

Job Efficiencies on the RAL Tier-1 Batch Farm

Wednesday, 15 February 2006 09:00 (20 minutes)

In preparation of the Grid for LHC start-up, and as part of the early production service (under the UK GridPP project), we calculate efficiencies for jobs submitted to the RAL Tier-1 Batch Farm. Early usage of the Farm was characterised by high occupancy, but low efficiency of Grid jobs, but improvement has been observed over the last six months. This behaviour has been examined by calculating overall efficiencies, defined as ratios of the total CPU time and the total elapsed wall time. This is done on a monthly basis for each virtual organisation (VO) and for the Farm as a whole. The generation of the statistics is fully automatic and is based on querying job parameters stored in a MySQL database. The data give an overview of how efficiently the Farm is being used, and identify VOs whose efficiency is low. Further information is gained from per-VO scatter plots of CPU time against efficiency for each job. In particular, these plots can identify classes of jobs that terminate due to CPU time and elapsed wall time limits being hit in the batch system. Many factors can lead to low job efficiencies, including local execution problems (e.g., high rates of disk I/O), and Grid-related problems (e.g., transferring remote data). As the efficiency data provide information about job execution on the Farm, they are therefore of use to both site administrators and end users.

Primary author: Dr HODGES, Matthew (RAL - CCLRC)

Co-authors: Mr ROSS, Derek (RAL - CCLRC); Mr TRAYLEN, Steve (RAL - CCLRC)

Presenter: Dr HODGES, Matthew (RAL - CCLRC)

Session Classification: Poster

Track Classification: Grid middleware and e-Infrastructure operation