Assessment Strategies in OSG

OSG-1032-v2

March 4, 2011

1. Introduction

OSG as an infrastructure is a highly instrumented facility with multiple information sources providing data relevant to assessing performance and effectiveness towards meeting fundamental OSG program goals. Key performance metrics have been used to programmatically assess areas of service operational stability, reliability, and throughput measures for jobs, CPU usage, data transfers and storage capacities. These (mostly facility-centric) measures have been crucial to guiding the development of the infrastructure into a mature, stable platform for science requiring high throughput access to distributed resources. In OSG' we intend to extend the program of metrics collection, synthesis, analysis and reporting in two key areas: 1) measurement of core processes relating to project execution leading to 2) measures which more broadly assess the impact of OSG in advancing the scientific and engineering goals of its stakeholder organizations.

Our approach will involve examining each principal objective of the program and identifying they key data sources and derived measures that can be used to "tell the story" and recognize the impacts. For example, OSG advances science through open distributed computing. This is a core objective. What are the key measures that support that assertion and how can better synthesize, report and visualize the data to assess the relevance and reach of the impacts?

The product of scientific research is knowledge and information. Classically one can measure scientific publications and citations, citations per author, etc, in refereed journals as an indicator of impact. In the larger picture of metrics usage as criteria for evaluating academic merit, researchers have developed measures such as h-index [1], a measure of productivity and impact of a researcher. But there are many others [2] (see bibliometrics Fig. 1) and a goal will be to present the correct data to scientometric researchers so as to accurately measure the utility and role of the OSG CI in producing science. Note that for OSG, this presents a particular challenge given the shared accounting metrics among authors and the (extreme) case of multiple authors.

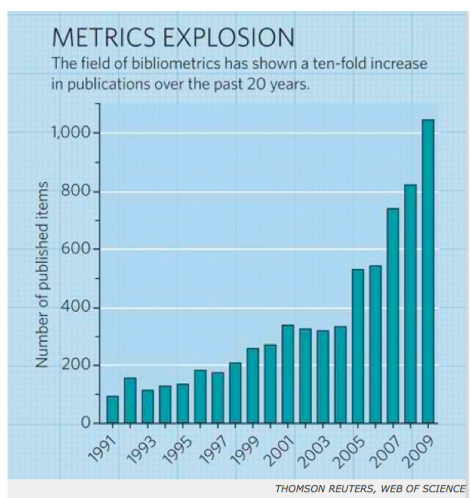

Figure 1 Publications in bibliometrics (Nature 465, 864-866 (2010))

Only a few years ago students could expect to access computing resources from only a few sources – their desktop, their departmental cluster, perhaps centralized facilities on campus. Today we have nationwide infrastructures such as TeraGrid for HPC-centric applications submitted via PI-allocations [3]. In DHTC, we see thousands of (VO-authorized) students accessing hundreds of sites (OSG proper, and the via federated EGEE/EGI) accessing shared resources employing a variety of tools, datasets and modalities. A levelling of the playing field has occurred in terms of capabilities for students among institutions collaborating in the same science program. Thus one could envision a user "capability index" as a productivity measure, tracked over the project and in historical terms. One could speak of "PhD throughput" – as a measure of the number of PhD's enabled by the OSG CI and in which fields.

All of these notions directly related to demonstrating to our stakeholders, management, users, and importantly our funding benefactors the effectiveness of the infrastructure in easy to access and understand terms. In particular can we respond to providing data on metrics and impact factors as required by agency initiatives such as STAR METRICS — *Science and Technology for America's Reinvestment: Measuring the Effect of Research on Innovation, Competitiveness and Science* — a multi-agency venture led by the National Institutes of Health, the National Science Foundation (NSF), and the White House Office of Science and Technology Policy (OSTP) [4]. The second phase of the program is focused on economic growth (as measured through indicators such as patents and business start-ups), workforce outcomes as measured by student mobility into the workforce and employment markers, and most significantly for OSG', scientific knowledge as measured by publications and citations.

For stakeholders, in particular users of the OSG CI whose work directly benefits, OSG' assessment services could provide and important role in peer-demonstration of the effectiveness of various approaches to computing [5].

2. Metrics Gathering and Synthesis

As mentioned the OSG CI is already a highly instrumented infrastructure. The challenge in OSG' will be to gather and pre-process key data metrics for query optimization relevant to impact analysis. We do not envision a need for building a new technology stack to acquire, synthesize and report this data; on the other hand we will work with the Software team and application developers to augment existing reporting/tracing/accounting capabilities as needed. We plan to create a (lightweight) repository of derived metrics data objects suitable for programmatic reporting and advanced assessment analyses and visualization. For example, a database that associates publications, usage, application-types, VOs and users, science topics would be constructed. There may be tweaks at various points our infrastructure, but these should be in the category of 'feature enhancement' as opposed to development. Along these lines we have opened discussions with Prof. Krishna P.C. Madhavan of Purdue University who specializes in developing cyberinfrastructure tools for measuring research behaviors and science impact, especially within the context of engineering education. In particular the development of mixed-used dashboards that combine evidence sources from integrated reporting systems, surveys, user experience analyses, interview and questionnaires.

3. Reporting and Visualization

3.1. For Scientific Users

Provide users with a clear picture of usage and benefit from the OSG CI, with comparison to peer and facility usage.

3.2. For Systems Administrators and Resource Providers

Provide systems administrators with view of performance and job behavior, local and peer sites. Provide resource providers with scientific impacts resulting from delivery of resources into the OSG CI. Show value of Campus Grid infrastructure to university communities and administrators.

3.3. For the Executive Team

Provide the ET the data, views, trend analyses needed for organizational efficiency, project management, budgets and processes.

3.4. For Agency Review

Providing the data and reporting to our funding agencies to gauge return on their investments.

4. References

[1] h-index, http://en.wikipedia.org/wiki/H-index, and http://arxiv.org/abs/physics/0508025

[2] See for example http://www.harzing.com/pop.htm#metrics

[3] Johan Bollen, Geoffrey Fox and Prashant Raj Singhal. How and where the TeraGrid supercomputing infrastructure benefits science. Journal of Informetrics, 2010 (doi:10.1016/j.joi.2010.09.004)

[4] STAR METRICS, http://nrc59.nas.edu/star_info2.cfm

[5] Assessing Assessment, Nature Volume: 465, Page: 845 (17 June 2010) doi:10.1038/465845a, http://www.nature.com/nature/journal/v465/n7300/full/465845a.html

Author:

Rob Gardner, University of Chicago

Revisions:

1	1/13/2011	RWG	Initial version
2	3/4/2011	RWG	OSG doc template format