



Contribution ID: 333

Type: oral presentation

Multicore job scheduling in the Worldwide LHC Computing Grid

Tuesday, April 14, 2015 3:30 PM (15 minutes)

After the successful first run of the LHC, data taking will restart in early 2015 with unprecedented experimental conditions leading to increased data volumes and event complexity. In order to process the data generated in such scenario and exploit the multicore architectures of current CPUs, the LHC experiments have developed parallelized software for data reconstruction and simulation. A good fraction of their computing effort is still expected to be executed as single-core tasks. Therefore, jobs with diverse resources requirements will be distributed across the Worldwide LHC Computing Grid (WLCG), making workload scheduling a complex problem in itself.

In response to this challenge, the WLCG Multicore Deployment Task Force has been created with the purpose of coordinating the joint effort from experiments and WLCG sites. The main objective is to ensure the convergence of approaches from the different LHC Virtual Organizations (VOs) to make the best use of the shared resources in order to satisfy their new computing needs and minimize any inefficiency deriving from the scheduling mechanisms. This should also be achieved without imposing unnecessary complexities in the way sites manage their resources.

Job scheduling in the WLCG involves the use of grid-wide workload submission tools by the VOs linked via Computing Element (CE) middleware to the batch system technologies managing local resources at every WLCG site. Each of these elements and their interaction has been analyzed by the Task Force. The various job submission strategies proposed by the LHC VOs have been evaluated, providing feedback for the evolution of their grid-wide submission models and tools. The diverse capabilities of different CE technologies in passing the resource request from the VOs to the sites have been examined. The technical features of the most common batch systems in WLCG sites have been discussed for a better understanding of their multicore job handling capabilities. Participants in the Task Force have also been encouraged to share their system configurations with the purpose of avoiding duplicated efforts among sites operating the same technologies.

This contribution will present the activities and progress of the Task Force related to the aforementioned topics, including experiences from key sites on how to best use different batch system technologies, the evolution of workload submission tools by the experiments and the knowledge gained from scale tests of the different proposed job submission strategies.

Authors: FORTI, Alessandra (University of Manchester (GB)); FILIPCIC, Andrej (Jozef Stefan Institute (SI)); LAHIFF, Andrew David (STFC - Rutherford Appleton Lab. (GB)); Dr PEREZ-CALERO YZQUIERDO, Antonio (Centro de Investigaciones Energ. Medioambientales y Tecn. - (ES)); ACOSTA SILVA, Carlos (Universitat Autònoma de Barcelona (ES)); WALKER, Christopher John (University of London (GB)); TRAYNOR, Daniel Peter (University of London (GB)); TEMPLON, Jeff (NIKHEF (NL)); ALEF, Manfred (Karlsruhe Institute of Technology (KIT)); GILA, Miguel (ETH Zurich); Dr WALKER, Rodney (Ludwig-Maximilians-Univ. Muenchen (DE)); Dr SKIPSEY, Samuel Cadellin; GADRAT, Sebastien (CC-IN2P3 - Centre de Calcul (FR)); DAL PRA, Stefano (INFN); HARTMANN, Thomas (KIT - Karlsruhe Institute of Technology (DE))

Presenter: FORTI, Alessandra (University of Manchester (GB))

Session Classification: Track 4 Session

Track Classification: Track4: Middleware, software development and tools, experiment frameworks, tools for distributed computing