A STATISTICAL ANALYSIS OF JOB PERFORMANCE WITHIN THE LCG GRID

M. Aggarwal[⊥], D. Colling[#], B. McEvoy[‡], G. Moont[†], O. van der Aa[*]

Abstract

In this paper the primary motivation is to analyze LCG performance by doing a statistical analysis of the lifecycles of all jobs submitted to it. In this paper we define metrics that will describe typical job lifecycles. The statistical analysis of these metrics enables us to gain insight into the workload management characteristics of the LCG Grid [2].

INTRODUCTION

The LCG is an operational Grid currently running at 136 sites in 36 countries, offering its users access to nearly 14,000 CPUs and approximately 8PB of storage [1] . Monitoring the state and performance of such a system is challenging but vital to successful operation. The four main LHC Virtual Organisations (VO) are alice, atlas, lhcb, and cms. The dataset consists of more than 3 million jobs submitted through Resource Brokers (RB). In this paper we have not considered jobs submitted by local users without using a Resource Broker. We know that some of the atlas production jobs are not submitted through Resource Brokers.

DATASET

The dataset has been obtained through direct queries to mySQL databases of RB from Sept 2005 to Jan 2006. Currently 28 RB are being monitored. The query structure consist of all jobs that had an event in the last minute. Events are retrieved from the Compute Element (CE) and Worker Nodes (WN) information. Using this data, a complete XML document, with a description of all the jobs, is generated. All finished jobs, are removed after 2 hours. A summary file is generated for all the finished jobs. In order to analyze the dataset, the data is imported into ROOT [3] to perform the analysis. In the analysis each job is characterized by a set of variables that we define in the next section.

A total of more than $3x10^6$ jobs where collected. This is shown on Fig. 1, on average ~20000 jobs are submitted each day.

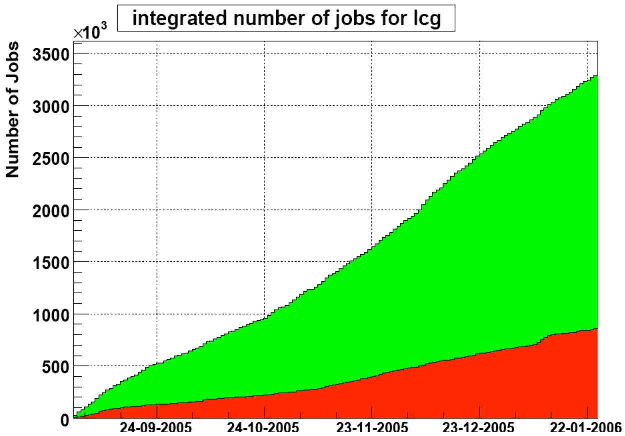

Figure 1: Integrated number of jobs over the analysis period. Successful jobs are shown in green while failed jobs are shown in red.

Variables definition

The variables that we discuss in this paper are defined below.

- *CE*: Fully qualified name of the Compute Element where the job has been submitted.
- *VO:* Virtual organization submitting the job.
- *RB:* Resource Broker.
- *JRT:* Job Run Time. Time in seconds the job has been running.
- *JTT:* Job Total Time. Time in seconds elapsed

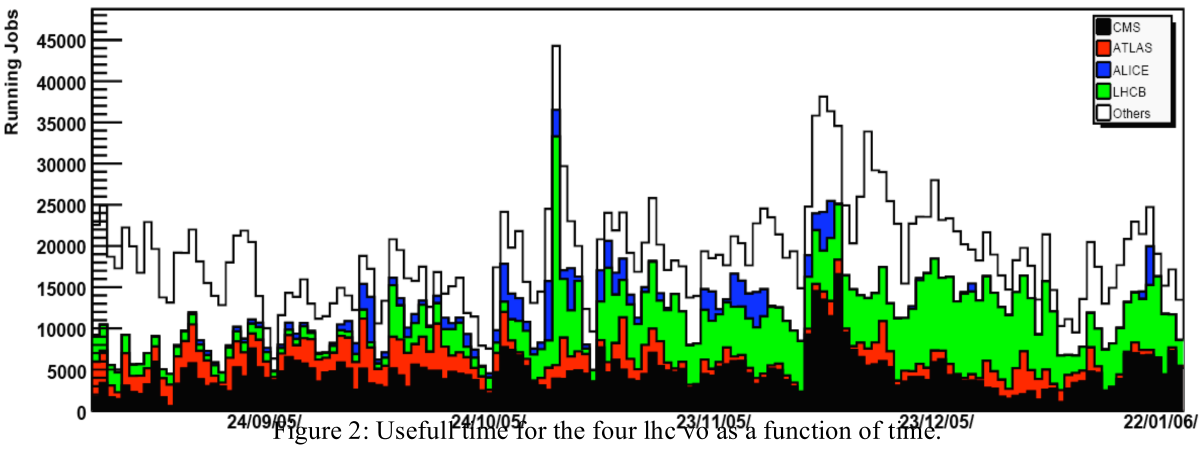

Figure 2: Useful time for the four lhc vo as a function of time.

from the event of the job submission to the final job state. The final job state can be (done/cancelled/failed/aborted).

- *EFF:* Efficiency computed as JRT/JTT. Measure the overhead in the job total time due to the LCG middleware. A low value will indicate that most of the total job time is spend in middleware activities.
- *UT:* Useful Time. The fraction of time used by successful jobs.
- *MT:* Match Time. Time for the resource broker to match the job requirements with a computing element.

JOB PERFORMANCE METRICS

The job performance metrics can be seen in two perspectives: the user of the system using the Grid resources, and the resource provider providing the compute resource to the Grid.

User perspective

The statistical analysis shows that for all the four VO's, the number of jobs, having a runtime between 10 to 20 hours, is substantial. CMS has the highest number of job processing hours within the period of measurement.

A measure of the Grid usage is given by the Grid Load which is the number of running jobs as a function of time. This is shown on Fig 2 for the four LHC vo and all the other vo grouped in one set. We observe that the lhcb activity has increased significantly over the end of last week while cms and atlas had a more steady profile. In total cms, atlas, lhcb and alice have respectively consumed $3.5x10^6$ h, $0.97x10^6$ h, $0.89x10^6$h, $0.43x10^6$h of computing time. The remaining VO's have consumed 1.3x106 hours.

The JRT distribution gives insight on the usage patterns by the VO. On Fig. 3 we clearly see a two peak structure for lhcb which is not seen for cms. This is due to the different nature of the lhcb and cms jobs. cms activity over last year has mainly been analysis by individual users. The lhcb strategy is to send a lot of short jobs that monitor the site status. Once a resource has been found to match the job requirements, the real job is being pulled to perform the actual computing. This explains the peak at low JRT for lhcb on Fig. 3.

Resource provider perspective

From the perspective of the compute element they are not equally targeted by jobs. lhcb is targeting all CE with their test jobs while cms is targeting a small subset of the available CE. This is because cms analysis is performed only on sites that have the dataset they need. Not all sites have those dataset. lhcb is mainly doing production, hence all sites can be used .

Generally we find that all the sites are not equally well utilized. Tier1 sites handle most of the jobs, while Tier2 sites are highly under utilized.

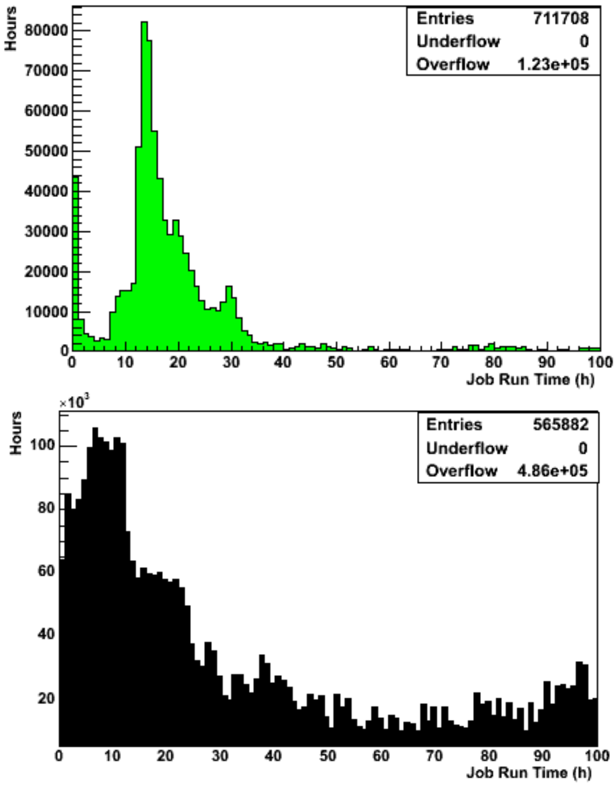

Figure 3: Distribution of job run time weighted by job run time. The distribution on the top is for the lhcb vo, the bottom one is for cms.

We can measure how well the provided resource has been used. The measure we use is the useful time (UT). This measures how well the job used the resources by computing the ratio of successful compute hours with the total compute hours. A value lower than one indicates the resource has been utilized to run jobs that where not successful. The UT is shown on Fig. 4 for the atlas and cms vo as a function of time. The mean value is above 80% but they are significant dips.

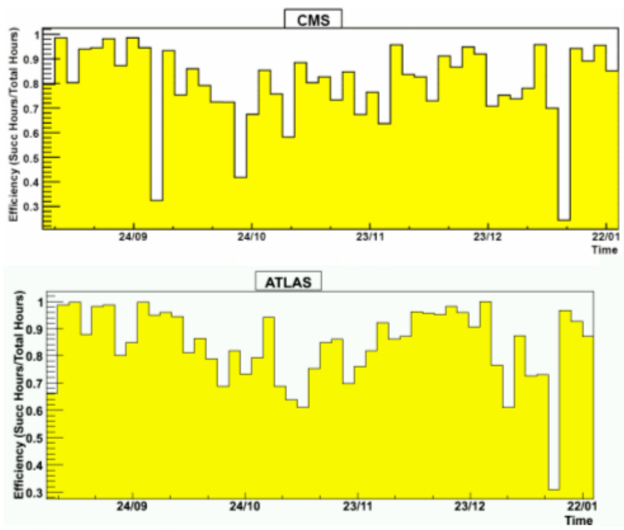

Figure 4: Useful Time as a function of time for the atlas and cms vo.

Resource broker load

An important feature of resource brokers is their throughput. We define the throughput as the number of jobs that can be submitted by unit of time without observing an increase of the match time of the RB. The match time is the time it takes for a job to be matched to a computing element. We have measured the throughput for several RB in the LCG grid. We have observed that a small fraction of the available resource brokers are being used to send the majority of jobs. Hence the load on those resource brokers can be high.

We can see on Fig. 5 that the minimum match time starts to increase when the load is higher than 10 jobs/sec. As the number of jobs per second increases from 10 jobs per second to 50 jobs per second, the minimum match time increases nearly linearly from 20 seconds to 140 seconds. Following our definition of throughput, this RB is capable of 10jobs/s. In particular the mean match Time is 36 seconds.

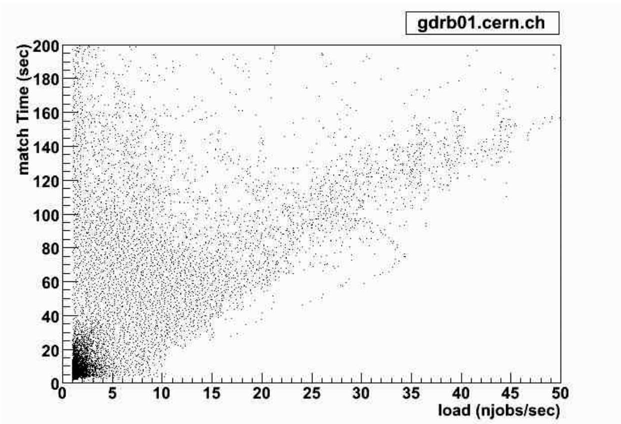

Figure 5: match Time as a function of the load on a cern resource broker (gdrb01.cern.ch)

CONCLUSIONS

We have discussed distinct usage patterns emerging from each VO from Sept 2005 to Jan 2006. We have measured that the RB minimum match time increases linearly when more than 10 per second are submitted.

REFERENCES

[1] http://www.gridpp.ac.uk

[2] Crosby P, Colling D,Waters D, Efficiency of resource brokering in grids for high-energy physics computing, IEEE Transactions on Nuclear Science, 2004, Vol: 51, Pages: 884 - 891, ISSN: 0018-9499

[3] Rene Brun and Fons Rademakers, *ROOT - An Object Oriented Data Analysis Framework*, Proceedings AIHENP'96 Workshop, Lausanne, Sep. 1996, Nucl. Inst. & Meth. in Phys. Res. A 389 (1997) 81-86. See also http://root.cern.ch/.

* o.van-der-aa@imperial.ac.uk, Imperial College London.
# d.colling@imperial.ac.uk, Imperial College London
⊥ m.aggarwal@imperial.ac.uk, Imperial College London
‡ b.mcevoy@imperial.ac.uk, Imperial College London
† G.Moont@imperial.ac.uk, Imperial College London