OPTIMISING MULTI-USER AND MULTI-APPLICATION HPC SYSTEM UTILISATION USING EFFECTIVE QUEUE MANAGEMENT

I.Kureshi, V. Holmes, S. Liang, and D. Cooke
University of Huddersfield, Queensgate, Huddersfield HD1 3DH, UK

ABSTRACT

As the evolutionary cycle of computers continues, more and more organisations are deploying large computational systems in their data centers to be used by many users – a paradigm that existed up to the early 90's. IT managers are moving away from the model of purchasing the most powerful desktop/workstation available and placing it under a user's desk. These types of systems are useful when being utilised for their intended purpose. However, for a majority of the time they will be used for internet and email. A more cost effective mechanism to meet computational requirements is to purchase a significantly powerful system such as a computer cluster and allow multiple users to access it.

These multi-user systems usually employ a batch queuing system to help manage jobs. Normally operating with a first come first served policy, coupled with a scheduling system, the batch system can be adapted to provide a more efficient service. Service level agreements (SLA) or fair use policies (FUP) can be effectively enforced through this scheme, meeting a basic quality of service (QoS) across the board to all users. When making scheduling decisions a job scheduler needs to know what resources exist and are they available, how many resources does the particular job require, and how long does the job require these resources for. Traditionally it is left to the end user to provide this information to the scheduler, but when some information is left out it could lead to a scheduler not having enough information and reverting to a first come first served method of processing. Worse case scenarios exist where a job gets stuck in a queue indefinitely or the system operates in such a way that jobs submitted first, which actually need the least resources for the least amount of time running at the end of the queue. This leads to end user frustration and a bad QoS.

The High Performance Computing Resource Centre at the University of Huddersfield aims to provide the academic and research community with an effective and robust HPC system. The system cannot be optimised for a single piece of code but has to be kept flexible to meet the diverse needs of the research community. Effective queue and scheduler management has provided guaranteed QoS to each end user. This poster will outline these optimisation methods using the Torque Batch Queuing System and the Maui Scheduling System.

Keywords research computing, HPC, high performance computing, parallel processing, distributed systems, middleware, job scheduler, batch queuing, PBS, MAUI, TORQUE, OSCAR