20th International Conference on Computing in High Energy and Nuclear Physics (CHEP2013)

Contribution ID: **438** Type: **Poster presentation**

Leveraging HPC resources for High Energy Physics

Monday, 14 October 2013 15:00 (45 minutes)

High Performance Computing (HPC) provides unprecedented computing power for a diverse range of scientific applications. As of November 2012, over 20 supercomputers deliver petaflop peak performance with the expectation of "exascale" technologies available in the next 5 years. Despite the sizeable computing resources on offer there are a number of technical barriers that limit the use of HPC resources for High Energy Physics applications. HPC facilities have traditionally opted for specialised hardware architectures and favoured tightly coupled parallel MPI-based workloads rather than the high throughput commodity computing model typically used in HEP.

However, more recent HPC facilities use x86-based architectures managed by Linux-based operating systems which could potentially allow unmodified HEP software to be run on supercomputers. There is now renewed interest from both the LHC experiments and the HPC community to accommodate data analysis and event simulation production on HPC facilities either from a dedicated resource share or from opportunistic use during low utilisation periods. If existing job scheduling and execution frameworks used by the LHC experiments could be successfully adapted for HPC use it would significantly increase the total amount of computing resources available.

A feasibility study into the use of LHC software in an HPC environment will be presented. The HECToR supercomputer in the UK and the SuperMUC supercomputer in Germany will be used as demonstrators. The challenges faced from the perspective of software execution and the interaction with existing LHC distributed computing environments will be highlighted using software typical in data analysis workflows. In particular, we will discuss how tighter restrictions on HPC worker node access could limit the functionality of existing middleware and how this can be potentially adapted for external job submission and more efficient job scheduling.

Summary

Primary author: WASHBROOK, Andrew John (University of Edinburgh (GB))

Co-author: Dr WALKER, Rodney (Ludwig-Maximilians-Univ. Muenchen (DE))

Presenter: WASHBROOK, Andrew John (University of Edinburgh (GB))

Session Classification: Poster presentations

Track Classification: Distributed Processing and Data Handling A: Infrastructure, Sites, and Virtualization