

Analysis of the Current Use, Benefit, and Value of the Open Science Grid

R. Pordes¹ for the Open Science Grid Executive Board, J. Weichel
Fermilab, PO Box 500, Batavia, IL 60510

E-mail: ruth@fnal.gov

Abstract. The Open Science Grid usage has ramped up more than 25% in the past twelve months due to both the increase in throughput of the core stakeholders – US LHC, LIGO and Run II – and increase in usage by non-physics communities. It is important to understand the value collaborative projects, such as the OSG, contribute to the scientific community. This needs to be cognizant of the environment of commercial cloud offerings, the evolving and maturing middleware for grid based distributed computing, and the evolution in science and research dependence on computation. We present a first categorization of OSG value and analysis across several different aspects of the Consortium’s goals and activities. And lastly, we presents some of the upcoming challenges of LHC data analysis ramp up and our ongoing contributions to the World Wide LHC Computing Grid.

1. Introduction

The Open Science Grid[1] architecture enables contributing members to: 1) provide access to their computing and storage resources and/or software developments, 2) benefit from the use of and support for the common software stack and operational services, and 3) store, access and process their data on the ensemble of resources made accessible.

During 2008 the scale of the OSG has increased both in size and use by the high energy physics communities, including both the Tevatron experiments for large scale data simulations and the LHC experiments for simulation, processing and analysis. The usage has remained more or less constant for the non-physics users, including LIGO, protein structure prediction development, and molecular dynamics, with significant cycles of use by each of the individual communities.

During 2008 we have made an initial assessment of the benefit from and “value” of OSG to its members[2]. This will provide us with a mechanism to compare alternatives in the future – especially emerging commercial clouds. We know we don’t have a complete assessment with this first attempt -- we regard it as a work in progress.

2. Analysis of the Use

The following table shows the resources currently available through the OSG infrastructure:

	Number	Comments
Compute Elements	93	Linux and Microsoft clusters with Condor, SGE, PBS, LSF batch systems through the Gram 2 or Gram 4 job submission interfaces.

¹ To whom any correspondence should be addressed.

		~20 of these clusters currently support the “glexec” security component which allows secure use of “glide-in” or pilot job management.
Storage Elements	25	Disk and tape storage, mainly used by the owning community. One community (DZero) continues to benefit from the use of “opportunistic storage” on sites sharing their disk caches with OSG. OSG supports Bestman, Bestman/xrootd and dCache disk storage implementations, and through these HPSS and Enstore mass storage systems.
Number of individual Users	2500	The average number of users per week is ~300.
% of total ATLAS & CMS job throughput	>30%	Resources accessible through the OSG support US LHC simulated event generation, processing and analysis.
Software Releases of the OSG stacj	1 Major 10 Minor	The first major production release (V1.0) of the software stack was in June 2008. This included the baseline services needed for initial LHC data taking.
2008 Science publications based on OSG contributions	>120	Many from the CDF and D0 Tevatron experiments, 8 from STAR nuclear physics experiment, a few from LIGO and OSG itself, and ~five from the non-physics communities.

Table 1: Snapshot of core OSG parameters

The following were the major improvements in the technologies deployed in 2008 on the OSG[3]:

- Initial use of opportunistic storage for science output for the DZero experiment.
- Early adoption in physics and generalization of “overlay” job scheduling or “pilot” technologies
- Resource service validation framework and probes for monitoring site configurations and service availability
- Resource selection and matchmaking services.

Figure 1 shows the usage of the OSG over the past year and Figure 2 is a map of resources in the US accessible from the OSG infrastructure.

services enable OSG's partners and users to do. One of the goals is to quantify, where possible, the value components. This was done by quantifying each area into a dollar figure or effort or schedule time that can be translated to dollars. As discussions took place with the stakeholders, it became clear that some important benefits from OSG are intangible and broadly scientific and sociological in nature. One cannot easily quantify intangibles such as "broader influence on the community", "benefit of sustained multi-disciplinary collaboration", "conducive cooperation". This makes quantification into dollars less meaningful. The fundamental working model of OSG is collaboration and cooperation with its partners and users to enable all parties to accomplish their goals more effectively. This day-to-day collaborative model means it is difficult to tease out the value of OSG per se. OSG also contributes to a reduction in risk, which has value in reducing contingency costs of the individual organizations. We did not cover this in the initial version of the document.

We identified five areas of benefit. OSG which are summarized in the following subsections:

1. Supports collaborative research from small to large scales
2. Provides a sustained US cyber-infrastructure for scientists
3. Contributes to computer science and software body of knowledge
4. Sustains and enhances US expertise
5. Creates an environment for opportunistic computing

3.1 Supports collaborative research from small to large scales

OSG has specific responsibilities to support the science and distributed computing systems of the US ATLAS, US CMS and LIGO collaborations. OSG provides value in the underlying distributed facility, common services, security and software, on which the collaborations' systems are overlaid. OSG also contributes value by providing effort for joint activities to:

- Make increasing use of common and externally supported components, software and resources not owned by the experiment;
- Make designs and implementations more general so they can more easily adapt to the inevitable change in external software and operating environments over the multi-decade lifetime of the experiment;
- Facilitate sociological buy-in and change within the collaboration to increase acceptance and understanding of the principles of the OSG, use of the services offered by OSG, and through this increase the value of OSG to the collaboration.

3.2 Provides a sustained US cyber-infrastructure for scientists

OSG fosters commonality across communities. It provides, maintains, and evolves common software, procedures, organizational standards, and resource use policies that reduce the threshold to adoption by owners of computing resources.

OSG provides software distributions and support. The value of a common, centralized software distribution and support infrastructure is that packaging is done only once, patches and fixes need be applied only once, users have an understood common environment when their data and jobs land on a resource, and all communities get immediate access to new capabilities and services. The overhead for each user community is that the software stack contains more than the particular services they need, that turnaround time for changes and redistribution are not under their direct control, and that they rely on external expertise rather than have it directly at hand internally. One value provided is integrating heterogeneous components and associated libraries – all needed by service providers in a distributed system – and building them successfully on all needed thirteen environments. This typically takes between 1 and 4 FTE weeks for each new major version of or new piece of software.

	# Major releases of dependent components a year
--	---

Condor and the Globus core grid middleware	~2 Condor, ~1 Globus
Three storage service implementations	~2 each (storage is an important focus of WLCG in particular)
Information and job management services	~2 for each of 4 components
Security infrastructure	~2 for each of 4 components
Underlying common toolkits	~10 in total
Total Effort expended	35 – 140 FTE weeks (0.67-2.7 FTE)

Table 2: Rate of and Effort on New Software Releases

OSG operates common services and procedures across all supported communities, as well as providing a central, common interface to OSG's partners. This obviates the need for the experiments to each have their own ticketing system and 24x7 operational support team. The expertise of the central security team enables the tens of sites to leverage incident alarm, analysis and response, risk assessment, and mitigation; as well as common policies and procedures. For around the clock operations each collaboration would require ~3 FTEs for security and monitoring services. OSG also provides central performance metrics, service availability monitoring of compute and storage elements, and accounting services. There is some overhead of interacting with a central operations group rather than community specific ones. We make an initial estimate that each community saves 33% of the 3 FTEs that each community would otherwise need.

OSG provides an at-scale integration testbed, documentation, and site support[4] which provides value both for the members of the OSG and for external software developers. The testbed implements a complete grid infrastructure mimicking the services used in production. The process of validating new software includes ensuring adequate documentation and generating the OSG specific documentation for configuration, testing, and diagnosis.

The OSG security team provides operational security management, security and risk assessment and incident response across the infrastructure. The OSG security team and operations take ownership of security alerts and follow them through until resolved with any needed mitigations made available. The effort put into analyzing and responding to each incident varies enormously from a few (3) FTE days to weeks (10 FTE days). On average we receive information about an incident once a month. We estimate that for any given incident 10% of the (80) sites are affected (to date this matches the experience) and the effort of the central team saves 50% of the effort that each site would independently have to supply if there was no OSG team. OSG writes policies and agreements and works on ~5 policies a year. A policy takes between 2 and 4 FTE weeks of OSG staff time to develop.

3.3 Contributes to Computer Science and the Software Body of Knowledge

OSG contributes to the computer science and software body of knowledge fostering a broad collaborative community with in depth expertise and understanding of distributed computing in all aspects. OSG is a close collaboration at all levels between computer and domain scientists.. The main examples here are the Condor project, Globus, gLite and other common middleware developments, and the various security and storage software projects with whom OSG works.

OSG improves software for distributed computing. OSG has partnerships with software development groups that include computer science components in their developments. For each of these groups, OSG provides a testing and integration environment and sustained usage in production that provides proven utility for the software development groups. OSG helps identify problems in distributed computing that need study and invention and has initiated projects, internally and externally, directed at improving grid computing. OSG fosters a sustained community of experts by providing an interchange across the software development groups themselves, as well as between the development groups and the user communities. OSG benefits software development and computer science through:

- Use at the largest scales in the field.
- Direct exposure to user requirements and feedback to ensure effective development.
- Exposure to a broader community of software experts for design and implementation discussions.
- Practical influence to encourage adoption of more general tools and technologies by a broader community of users.
- Established process for distribution, testing, and support of the software to a wide audience.

3.4 Sustains and Enhances US Expertise

OSG sustains and enhances US expertise by contributing to national and international collaborative scientific efforts and advancing the state of the art of large scale distributed computing and cyber-infrastructure. Some of these contributions are presented in the previous section. Others are OSG: Gives universities easy access to distributed computing technologies to foster innovation; Trains students and the workforce in distributed computing technologies use and support - the OSG has given 1-3 day classes to 250 students and supported University classes at RIT and University of Missouri with >10 students each; Enhances US reputation via international participation, and increases ability of US to compete and contribute internationally.

3.5 Creates an Environment for Opportunistic Computing

OSG creates an environment for opportunistic computing by providing an infrastructure that enables users to use resources that their community does not own. OSG facilitates more effective use of the total installed base of resources². This “opportunistic use” results from the collaborative nature of the OSG Consortium. The large resource owners support this use pattern in return for the other perceived benefits offered by the OSG and to increase the broader return on their hardware and administrative investments. Small communities or individual researchers gain access to 2 orders of magnitude more computing throughput than they are able to achieve locally with their available funds We have developed an overall cost/delivered capacity scenario for a large and medium size resource center. The following table summarizes the costs using Fermilab and the University of Nebraska as representative centers. We were not, and cannot be, completely rigorous here. Our attempt is to provide some measure for discussion and comparisons.

	Fermi	Nebraska	Oct 1, 2007 - Oct 31, 2008
Cost Item	\$K/Year	\$K/Year	Notes/Definition
Facility \$/ CPU	\$92	\$117	Used as sanity check of Facility and Power Costs
System \$/ CPU	\$128	\$150	Used as sanity check of System Costs
Staff \$/ CPU	\$191	\$840	This is expected by economy of scale. (Nebraska doubled their capacity this year with no staff increase. With that increased capacity for next year, the staff cost/CPU will be \$388.)
Total Cost/ CPU Hour	\$0.047	\$0.126	For FY 2008
2009 \$ / CPU Hour	\$0.040	\$0.071	Estimate for FY 2009 (Assuming full year costs with current capacity)

Table 3: Summary of Costs at Large and Medium Facility

² There are 2 modes of such facilitation: brokering of agreements between users and resource owners to define commitments for shared use; and dynamic use of “at-the-time available” cycles and storage without prior commitment.

Another validation of these estimates is to look at the commercial market for computing cycles sold by companies like Amazon³ (Amazon Elastic Compute Cloud). They typically charge as low as \$0.10 per CPU hour. We would expect this to be roughly twice the actual cost of those hours. This matches very closely with Fermilab's cost of \$0.04 to \$0.05 per CPU hour. These costs can be used to quantify the value of opportunistic computing by the VOs. Overall, for the year November 15, 2007 through November 12, 2008, there were a total of 42.3M Opportunistic CPU hours delivered. This represents about 30% of the total CPU resource used in OSG during the period.

3.6 Summary of Quantitative Value

	Savings: Effort or \$ (annual)
Tangible Value:	
Direct benefit to 3 key stakeholders middleware, applications	5.3 FTE
Direct contributions to the WLCG	2.0 FTE
Up-front cost savings for new users to adapt to using OSG	0.8-2.5 FTE
Software packaging and distribution	8 -24 FTE
Operating the production infrastructure	7 FTE
Integration testbed and documentation	6-12 FTE
Documentation, training for administrators of new sites	2.5 FTE
Security incident response	0.4-1.3 FTE
Policy development	1.2-2.4 FTE
Direct benefit in tools developed for site administrators	0.5 FTE
Effort savings & direct benefit (intangible benefits discussed in the text)	33.7-59.5 FTE
Economic Value of Resource Sharing: (Equivalent value of opportunistic use)	\$2.1M - \$5.5M

Table 4: Summary of OSG Value and Benefit

4. Future Plans

OSG, working with the LHC experiments in the US, is committed to be an effective contributor to the World Wide LHC Computing Grid Collaboration for the foreseeable future. In the near term we plan to extend and adjust our services and support to meet the needs of the US LHC Tier-3 university facilities. There are currently about twenty sites already collaborating with the OSG, and there are expected to be several tens more in the next year or two. The scale and needs of these groups will vary widely and be significantly different in scope from the existing US LHC Tier-1 and Tier-2 sites. For example, some groups may act as end points for receiving experiment datasets but otherwise operate entirely independently from the rest of the experiment's distributed system, while others will be full partners both with other university departments locally and the OSG nationally.

During the next year we will also spend effort on better understanding the needs and technologies for the support of short-lived or "ad-hoc" virtual organizations as a next step in the engagement of the Spallation Neutron Source communities.

The core OSG services, technologies and policies allow for the use of virtualization, glide-ins or cloud technologies. Member communities in the are actively exploring the integration and interface issues to use commercial cloud computing, such as Amazon EC2, as an additional resource accessible through the OSG infrastructure. There are also independent evaluations underway of the use of virtual machine

³ Using Amazon as a comparative point to determine if we had done something wrong in the cost analysis. Clearly, they don't provide the exact environment that exists in OSG and they separately charge for storage and data movement that we have included into the CPU core hour. The fact that the numbers are comparable (assuming that Amazon does profit from the service) is all that we are trying to show.

technologies to improve the ease of use and reliability of running jobs across a heterogeneous infrastructure.

5. References

- [1] Project supported by the Department of Energy Office of Science SciDAC-2 program from the High Energy Physics, Nuclear Physics and Advanced Software and Computing Research programs, and the National Science Foundation Math and Physical Sciences, Office of CyberInfrastructure and Office of International Science and Engineering Directorates.
- [2] OSG Doc 813.
- [3] OSG Annual Report, 2008 <https://osg-docdb.opensciencegrid.org:440/cgi-bin/RetrieveFile?docid=815&version=1&filename=OSG%20DOE%20Report%20Dec-2008%20Final.pdf>
- [4] OSG Integration project <https://twiki.grid.iu.edu/bin/view/Integration/IntegrationProject>