



Contribution ID: 103

Type: **Poster**

ASC: an Adaptive Scheduling Controller

Monday 12 April 2010 18:21 (3 minutes)

The deployment, management and TCO of large computing environments always involve huge investments. These systems, once in production, have to meet the needs of users belonging to large and heterogeneous communities: only an efficient and effective use of these systems can repay the investment made. In this context, we report the experience made to design, implement and validate an adaptive scheduling system able to reconfigure itself when a lack in utilization efficiency occurs.

Detailed analysis

The adaptive scheduling controller (ASC) relies on top of a Maui-Torque scheduling system and has been developed following the following steps: (i) we have identified a set of Maui key-parameters, related to a combination of fairshare, reservation, preemption and backfill mechanisms, used to achieve an efficient and effective use of the system; (ii) we have evaluated the system behavior with respect to some key-statistics (queue waiting time, jobs throughput, resource usage, and so on); (iii) we have developed a control loop that uses information about the key-statistics and the desired performance profile in order to dynamically define a new set of Maui key-parameters values.

The default profile of the ASC control loop, based on automated log analysis and neural network techniques, can be chosen among a set of available profiles, each one identifies a target class of applications/users (e.g., parallel jobs, multi-thread jobs, concurrent jobs, and so on).

Conclusions and Future Work

Here we describe the work made to devise an adaptive scheduling controller (ASC), which aims to gain a balanced, efficient and effective use of the computing resources by heterogeneous communities.

Actually ASC has been deployed for Maui-Torque scheduling system, but it can be easily implemented on the top of other scheduling systems (e.g., LSF, PBS/Moab, and so on).

Impact

This work has been deployed and validated on computational resources of the University of Naples Federico II, acquired in the context of PON "S.Co.P.E." Italian National project. The resources are shared among three different contexts all based on gLite middleware: EGEE, Southern Italian and metropolitan GRIDs.

Due to the heterogeneity of the user community, the computational resources are used both for traditional GRID jobs and for HPC applications.

The adaptive scheduling represents an appealing solution to gain the needed trade-off among the needs of these, usually contrasting, class of applications.

We have validated the system by tests both with "driven load" and real production load (i.e., 100 Kjobs/month on about 2000 CPU and HPC jobs required about a 10% of resource usage).

The whole user community has experienced a good level of satisfaction. In particular, HPC community, usually penalized by a general-purpose scheduler configuration, registered an improvement.

Keywords

adaptive systems, job scheduler, resource management systems, log analysis, neural networks.

URL for further information

<http://www.scope.unina.it/C2/scheduling/default.aspx>

Authors: Dr DORIA, Alessandra (INFN); Dr DI MARTINO, Catello (University of Naples Federico II); Dr BOTTALICO, Davide (University of Naples Federico II); Dr BARONE, Giovanni Battista (University of Naples Federico II); Dr D'ANGELO, Giovanni (University of Naples Federico II); Dr CARRACCIUOLO, Luisa (CNR); Dr BOCCIA, Vania (University of Naples Federico II)

Co-authors: Dr ESPOSITO, Christian (University of Naples Federico II); Dr BUSIELLO, Gianluca (CEINGE); Dr CANTELE, Giovanni (CNR); Dr VITAGLIANO, Giuseppe (University of Naples Federico II); Dr POLLIO, Maurizio (University of Naples Federico II); Dr PETRILLO, Mauro (CEINGE); Dr PARDI, Silvio (INFN)

Presenter: Dr BARONE, Giovanni Battista (University of Naples Federico II)

Session Classification: Poster session

Track Classification: Software services exploiting and/or extending grid middleware (gLite, ARC, UNICORE etc)