



Contribution ID: 451

Type: **Poster presentation**

Maximising job throughput using Hyper-Threading

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

The WLCG uses HEP-SPEC as its benchmark for measuring CPU performance. This provides a consistent and repeatable CPU benchmark to describe experiment requirements, lab commitments and existing resources. However while HEP-SPEC has been customized to represent WLCG applications it is not a perfect measure.

The Rutherford Appleton Laboratory (RAL), is the UK Tier 1 site and provides CPU and disk resources for the four largest LHC experiments as well as to numerous other experiments.

Recent generations of hardware procurement at RAL have included CPUs with Hyper-Threading. Previous studies have shown that as the number of logical cores being used increases, the measured HEP-SPEC will also increase but by increasingly smaller amounts. The more jobs that are running, the higher the chance that there will be contention on other resources which will cause jobs to slow down. It is therefore not obvious what is the optimal number of jobs to run on a Hyper-Threaded machine.

This paper details the work done to maximize job throughput at RAL. Over the course of several months different machine configurations were tested at RAL to see the impact on real job throughput. The results have allowed RAL to maximize job throughput while also accurately reporting the available HEP-SPEC and provided useful information for future procurements.

Summary

Primary authors: DEWHURST, Alastair (STFC - Science & Technology Facilities Council (GB)); ZILASKOS, DIMITRIOS (STFC)

Presenter: ZILASKOS, DIMITRIOS (STFC)

Session Classification: Poster presentations

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