

Grid Monitoring: From Small-Scale to Large-Scale Grids

Serafeim Zaniolas, Rizos Sakellariou
University of Manchester

Monitoring is the process of collecting data about entities of interest. In the context of grids, such data are typically about hosts, networks, services, datasets and running applications. Due to the dynamic nature of grids, monitoring services are needed to query the state of grid resources (resource enquiry) and to match user requirements with available grid resources (resource discovery). As such, monitoring is an important part of any grid information service.

This talk will introduce the concepts and requirements for grid monitoring systems, and outline the Global Grid Forum's recommendation on a Grid Monitoring Architecture (GMA). A taxonomy that is based on the GMA, will be presented to classify grid monitoring and information systems, with respect to a variety of properties; most notably a system's ability to scale, taking into account any provision for distribution and replication of components. Using the taxonomy, several grid monitoring and information systems will be classified, including Globus Monitoring and Discovery Service (MDS), Network Weather Service, Gridlab Mercury, and R-GMA.

The survey and classification of existing systems suggests that they all strive to deliver reliable results by means of online query resolution (i.e., resolving user queries against recently acquired data). The partitioning of services to accommodate administrative autonomy, combined with online query resolution, leads to low performance in the resolution of cross-administrative or large-scope resource discovery queries. For instance, the execution of the query “search throughout the Grid for the cheapest resource with given characteristics” would require a significant amount of time, even using the most scalable of the existing systems.

To deal with this lack of scalability, we envisage large-scale information services that will collect and maintain data about resources throughout the grid. An example of such a service is grid search engines that, similarly to web search engines, will deliver instant resolution of grid-wide resource discovery queries, at the cost of some results being stale. Another example is directories (as in Yahoo!) of grid resources, where end-users can lookup for resources of interest in predefined categories, such as application-domain, QoS, physical location, etc.

The realisation of such large-scale information services requires a monitoring infrastructure that will provide the mechanisms to maintain an up-to-date collection of data about resources throughout the grid, and at the same time hide the existing systems' heterogeneity in data models, representation, etc. The talk will propose an architecture for such a worldwide monitoring infrastructure, outline the challenges involved and discuss the relevant design space.

This work builds on the following publications:

- Serafeim Zaniolas and Rizos Sakellariou. "A Taxonomy of Grid Monitoring Systems" *Journal of Future Generation Computer Systems*, 21(1), January 2005, pp. 163--188.
- Serafeim Zaniolas and Rizos Sakellariou. "Towards a Monitoring Framework for Worldwide Grid Information Services" 10th International Euro-Par Conference, Vol. 3149 of *Lecture Notes in Computer Science*, Springer-Verlag, Pisa, Italy, August 31--September 3, 2004, pp. 417--422.