Contribution ID: **22**   Type: **Oral**

Modeling the EGEE latency to optimize job time-outs

Tuesday, 12 February 2008 11:10 (20 minutes)

Jobs submitted to a production grid infrastructure are impacted by a variable delay resulting from the grid submission cost (middleware overhead and queuing time). The actual execution time of a job will depend on the process execution time, which can be known through benchmarks, and the variable grid latency duration, which is difficult to anticipate due to the complexity of the grid infrastructure and the variable load patterns it is enduring. We aim at estimating the grid latency through a probabilistic approach that is well adapted to complex system modeling. We derived a model of the expected execution time of a job function of the time-out value in a time-outing and resubmission setting. To follow on the variable load conditions, a monitoring service sends regular probe jobs to the infrastructure and measures their latency duration. This information is injected into the model and a numeric minimization provides the time-out value that minimizes the expected execution time.

3. Impact

A study on the EGEE grid workload pattern has shown that the latency endured by jobs follows heavy-tail distributions. Consequently, a non-negligible fraction of the jobs duration is likely to encounter very long latencies which are penalizing multi-job applications dramatically. Setting up the time-out strategies protects the application from these faults while introducing a very light overhead. The time-out estimation service is currently a prototype deployed on top of the EGEE middleware. The monitoring activity uses the workload management system to submit and monitor the jobs durations. The model computation is a lightweight numerical integration that can be integrated in any application. When their time-out expires, the application has to cancel and resubmit the faulty jobs to avoid abnormal computation times. An interesting perspective would be direct access to the RB logs to avoid application-level probing of the infrastructure.

URL for further information:

http://www.i3s.unice.fr/~glatard/publis/ccgrid07.pdf

4. Conclusions / Future plans

The model was tested on the EGEE production infrastructure using thousands of probe jobs over hours of execution. A 2.5% faulty jobs ratio was measured. Recovering from these faults by time-outing protects the application from unbounded execution time. The model can be adapted to more or less reliable system by varying the outlier ratio. Simulation of a fault-less system (e.g. cluster) with similar load conditions than the grid show that a minimum speed-up of 1.36 is achieved.

Provide a set of generic keywords that define your contribution (e.g. Data Management, Workflows, High Energy Physics)

Workload management, time-outing strategy

1. Short overview

Applications submitting a large number of jobs to the grid infrastructure have to consider and recover from faulty jobs that are due to system failures or abnormally long job durations. A simple time-outing and re-submission strategy protects the application from very long durations in case outliers happen. However, determining the time-out value is not straight forward, especially for shorter jobs, as their execution time significantly depends on the grid workload conditions.

Primary authors: Dr LINGRAND, Diane (UNSA); Dr MONTAGNAT, Johan (CNRS); Dr GLATARD, Tristan (CNRS)

Presenter: Dr LINGRAND, Diane (UNSA)

Session Classification: From research to production grids: interaction with the Grid'5000 initiative

Track Classification: Existing or Prospective Grid Services