



Contribution ID: 199

Type: **Parallel**

## Multi-core processing and scheduling performance in CMS

*Monday, May 21, 2012 5:00 PM (25 minutes)*

Commodity hardware is going many-core. We might soon not be able to satisfy the job memory needs per core in the current single-core processing model in High Energy Physics. In addition, an ever increasing number of independent and incoherent jobs running on the same physical hardware not sharing resources might significantly affect processing performance. It will be essential to effectively utilize the multi-core architecture.

CMS has incorporated support for multi-core processing in the event processing framework and the workload management system. Multi-core processing jobs share common data in memory, such as the code libraries, detector geometry and conditions data, resulting in a much lower memory usage than standard single-core independent jobs.

Exploiting this new processing model requires a new model in computing resource allocation, departing from the standard single-core allocation for a job. The experiment job management system needs to have control over a larger quantum of resource since multi-core aware jobs require the scheduling of multiple cores simultaneously. CMS is exploring the approach of using whole nodes as unit in the workload management system where all cores of a node are allocated to a multi-core job. Whole-node scheduling allows for optimization of the data/workflow management (e.g. I/O caching, local merging) but efficient utilization of all scheduled cores is challenging. Dedicated whole-node queues have been setup at all Tier-1 centers for exploring multi-core processing workflows in CMS.

We will present the evaluation of the performance scheduling and executing multi-core workflows in whole-node queues compared to the standard single-core processing workflows.

**Primary authors:** Dr EVANS, Dave (Fermi National Accelerator Lab. (US)); Dr HERNANDEZ CALAMA, Jose (Centro de Investigaciones Energ. Medioambientales y Tecn. - (ES); Dr FOULKES, Steve (Fermi National Accelerator Lab. (US))

**Presenter:** Dr HERNANDEZ CALAMA, Jose (Centro de Investigaciones Energ. Medioambientales y Tecn. - (ES)

**Session Classification:** Distributed Processing and Analysis on Grids and Clouds

**Track Classification:** Distributed Processing and Analysis on Grids and Clouds (track 3)