Sustaining and Extending the Open Science Grid:
Science Innovation on a PetaScale Nationwide Facility
(DE-FC02-06ER41436)
SciDAC-2 Closeout Report

Lead Institution: University of Wisconsin-Madison
Madison, WI 53706

Lead Principle Investigator and DOE Contact: Miron Livny
Computer Science Department
1210 West Dayton Street
Madison, WI 53706
miron@cs.wisc.edu
608-262-0856

Office of Science Programs Address: High Energy Physics, Nuclear Physics, and Advanced Scientific Computing Research

Office of Science Program Office Technical Contacts: Lali Chatterjee, Ted Barnes, and Rich Carlson

Participating Institutions and Principle Investigators:

Boston University, James Shank
Brookhaven National Laboratory, Michael Ernst
California Institute of Technology, Kent Blackburn*
Clemson University, Sebastien Goasguen
Columbia University, Michael Tuts
Cornell University, Lawrence Gibbons
Fermi National Accelerator Laboratory, Ruth Pordes*
Harvard Medical School, Piotr Sliz
Information Sciences Institute (USC), Ewa Deelman
Indiana University, William Barnett
Lawrence Berkeley National Laboratory, Doug Olson
Renaissance Computing Institute (UNC), John McGee
SLAC National Laboratory, Robert Cowles
University of California at San Diego, Frank Wuerthwein
University of Chicago, Robert Gardner
University of Florida, Paul Avery*
University of Illinois, Shaowen Wang
University of Iowa, Shaowen Wang
University of Nebraska – Lincoln, David Swanson
University of Wisconsin, Miron Livny*

* Project Principle Investigator

1 Project Goals

Under this SciDAC-2 grant the project's goal was to stimulate new discoveries by providing scientists with *effective and dependable access to an unprecedented national distributed computational facility*: the Open Science Grid (OSG). We proposed to achieve this through the work of the Open Science Grid Consortium: a unique hands-on multi-disciplinary collaboration of scientists, software developers and providers of computing resources. Together the stakeholders in this consortium sustain and use a shared distributed computing environment that transforms simulation and experimental science in the US. The OSG consortium is an *open* collaboration that actively engages new research communities. We operate an *open* facility that brings together a broad spectrum of compute, storage, and networking resources and interfaces to other cyberinfrastructures, including the US XSEDE (previously TeraGrid), the European Grids for ESciencE (EGEE), as well as campus and regional grids. We leverage middleware provided by computer science groups, facility IT support organizations, and computing programs of application communities for the benefit of consortium members and the US national CI.

2 Project Outcomes & Findings

Since its inception, the Open Science Grid (OSG) has evolved into an internationally recognized key element of the U.S. national cyberinfrastructure enabling scientific discovery across a broad range of disciplines. This has been accomplished by a unique partnership that cuts across science disciplines, areas of technical expertise and U.S. research institutions. Building on novel software and shared hardware capabilities, the OSG has been expanding the reach of high throughput computing (HTC) to a growing number of science communities. This OSG project, "Sustaining and Extending the Open Science Grid: Science Innovation on a Petascale Nationwide Facility," was performed under grant numbers DOE DE-FC02-06ER41436 and NSF PHY-0621704.

The OSG now provides a comprehensive, dependable and cost-effective suite of distributed high throughput computing services, which underpins the distributed computing infrastructures of a wide spectrum of science projects. The largest OSG science stakeholder has been the Large Hadron Collider (LHC) program at the European Organization for Nuclear Research (CERN), comprising the U.S. ATLAS, U.S. CMS and ALICE-USA contributions to the World Wide LHC Computing Grid (WLCG). The global shared computing infrastructure enabled by the OSG services has facilitated a transformation in the delivery of results from the LHC and other advanced experimental facilities – enabling the public presentation of results days or weeks after the data is acquired, rather than the months or years it took before. The U.S. LHC scientific program embraces OSG as a major strategic partner in developing, deploying and operating their novel and cost effective DHTC infrastructure.

As a consortium of researchers, software developers and facility operators, the OSG contributes to the DOE's and the NSF's goals in serving science, research and workforce development through innovation in computing and software technologies. The OSG provides an intellectual hub for the entire DHTC community and drives the development of novel frameworks, educates and trains, forms new collaborative efforts, supports the development of new tools and serves as a testing and evaluation laboratory. We contribute to the NSF eXtreme Digital (XD) program as a Service Provider (SP) to the Extreme Science and Engineering Discovery Environment (XSEDE) project and to the DOE Scientific

Discovery through Advanced Computing (SciDAC) program as a promoter and adopter of advanced computational technologies and methods.

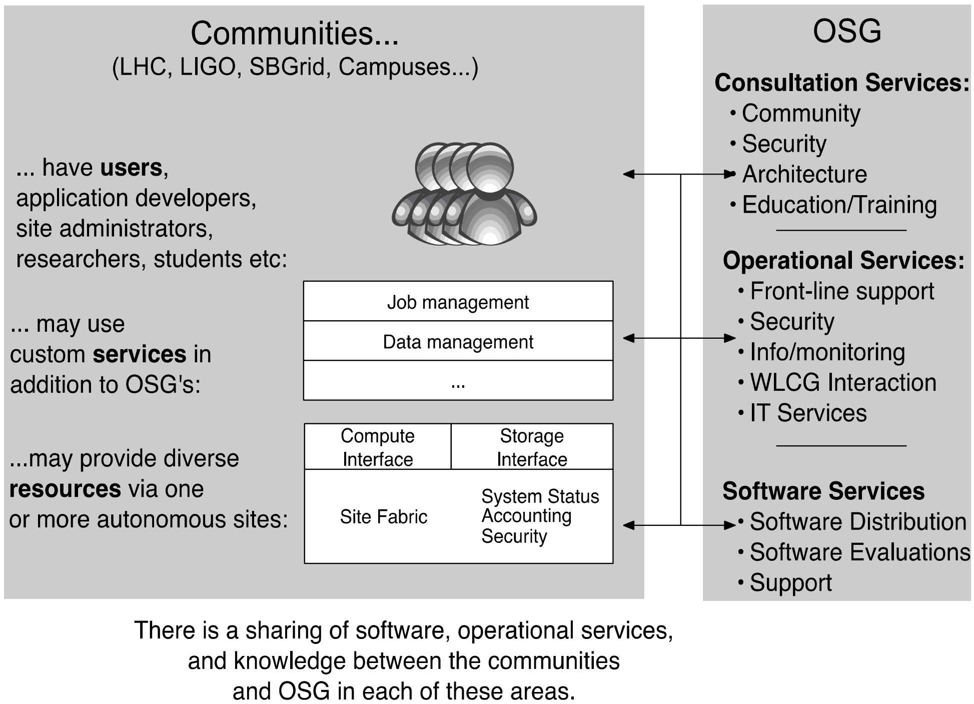

Figure 1: OSG's Fabric of Services & Community Focused Architecture

Today, the OSG fabric of services is composed of three groups—software services, support services like education, training, consulting in the best practices of DHTC, and an infrastructure of DHTC services (referred to as production services) for those who would like to join the OSG DHTC environment (Figure 1). Services in the first two groups serve the broader community that builds and operates their own DHTC environments (e.g. LIGO), as well as supporting the DHTC environment of the OSG.

The past focus in OSG toward meeting the strict demands of the LHC scientific community has led OSG to actively drive the frontiers of Distributed High Throughput Computing (DHTC) and massively Distributed Computing and to the development of a *production quality facility* that supports a broad variety of researchers; the OSG distributed facility, composed of laboratory, campus, and community resources, has demonstrated that it is meeting the current and future needs of scientific computing at all scales. It provides a broad range of common services and support, a software platform, and a set of operational principles that organize and support scientific users who access DHTC resources that are accessible via the mechanisms of Virtual Organizations (VOs) and Campus Grids (CGs).

The LHC relies extensively on the fabric of services provided by the OSG project. The U.S. ATLAS and U.S. CMS communities have invested heavily in this fabric of services through collaboration and direct contribution – human and computing resources, operational coherence, testing, integration and documentation. In a recent document the U.S. LHC management stated “It is vital to the LHC program that the present level of service continue uninterrupted for the foreseeable future, and that all of the services and support structures upon which the LHC program relies today have a clear transition or continuation strategy.[1]"

High Throughput Computing technology created and incorporated by the OSG and its contributing partners has now advanced to the point that scientific user communities (VOs) are simultaneously utilizing more geographically distributed HTC resources than ever before. Typical VOs now utilize ~20 resources with some routinely using as many as ~40 simultaneous resources. The overall usage of OSG has grown steadily during the last five years and now stands at ~55M hours per month as shown in Figure 2.

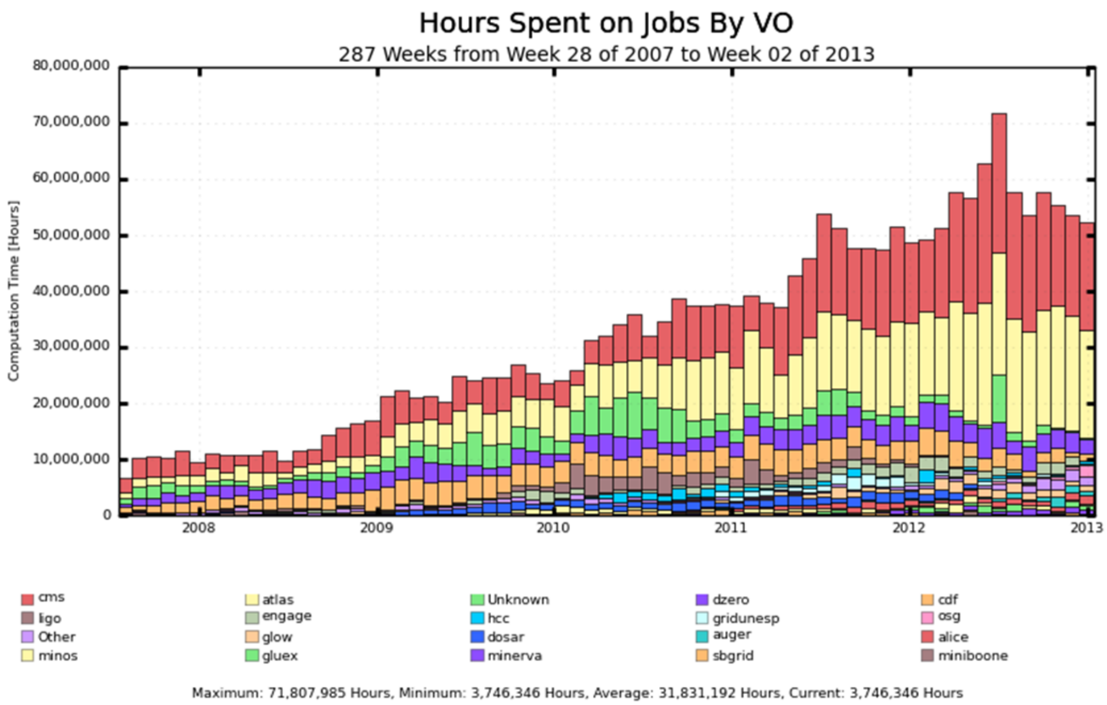

Figure 2: OSG Usage (hours/month) from July 2007 to January 2013

During stable, normal operations, OSG now provides ~1.8M CPU wall clock hours a day (~76,000 cores) with peaks occasionally exceeding 2.1M hours a day; approximately 25% of this capacity is available on a daily basis for resource sharing. Based on transfer accounting, we recorded ~320 Petabytes of data movement (both intra- and inter-site) during the last year; of this, we estimate 25% is GridFTP transfers between sites and the rest is via LAN protocols. In the last year, non-HEP researchers accessed ~49M hours on an opportunistic basis; this usage was often greater than 1 million hours per week even during peak usage volumes of LHC activity.

Key results from this project include: an effective DHTC infrastructure providing single sign-on for use of its services; round-the-clock, dependable services that facilitate effective and secure sharing of the resources; a high quality DHTC software stack used by LIGO, the WLCG, TeraGrid, and other smaller organizations and campuses; and a "home" to the DHTC community which pioneered the concept of federated national grids and distributed resource management overlays. Other accomplishments include

[1] http://osg-docdb.opensciencegrid.org/cgi-bin/ShowDocument?docid=983

joint activities with science groups, engagement with end-users and training of students through the residential summer school[2].

OSG continue to explore and innovate methods to make DHTC easier to use and more generally available to researchers in the USA. Some noteworthy facets of that work include:

1. Integration of cloud based resources,
2. Improved and easier to use Identity Management,
3. Better understanding of total available capacity in the OSG production fabric,
4. Enabling campus researchers to adapt their computing to run “anywhere” in the national CI

We endeavor to strengthen the services provided to the HEP community and thus enable their work at ever increasing scales of data transport and computation. In addition, we are working to better understand the computation needs and challenges of researchers on campus and develop paths that enable improved access to DHTC in the national CI.

3 Contributions to Science

In 2012, 474 scientific papers were published that depended on use of OSG services and software, many of which are LHC results and 20% of which are non-physics. The number of users of the OSG has risen substantially over the past five years, with more than 2000 end-users accessing the OSG computing resources. More than 160 students and 80 system administrators have attended technical training and education, and the number of university resources accessible through the OSG has risen from 40 to over 100. Prior annual reports have documented an ever increasing impact on science where the number of scientific publications using OSG provides one relevant metric: the June 2011 report lists 260 publications[3]; the June 2010 report lists 367 publications[4]; the June 2009 report list 111 publications[5]; and the June 2008 report lists 90 publications[6].

ATLAS and CMS depend on the OSG to deliver the operational services, and an intellectual community to sustain and enhance their national CI, and to guarantee its seamless integration into the evolving global computing landscape. Much of what is presently taken for granted by the LHC community when analyzing petascale datasets, in particular the seamless use of the huge OSG DHTC resources by hundreds of individual researchers and groups for data analysis, was bleeding edge technology research only a few years ago.

LIGO benefited from OSG software infrastructure to operate the LIGO Data Grid and from OSG services to share LIGO Data Grid computing resources with other communities. LIGO also benefited from opportunistic computing, with the OSG project providing the “embedded user support” effort to port LIGO applications to a DHTC environment.

The Tevatron Run II experiments have increasingly adapted their legacy infrastructure to integrate with and depend on the services provided by the OSG. Over the next several years these experiments will

[2] http://osg-docdb.opensciencegrid.org/cgi-bin/ShowDocument?docid=813
[3] http://osg-docdb.opensciencegrid.org/cgi-bin/ShowDocument?docid=1054
[4] http://osg-docdb.opensciencegrid.org/cgi-bin/ShowDocument?docid=965
[5] http://osg-docdb.opensciencegrid.org/cgi-bin/ShowDocument?docid=860
[6] http://osg-docdb.opensciencegrid.org/cgi-bin/ShowDocument?docid=770

continue analysis of the data acquired before the Tevatron shutdown and provide the LHC community valuable physics validation and in-sights. They are joined by several (intensity and cosmic frontier) particle physics experiments that have adopted OSG services.

Apart from the core constituencies, a number of additional science communities from physics, biology, chemistry, mathematics, medicine, computer science, and engineering have benefitted from the DHTC services and software provided by the OSG. The structural biology community (SBGrid) centered at Harvard Medical School actively leverages OSG; this is a collaboration of more than 140 X-ray crystallography, NMR and electron microscopy laboratories, including groups at more than 50 academic institutions in 12 countries. Several communities — including DES, GlueX, IceCube, SCEC, NEES and LSST — have started their large scale distributed computing with OSG in mind or are exploring OSG as part of their preparations for the future.

Appendix 1

Institutional Contributions

Institution	**Functional Area**
Boston University	ATLAS Extensions
Brookhaven National Laboratory	ATLAS Extensions, Technology Investigation , Software
California Institute of Technology	Co-PI, LIGO Extensions, User Support
Clemson University	New User Engagement
Columbia University	Outreach (South Africa)
Cornell University	Software
Fermi National Accelerator Laboratory	Co-PI, Executive Director, Security, Software, User Support, Project Management
Harvard Medical School	SBGrid Community Support
Information Sciences Institute (USC)	User Support
Indiana University	Operations, Communication
Lawrence Berkeley National Laboratory	Security, Software
Renaissance Computing Institute (UNC)	New User Engagement
SLAC National Laboratory	Security
University of California at San Diego	CMS Extensions, Scaling Analysis
University of Chicago	Integration Testing, Education, Software
University of Florida	Co-PI, Education
University of Illinois	Security
University of Iowa	User Support
University of Nebraska – Lincoln	Metrics, Technology Investigation, User Support
University of Wisconsin	PI, Software