Benchmarking and Normalization of Compute Resources in PL-Grid Infrastructure

Wojciech Ziajka, Marcin Radecki, Paweł Szepieniec and Tomasz Szepieniec

AGH University of Science and Technology, ACC CYFRONET AGH
ul. Nawojki 11, 30-950 Kraków, Poland
emails: {w.ziajka,m.radecki}@cyfronet.pl

Keywords: compute resources, normalization, benchmarking, grid accounting

1. Introduction

Compute resources of PL-Grid Infrastructure [1] are provided to users as a consistent pool with unified access, however they are located in different compute centres and are of course heterogeneous. It is clear that given CPU consumption on some less powerful system does not have the same value for the user as an equal consumption on an up-to-date processors [4]. To solve such issue a usage normalization need to be defined. Normalized metrics are indispensable to express the users demand for compute power in heterogeneous environment. A process of establishing normalization factor for a given compute resource is usually based on some benchmark. Using this factor for a given resource the users can evaluate demanded compute hours, especially if they do not have direct access to it. Of course this is as much precise as normalization factors reflects parameters essential for the users' application.

Measuring a CPU peak performance differs from measuring a production system where compute power available for a user may wary with time, especially on multi-core nodes. The reasons may be that other users processes heavily use the memory bus or the CPU frequency going down when temperature raises. To normalize usage consumed by each job we need to establish the factor for each node, but it is not feasible in Grid to benchmark each node.

In this paper we present our approach to normalization of accounting data and benchmarking of PL-Grid Infrastructure nodes. We show how we calculate and use normalization factors. Some initial results from running a prototype in the production infrastructure are also presented.

The largest Grid infrastructures dealt with the normalization and benchmarking in different ways. In European Grid coordinated by EGI.eu an single value of average core power is published for a whole site [2]. They use HEP-SPEC for benchmarking. In US infrastructure called XSEDE, users express their compute needs in so called Service Units and use static conversion between XSEDE systems [3].

2. Description of a problem solution

To avoid a need of benchmarking each node in Grid we introduce a concept of a *node class.* A node class is constituted by all machines at a site that have the same key performance characteristics. A normalization factor is established for a node class. We run a benchmark and relate it to the result on some reference machine to establish the factor. Using historical results we calculate the current factor value for a node class based on an exponential moving average. Our priority for normalization factors is to reflect the users' experience with the infrastructure, thus we run benchmarking jobs in the same way and conditions as users run their jobs. The benchmarks are submitted to all systems from a module in Benchmarking System. Our benchmark code is based on High-Performance Linpack which is also used to rank machines for Top500 list. Benchmark results are collected back through a Message Bus (JMS) and normalization factors are computed in Benchmarking System. Module called BS_Agent reports a node class for each node in the Grid.

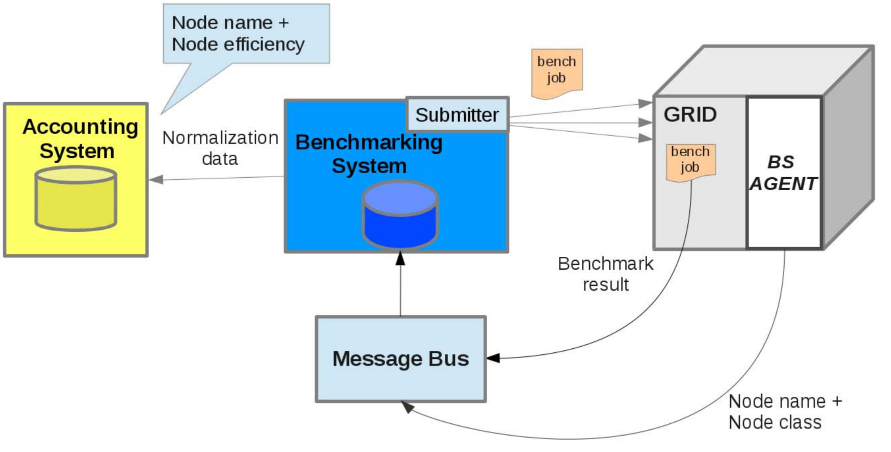

Fig. 1. Interactions between Accounting, Normalization and Benchmarking Systems in PL-Grid Infrastructure.

3. Results

A prototype of normalization system has been running for 3 months in PL-Grid production infrastructure. The results show that normalization factors for the same node class, even located in different sites, are quite similar. As expected we can observe power variations in benchmarked node classes. Frequently this is related to another users activity on a multi-core node. To picture the scale of that effect we show how a number of benchmarks running on the same node influences the result of an individual benchmark.

4. Conclusions and future work

Benchmarking and normalization of production Grid infrastructure is troublesome but can be dealt with by defining some compromise between system accuracy and what is sufficient for the users. In future we plan to extend the benchmark to be more comprehensive: include RAM and I/O characteristics. Another experiment on how much improvement may be made by tuning compile-time options is on our agenda.

References

1. M. Bubak, T. Szepieniec, K. Wiatr (eds): Building a National Distributed e-Infrastructure - PL-Grid, LNCS Vol. 7136, Springer, 2012.
2. EGI wiki, Transition to HEP SPEC - new CPU Benchmark, `https://wiki.egi.eu/wiki/HEP_SPEC06`, (accessed Sept. 2012).
3. XSEDE Service Unit Converter, `https://www.xsede.org/su-converter`, (accessed Sept. 2012).
4. R. Buyya, K. Bubendorfer: Market-Oriented Grid and Utility Computing. Wiley, 2009.