

Multicore job management in the Worldwide LHC Computing Grid

Tuesday, 20 May 2014 15:10 (20 minutes)

After two years since the very successful first run of the Large Hadron Collider finished, data taking is scheduled to be restarted in early 2015. The experimental conditions for this second run include higher collision energy and beam luminosities, both leading to increased data volumes and event complexity. In order to process the data generated in such scenario, and also best exploit the multicore architectures of current CPUs, the LHC experiments have been developing parallelized data analysis and simulation software. However, workload scheduling in these conditions becomes a complex problem in itself, as computing jobs with a broad range of resources requirements have to be efficiently distributed across the multiple sites which make up the Worldwide LHC Computing Grid. A WLCG Task Force has been created with the purpose of coordinating the joint effort from experiments and WLCG sites. This contribution will present the activities of the Task Force, including the experiences from sites on how to best use the different batch system technologies, the development of advanced workload submission tools by the experiments and the real-size scale tests of the different proposed strategies.

Description of work

Job scheduling in a distributed resources environment such as the WLCG involves the grid-wide workload submission tools used by the LHC experimental collaborations, known as Virtual Organizations (VO) in this context, and the batch system technologies in charge of the allocation of the local resources, which are deployed at every WLCG site. The objective of the WLCG Multicore Deployment Task Force is to explore, develop and propose ways to connect both elements in order to fulfill the computing needs of the different VOs, which now require sites to be able to run their newly developed multicore applications in addition to the more usual single-core software. Furthermore, the best use of the resources must be ensured, avoiding CPUs being idle when there is work to be done and minimizing CPU inefficiencies which may be originated by the scheduling mechanisms. All this should be achieved without imposing on the sites unnecessary complexities in the way they manage their resources and maintaining a high rotation of jobs from the different users in multi-VO sites.

Wider impact and conclusions

Apart from the main objective of satisfying the new computing needs of the LHC VOs, this task force has the mandate of providing the necessary coordination in order to avoid duplicated efforts in the development of new grid-wide submission tools, as well as ensuring the convergence of approaches from different VOs to best use shared resources. Additionally, a better understanding of the technical capabilities of existing batch systems and schedulers is expected, as the participants develop and present the best system configurations, which may be shared between sites operating the same technologies.

URL(s) for further info

<https://twiki.cern.ch/twiki/bin/view/LCG/DeployMultiCore>

Primary authors: FORTI, Alessandra (MANCHESTER); Dr PEREZ-CALERO YZQUIERDO, Antonio (CIEMAT)

Presenter: Dr PEREZ-CALERO YZQUIERDO, Antonio (CIEMAT)

Session Classification: Porting new applications to EGI

Track Classification: Porting applications to the grid and cloud platform (Track Leaders: G. Sipos, D. Wallom)