. Records are published to a central collector, implemented through a MySQL database; the central collector runs also a web interface that make possible to perform queries about collected data. The query mechanism is very flexible, allowing to focus from single user/VO/resource/day to any combination of these parameters. Supported batch systems are currently Torque and LSF. The publisher implementation, which relies more on the batch system rather than the middleware, makes Gustav usable for CPU accounting of any grid middleware, like gLite, UNICORE or ARC, as long as it supports one the above mentioned batch systems. Beside support to most widely used batch systems, it allows also interoperability with more widely used tools, making Gustav records processable by them through the adoption of standard formats for records.

Conclusions and Future Work

In a distributed computing infrastructure, based on shared resources, an accounting system is a fundamental requirement. Gustav aims to reach this goal in a simple and reliable way, with a simple installation and an easy to use web interface. It's independent of the middleware used to run the Grid infrastructure, and supports two of the most common LRMS. Efforts will be made in the near future to enhance its capability to show collected data in a graphic way, and plugins for other LRMS such as Condor and SGE are being considered. The deployment of Gustav on Asian sites has been also planned.

Impact

Gustav is currently deployed on the regional Grid of the COMETA Consortium in Italy, where it proved to work well on a gLite-based infrastructure made up of a dozen of sites, tens of VOs and some hundreds of users. The installation performed by local system administrators ran smoothly, and the whole system has been deployed very quickly. Several tools were already available for accounting purposes, like DGAS/HLRmon or APEL. They have a clean and robust design, suitable for larger scale infrastructures with some thousands of sites and users. These tools provide also exhaustive aggregated usage information. However their deployment is not trivial, and can require a large effort for regional Grids or single sites requiring a much simpler accounting system. Furthermore, given the small sizes of infrastructures whom Gustav has been designed for, it is not required to aggregate records, which are stored with their full details, allowing in this way a finer granularity on queries.

Keywords

CPU usage accounting, gLite, UNICORE, ARC, grid, PBS, Torque, LSF

URL for further information

<https://gustav.consortio-cometa.it>

Authors: Dr GIORGIO, Emidio (Unknown); Mr RICCIARDI, Gianni Mario (KISTI institute)

Co-authors: Prof. BARBERA, Roberto (University of Catania & INFN); Dr HWANG, Soonwook (KISTI institute)

Presenter: Dr GIORGIO, Emidio (Unknown)

Session Classification: Poster session

Track Classification: Software services exploiting and/or extending grid middleware (gLite, ARC, UNICORE etc)