

Monitoring the Grid at local, national, and global levels

Tuesday, 6 September 2011 14:00 (25 minutes)

Monitoring the Grid at local, national, and global levels
The GridPP Collaboration

The World-wide LHC Computing Grid is the computing infrastructure setup to process the experimental data coming from the experiments at the Large Hadron Collider located at CERN.

GridPP is the project that provides the UK part of this infrastructure across 19 sites in the UK. To ensure that these large computational resources are available and reliable requires many different monitoring systems. These range from local site monitoring of, for example, the hardware and of batch system utilization, to UK-wide monitoring of Grid functionality and ultimately the worldwide monitoring of resource provision and usage. In this paper we describe the monitoring systems used for the many different aspects of the system, and how some of them are being integrated together.

Local site monitoring covers, cluster load, batch system utilization, network bandwidth monitoring and fault condition monitoring. The most common software used to monitor a cluster is Ganglia, this system can be easily installed on all clients allowing data to be collected on a master node and displayed via a web server. Monitoring specific to the batch system used at a site is also typically used. Many GridPP sites use the torque batch system (developed from PBS). This can be monitored with pbswebmon, which provides a graphical way to monitor the occupancy of the cluster, and the different user's job shares and efficiencies. Another tool is Nagios, which provides a very powerful framework that can be used to monitor the status of systems. The Nagios system can be configured to run tests at intervals and carry out actions dependant on the results. This can be emailing a warning message or running an event handler that takes remedial action to solve a problem. One of the advantages of Nagios is that if all is well it does not bother you and there is no need to actually look at a status Web page. It can let you know (via email, web or SMS) when there is a problem. Network health, usage and bandwidth is monitored at many sites with cacti and/ or Network Weathermap. Available bandwidth between sites in the UK is monitored by each site having a dedicated 'Gridmon' test box that performs a matrix of iperf and other tests between the UK sites. The results are stored on a central database with a web frontend.

Other UK wide testing includes a GridPP developed summation of relevant WLCG tests coupled with dedicated UK tests developed by Prof. S. Lloyd at QMUL and the UK regional Nagios based Service Availability Monitoring (SAM). This service queries a central database (GOCDB) and Grid information services to create a list of sites and systems to be tested. The services offered are tested and the results of the tests are sent via an active MQ message bus to the EGI Central Operations Dashboard. Each region has an operator on duty that can raise alarm tickets against sites that have failed critical tests.

Systems Administrators are often overwhelmed by the number of different web sites and monitoring systems they should track. Attempts to integrate output from several systems into a site dashboard have been made at the Tier 1 and some of the larger sites. These systems will be described.

Primary author: Mr GRONBECH, Peter (Particle Physics-University of Oxford)

Presenter: Mr GRONBECH, Peter (Particle Physics-University of Oxford)

Session Classification: Tuesday 06th - Computing Technology for Physics Research

Track Classification: Track 1: Computing Technology for Physics Research