



Contribution ID: 89

Type: Poster

Performance evaluation of the Estimated Response Time strategy: tools, methods and an experiment

Monday, April 12, 2010 6:06 PM (3 minutes)

An extensive body of research focuses on economic and intelligent scheduling models. Conversely, the gLite matchmaking engine adopts an agnostic approach for estimating the waiting time of incoming jobs, derived from the Copernican principle: "this job is not special". An open question is the evaluation of this minimalist strategy. This work reports on the creation of the software tools required for this evaluation, the methodology and criteria, and presents preliminary results using the Grid Observatory logs of the Information System (IS) and PBS scheduler of GRIF/LAL.

Detailed analysis

gLite maps jobs to Computing Elements (CEs) by considering the categorical requests and breaking ties along the Estimated Response Time (ERT), an estimation of the queuing delay published by the CEs in the IS. Evaluating the accuracy of the ERT computation requires comparing it along an extended period with the actual response time (ART). The ART is readily available from the logs of the local schedulers. We created a logging system of the IS for the Grid Observatory (GO). To cope with the massive redundancy, IS logs are stored in compressed diff format, with a reference snapshot each day, and diff at 15 mn period. For the evaluation, Perl scripts use regular expressions to extract ERT from the patched (reciprocal of diff) IS logs, and ART from the scheduler logs. Next, we use the Object capability of Matlab to create an efficient data representation (CE, VoView, Site, with matched ART and ERT), with flexible functionalities (subsampling, elementary or advanced statistics, plotting, outliers management). A first usage of this representation is to exhibit rare but repetitive abnormal patterns where relatively long-lasting spikes of load are correlated with low ERT.

Conclusions and Future Work

The evaluation criteria should integrate both the statistics about the reliability of the estimation (BQP concept), and a parameter measuring the operational impact of the prediction. We are exploring a description by the receiver operating characteristic (ROC) curve that segments the ERT along the load regimes. In this work, the analysis is aggregated over the whole available history. Future work will combine the segmentation method developed in the GO and this description in order to refine the performance evaluations along time-differentiated activity regimes.

Impact

The process described so far is not fully satisfactory: 1) the fixed frequency of IS logs is not adapted to the varying job arrival rate, 2) the IS ERT might not be the value actually used by the WMS, and 3) nor of the best one that the CE could provide. Recording the used ERT in the LB would provide much more accurate information with respect to 1) and 2). With the information we had, we decided to estimate the ERT of a job by the last published ERT for its target CE. Thus, we get irregular time-series, with a (ERT, ART) pair at each job arrival. The tool allows easy search for correlation at different lags. We observed a strongly differentiated behavior between the two most heavily loaded queues: Atlas has a nearly flat correlation landscape, while

Biomed shows a peak at lag 20. The main issue for evaluation lies in defining the objective function. The Batch Queue Predictor (BQP) initiative accurately pointed that synthetic indicators e.g. Mean Squared Error, or the correlation coefficient, are misleading: for instance, the (accurate) estimation of a null ERT for an empty CE is not very informative, and should not mask the above-mentioned spikes.

Keywords

scheduling, estimation

URL for further information

www.grid-observatory.org

Authors: Mr CADY, Alain (LAL); Prof. GERMAIN-RENAUD, Cecile (LRI)

Presenter: Mr CADY, Alain (LAL)

Session Classification: Poster session

Track Classification: Scientific results obtained using distributed computing technologies