A job's life

Managing a job's passage through the Large Hadron Collider Computing Grid

The next generation of Particle Physics experiments will run on the Large Hadron Collider (LHC) currently being built at CERN near Geneva. These experiments will start in 2007 and will produce vast quantities of data. Similarly vast quantities of simulated data will be required to understand this real data.

The LHC Computing Grid (LCG) is the largest functioning Grid in the world consisting of more than 5,000 CPUs with almost 4,000TB of storage distributed at 73 sites around the world. GridPP provides the UK sites for the LCG. The LCG is designed to satisfy the needs of the Particle Physics community as they head towards the LHC generation.

In this presentation we examine the progress of a job as it travels around the LCG. Firstly the user describes the needs of their job using a Job Description Language and submits it from a User Interface (UI) node to a Workload Management (WM) node. The WM then matches the job's needs to the sites available and submits the job to the most suitable site. This matching process can take into account the location of the data that the job needs to analyse as well as the characteristics of the site (number of free CPUs, free storage space etc). The job then runs on the selected Compute Element (CE).

After the job has completed, the small output files from the job are transferred back to WM ready to be collected by the user, while larger files may stored on a Storage Element (SE) and their location registered in the Replica Location Service (RLS). All changes to the state of a job are recorded in a database at WM node. The user can interrogate this database in order to ascertain the current status of their jobs.

The LCG is built on many of the standard grid tools including Globus Toolkit 2 and CondorG.

Characterising the performance of a distributed Grid is notoriously difficult. However, we have made some crude attempts to quantify the performance of the LCG and its predecessor the European DataGrid (EDG). These are focussed around the efficiency of the Grid. Many thousands of jobs with different characteristics have been run over the EDG and LCG, in this presentation we show some of the results of these studies of their efficiencies.

Finally, we will describe the demonstration of the LCG that can be seen at the PPARC booth.

PPARC GridPP
UK Computing for Particle Physics

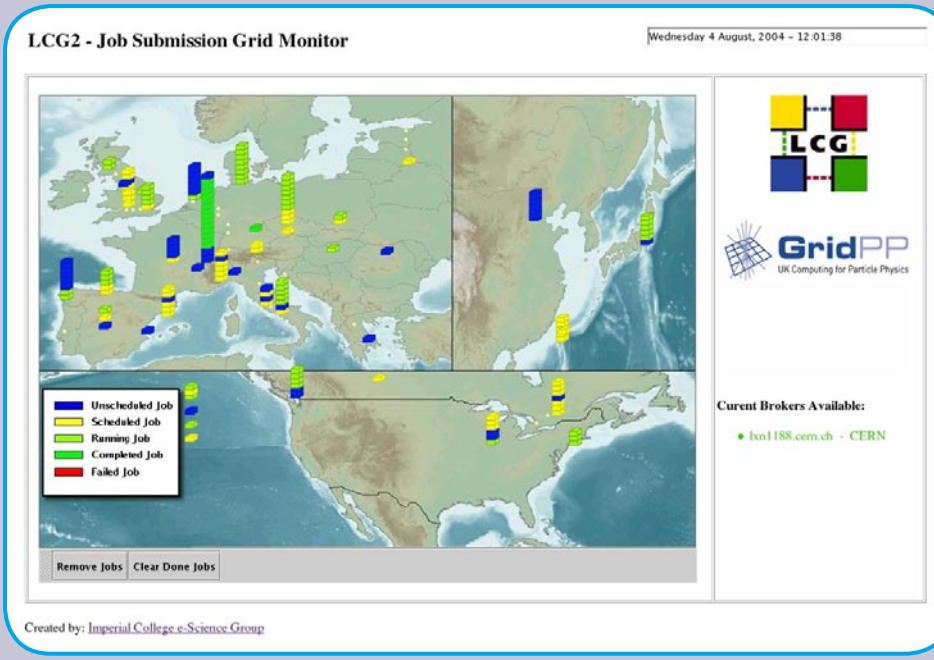

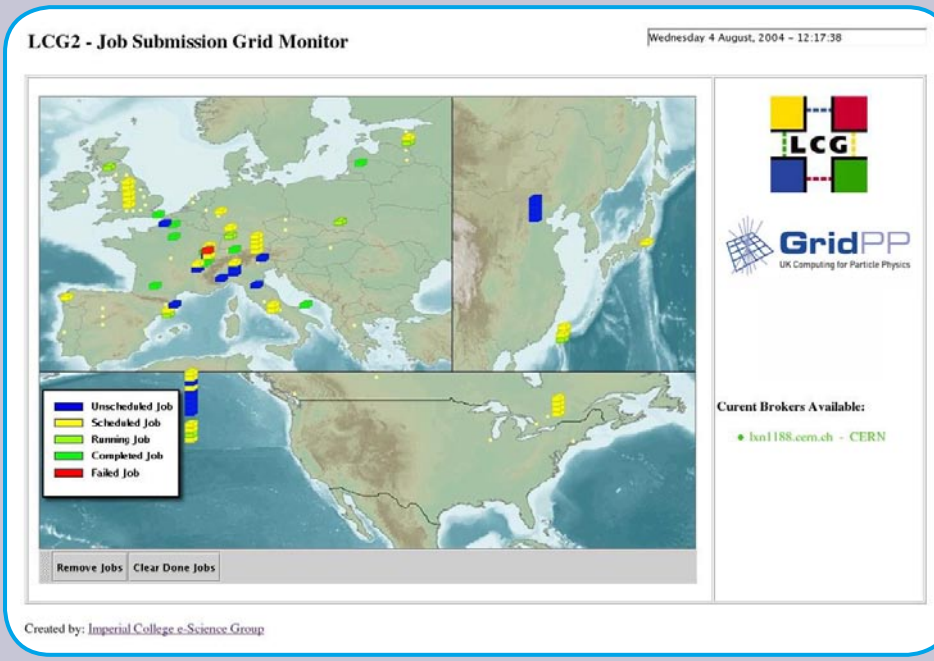

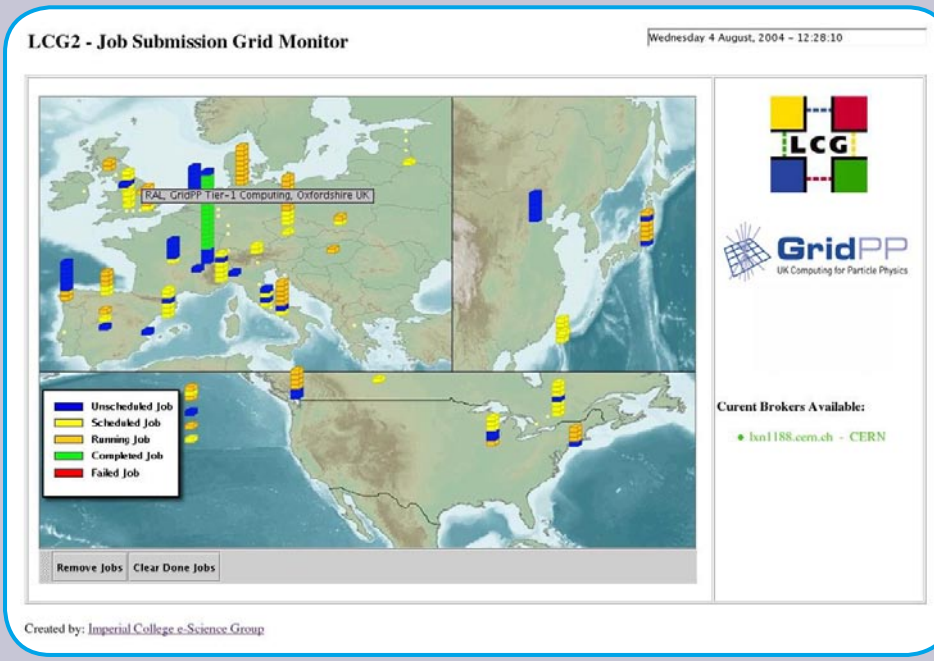

Monitoring the LCG. Jobs are submitted to the Resource Broker at CERN (blue), are distributed to sites on LCG and then scheduled (yellow). The jobs then run (flashing) and are returned to CERN either successfully completed (green) or failed (red).