## Open Science Grid Year 1 and 2 Summary

Open Science Grid (OSG) is a large-scale collaboration that is advancing scientific knowledge through high-performance computing and data analysis, by operating and evolving a cross-domain, nationally distributed cyber-infrastructure. OSG provides a collaborative distributed high-throughput computing (DHTC) eco-system to U.S. researchers, consisting of a consortium of contributing communities (science users, resource providers, software providers), the OSG project which is jointly funded by the DOE and NSF, and satellite projects. This summary provides a brief status report for the first two years of the OSG renewal project, based on the February 2013 annual report[1] and the March 2014 annual report[2].

OSG provides a production-quality facility, connecting a broad variety of researchers to distributed computing centers at university campuses, national laboratories, and other community resources. The OSG environment is based on a range of common services, user and site support activities, maintenance and distribution of a broad software stack, and operational principles. This way it helps producing scientific knowledge through distributed high-throughput computing.

The goal of OSG is to enable science. A useful metric of its impact is quantity of publications enabled by OSG resources. A list of science publication enabled in 2013 via leverage of OSG is available in the annual report. In summary, OSG was credited having enabled 84 ATLAS, 97 CMS, 97 CDF, 26 Dzero publications, and many more from a broad range of science fields like HEP experiments and theory, other physics areas including nuclear physics, computational biology, medical research, computer science and others, for a total of 448 publications credited in 2013. The summary count of publications in the first two years of the project is in Appendix 1.

OSG serves the needs of the Large Hadron Collider (LHC) experiments by providing a platform for their contribution and consumption of distributed computing resources. The LHC experiments continue to rely on OSG for their computing needs in the U.S., making up more than 2/3 of the total usage in OSG, or about 448M computing hours out of a total of 681M hours provided in OSG during 2013. The LHC experiments not only rely on OSG for computing resources, but also for operational, consulting, and software services. In addition, OSG is the conduit in the U.S. for collaboration and joint initiatives among the LHC experiments and the Worldwide LHC Computing Grid (WLCG) project.

In addition to the U.S. LHC experiments, the next biggest use of OSG is access to opportunistic computing resources by a broad and diverse range of U.S. based researchers. The OSG project focused on growing and broadening its eco-system, opening both to new organizations like Intensity Frontier experiments, and to individual researchers at campuses. In the last year, non- HEP researchers accessed ~87M CPU hours of opportunistic computing, a 60% increase with respect to the prior year. The opportunistic usage included medium to large data processing activities with increased complexity, including Intensity Frontier experiment workflows. OSG provided the technology and operated the infrastructure to support this, through newly developed mechanisms of the *OSG Direct* system used by the catch-all OSG VO, and the recently developed *OSG Connect* mechanism for OSG Campus Grids.

---

[1] http://osg-docdb.opensciencegrid.org/cgi-bin/ShowDocument?docid=1146

[2] http://osg-docdb.opensciencegrid.org/cgi-bin/ShowDocument?docid=1184

An example how OSG opportunistic resources deliver to the HEP community at large is the Snowmass Summer Study, which relied on detailed simulations of a "generic" detector for hadron colliders. The data was produced and used on OSG by a diverse group of theorists and experimentalists across the field, to predict the physics reach of a physics program at high-luminosity and high-energy colliders. The knowledge gained provided vital input to the P5 report. To simulate these data samples, OSG provided more than 11M CPU hours, opportunistically, based on a broad effort with contributions from 32 different computing centers. Half of the cycles were delivered by HEP lab computing, the other half by universities. Less than half of the resources came from the U.S. LHC Tier-1 centers at BNL and FNAL, the majority of CPU hours was obtained in about equal parts from U.S. LHC Tier-2 centers and a large number of other non-LHC computing centers. Almost 2 million CPU hours came from computers that were non-U.S.HEP, including a small contribution from international partners.

Examples of OSG use by experiments from outside the HEP Energy Frontier are the NOvA experiment and LBNE who go newly connected to OSG with help from our user support group. NOvA used both OSG opportunistic resources and dedicated NOvA resources at Harvard and SMU that were newly connected to OSG, to run event simulations. LBNE successfully ran a simulation workflow demonstrator for its DOE computing review to get first experience with its newly developed computing model, in which OSG is expected to play an important role.

This demonstrates the value of OSG for HEP, as a virtual facility enabling new and agile science users to get large and valuable additional benefits from already existing computing facility infrastructures, across DOE and NSF offices. This approach also enables tapping into additional resources from non-HEP sources. Large amounts of computing resources are being made available through OSG's integration and sharing of lab and university computing infrastructure. These opportunistic computing cycles that otherwise would go wasted make it possible for researchers to accomplish their goals in a compressed time, addressing peak-demands e.g. before conferences or during processing campaigns.

The OSG operations team continued to provide highly reliable common services, effective ticket handling, and support mechanisms that are vital to a production grade computing fabric. We continued our partnership with Worldwide LHC Computing Grid and European Grid Initiative operations groups on behalf of the USLHC experiments. This included working on change management for our shared services (Ticketing, Monitoring, and Accounting) and attendance and collaboration at major peering partner events.

Another important facet of OSG support for stakeholders is the software stack that provides the platform for DHTC. The OSG software packaging has completed the transition to native packaging; all software integrated and distributed by OSG for stakeholder installation is now available as RPMs. More than 80% of OSG sites have successfully transitioned to this new software base. Driven by community concerns about the strength of a widely used X.509 cryptography algorithm (SHA-1), the broader PKI community planned a move to a newer, more robust algorithm (SHA-2). Not all software that uses X.509 certificates was known to handle SHA-2 signed certificates properly; thus, the Software team evaluated and tested all OSG provided software that depends on PKI for SHA-2 compliance and, for the few components that were not ready, develop patches to bring the software into compliance. The project was completed well ahead of the August 2013 target for initial deployment. Another significant software update was the transition of all OSG software components that use the Java platform from Java version 6 to 7. Support for Java 6 implementations ceased early in 2013, and due to the frequency and urgency of Java platform security and bug updates, it was critical for the well-being of the OSG infrastructure to move to a supported Java implementation. The Software team updated Java components as needed, helped test all Java software in OSG, planned the deployment process for the transition and wrote a monitoring tool for Java installations. And this included contribution of a great deal of effort needed for updating the BeStMan and jGlobus software components to work in Java 7.

To further improve quality, we formed an independent release team in summer 2013 to further assure the quality of the OSG software by testing it in typical use cases as well as testing that the changes introduced by the Software team have their desired effect. This Release team has established a fixed release schedule which greatly benefits the sites and OSG staff.

The OSG project leverages its geographic and institutional diversity to bring DHTC computing to a broad and growing cross-section of US campus researchers, both as individuals and groups. OSG assists researchers and groups affiliated with OSG Campus Grids. These users typically do not bring resources to the OSG pool. They use both local and remote resources and are motivated to accomplish their computing efficiently and timely. The major OSG Campus Grids include GLOW (University of Wisconsin-Madison), the HCC (University of Nebraska), UC3 (University of Chicago), and UCSD (University of California at San Diego). Collectively, in 2013 they provided ~30M hours of computational access to their local researchers.

The OSG continued to expand access to OSG DHTC computing resources for U.S. researchers who are not already affiliated with an existing OSG community; this is being addressed on three distinct fronts:

1. <u>OSG-XD:</u> The OSG continues to be a service provider to XD/XSEDE. We provide the first line of support for these users and provide tutorials on structuring work for high-throughput computing (HTC) environments. The OSG-XD system can access the >100 sites in OSG by using the glideinWMS job management system; we provide a virtualized interface for XSEDE users who can view the OSG as one pool of compute resources. We continue to collaborate with XSEDE staff in refining the system and providing documentation and support to users.
2. <u>OSG-Direct:</u> We continue to provide access to DHTC resources to researchers who contact OSG directly. These researchers learn about us via referrals from other OSG users, and they usually connect with us via a web-form on our website or via email. They are provided access to DHTC via login nodes established by OSG.
3. <u>OSG-Flocking:</u> We have become aware of campus communities who need access to more HTC computing resources than are locally available but for a variety of reasons (e.g., effort, expertize, focus, etc.) are not able to form and operate their own VO. Thus, we now provide a "bulk" connection method for these communities; they can set up a local submit host and overflow their jobs to the OSG via such a connection.

OSG serves as a Level 2 service provider to the NSF XD project using opportunistic cycles. Working with the XSEDE project and participating in the XD resource allocations process (XRAC), in 2013 OSG enabled 23 PIs to use ~26M hours of computing.

In addition, several research groups and PIs directly approach OSG for access to computing resources. Through this *OSG Direct* channel, in 2013 OSG provided ~21M hours to 14 user groups.

All these activities for the LHC and for opportunistic users have benefited from improved capabilities in the OSG services and processes. The OSG project has worked at reducing obstacles to access and contribute resources. Some examples:

OSG is now a certificate and credential provider, after transitioning and re-engineering this service previously provided by the ESnet DOEGrids Certificate Authority (CA). Since February 2013, we have seen the effective operation of the OSG PKI, successfully taking over service from the DOEGrids PKI, which ceased issuing new certificates on March 23, 2013. The new OSG Public Key Infrastructure (PKI) CA is using a commercial vendor, DigiCert Inc., as the underlying certificate provider. OSG provides seamless support of strong authentication for OSG users and services and for the DOE scientific community at large.

In parallel, the project has made progress in further developing its trust relationship model, to loosen the dependency on end-user PKI certificates. In this new model, instead of burdening individual researchers with an additional authentication layer, OSG achieves cyber security through a certification and auditing process based on the ability of the OSG job submission services to trace the identity of the originating science user at a campus or lab, requiring no additional individual grid certificates which scientists consider a major obstacle. Another important alternative in the identity management area was to enhance our ability to support federated identities from university campuses and DOE laboratories by enabling sites to accept CILogon certificates for user access. We have worked with sites to understand their risk model regarding this CA and created a better understand of the pros and cons of accepting this CA; thus we have transition six sites that accept certificates from this CA and these sites represent 40% of the total production in OSG.

OSG is in the process of deploying the HTCondor-CE as the next generation of Compute Element gateway technology, to improve reliability, scalability, and fault propagation. HTCondor-CE shares underlying technology with both BOSCO (developed in collaboration between HTCondor and OSG), and the CREAM CE (developed by the European Middleware Initiative, EMI) and in the longer-term reduces the OSG's unique dependence on GRAM.

The OSG Connect service has made it easier for campuses to connect local researchers with the full array of high-throughput computing resources in OSG. The service enables campus researcher to access OSG using their local campus credentials. This opens OSG access to large campus communities, with examples at Duke University where nuclear physics and HEP groups make use of these mechanisms, in addition to non-physics applications.

The OSG Application Software Installation Service (OASIS) was deployed at the OSG sites as a mechanism to transparently deploy VO software distributions across OSG resources. This service is based on CVMFS, the CERN Virtual Machine File System software, and accelerates the use of distributed computing resources in particular for small and medium size VOs.

In a distributed system where jobs can "land" on any of the ~100 sites, we face the challenge of making large data sets accessible "anywhere" for those analysis jobs. We are trialing a solution call OSG public storage to enable this data delivery for some use cases. This capability (based on the integrated rule-oriented data-management software, iRODS) allows easier movement of large data sets to and from OSG sites, and this has included the needed integration with OSG storage elements. This solution is currently in limited deployment for OSG-Direct users.

OSG is also working on integrating tools from the community to increase network awareness. We are working on a process to collect and store metrics associated with network connectivity in OSG. Using network management probes at OSG sites through the PerfSONAR-PS, we are prototyping a dashboard with the goal of gathering, storing, and publishing current and historical network metrics for alarming, troubleshooting, and higher-level service integration.

In summer 2013, the OSG ran another of the successful week-long student schools. More than 22 attendees learned about the principles of distributed high-throughput computing and how to build and execute their applications on local and remote resources. This user school was conducted again in

July 2014.

The OSG Consortium continued its partnership with four satellite projects that are delivering benefit to the OSG community's future needs, including in networking and middleware.

In summary, the OSG project has delivered on continuing excellent support for the distributed high-throughput computing environment that the LHC and other experiments can rely on, and has increased its reach to start including Intensity Frontier experiments and a broad spectrum of science users at labs and university campuses.

## Appendix 1

### Publications enabled by use of OSG computing in year1 and year2

| **Research Community** | **Number of Publications** |
|---|---|
| ATLAS | 215 |
| CDF | 163 |
| CMS | 220 |
| Dzero | 66 |
| GLOW | 117 |
| GridUNESP | 35 |
| HCC | 6 |
| Minos | 3 |
| NEES | 1 |
| RENCI/Engage | 3 |
| SBGrid | 4 |
| STAR | 15 |
| UC3 | 9 |
| UCSDGrid | 11 |
| University of Notre Dame | 1 |
| Various PIs via OSG-Direct | 46 |
| | |
| **Total** | **915** |

## Appendix 2

### Status of Milestones in OSG Proposal

The OSG Program of work is updated annually based on the needs of the stakeholders and strategic goals that are informed by inputs from the OSG Council and US LHC community. Shown below is the status of specific milestones that were presented in the proposal submitted in early 2012.

**Annual Goals in Sustaining the OSG:**

a) Increase in CPU usage that is more than Moore's Law showing growth to meet the scaling needs of the users.

*OSG CPU usage has grown as driven by the work programs of stakeholders; during the last year, the usage of OSG resources by VOs has averaged ~14M hours per week. Peak weeks have grown from about 14M hours to recent usage of 18M hours in a single week.*

b) Full compliance with the operations service SLAs.

*OSG operational services met or exceeded expectations for service availability and reliability as defined by Service Level Agreements.*

c) Doubling of accounted data movement to meet growth in data intensive science.

*OSG's accounted data movement is dominated by large-scale data intensive activities at the USLHC Tier-1 and Tier-2 sites. We observed peaks through early 2013 as LHC Run-I work finished, but data transfers are currently below normal. We expect to see a doubling during LHC Run-2.*

d) >=2 additional communities using DHTC services in production; >2 tutorials and documentation for new capabilities.

*A broad set of new communities have been assisted in achieving production on OSG as outlined in section 2.2.5 of the 2014 OSG annual report. In addition, various teams have connected their resources to OSG as new sites; these include: 1) PNNL (for the Belle & Belle-2 experiment), 2) Oregon State (NEES), 3) Ohio Supercomputing Center (Belle-2); and 4) UMD-IGS. In addition, we continue to provide tutorial and outreach to communities interested in leveraging OSG; this included a tutorial on use of HTC resources at the XSEDE-12 conference.*

**Transform computing on campuses through new DHTC technologies:**

*Year 1:*

*a)* Deploy technology to account usage of users, jobs and data movement to campuses.

*Delayed to year2. In year2, a Field-of-Science project accounting system was designed and deployed in partnership with the Gratia project. This system enables campus PIs and users to register their projects in OIM and query for usage reports showing usage across OSG.*

b) Release campus infrastructure software distribution Production Version 1.

*Bosco is an integrated software package that provides an easy way for a user to submit HTC jobs to clusters available on the campus or directly from their laptop or desktop computers. The Bosco team released several early releases of Bosco for user testing that include a full set of documentation and the creation of the Bosco website for users. Experience with initial users indicated that this method was too complex and led to a fundamental change in direction which was implemented in year2.*

*Year 2:*

a) Assess metrics that encapsulate a measure of adoption and increase in usage.

*Due to the re-plan associated with Bosco software development and deployment; the decision was made to move to providing a service in addition to products. We have diversified our offerings to campuses by offering OSG-Direct (requires software installation), OSG-Connect (service we operate), and OSG-XD (service we operate). We track usage across all three independently. See answer to b) below and presentations at review for more details.*

b) 20% increase in each of new, and usage by existing, users of campus technologies.

*With the adoption of OSG-Direct and OSG-Connect as services for campus users, we have seen growth far exceeding 20% in Campus researchers using OSG-DHTC. In Year2, the total usage via OSG-XD, OSG-Direct, and OSG-Connect was >50M wall hours.*

c) 2 production versions of campus software to extend the capabilities and scalability.

*Cancelled; replaced by OSG-Connect as a service. This service is deployed to Duke and work is in-progress to deploy this at U-Chicago and U-Michigan. The same service is being evaluated by both US-ATLAS and US-CMS to better integrate their Tier-3s.*

*Year 3:*

a) 20% increase in new, and usage by existing, users of campus technologies.

*In-progress; early current usage trends suggest a growth >30% usage as compared to year2.*

b) Transition of appropriate campus support services to sustained OSG operations.

*In-progress; the Campus Grids team is working with operations to transition certain components of the OSG-Connect service to the OSG operations team.*

**Transformation of our core communities computing capabilities to exascale science:**

*Year 1:*

a) 10% of resources support end-to-end capability to be schedulable as HTPC and simultaneously usable and available by HTC/single processor job.

*Delayed; the testing of HTPC was completed but the demand for multi-core jobs did not materialize as expected. At present about 6 sites are configured for HTPC.*

*Year 2:*

a) 20% of resources support end-to-end capability to be schedulable as HTPC and simultaneously usable and available by HTC/single processor job.

*Delayed; the testing of HTPC was completed but the demand for multi-core jobs did not materialize as expected. At present about 6 sites are configured for HTPC.*

b) 90% of LHC workload and 10% of non-LHC workload supports remote I/O capabilities.

*Nearly 100% of the CMS workload supports remote I/O capabilities. However, it is worth noting that remote I/O was never meant to become the dominant access mechanism. It has always been intended to be used as needed to improve efficiencies and provide failover. The actual scale of use today is 10-20%, which is considered adequate. Several of the OSG and GLOW VO users utilize the network of HTTP caches put in place at OSG sites to distribute software and data to their jobs (either through directly downloading the files or indirectly through CVMFS). We do not have a precise accounting of how many workloads do this, but we encourage all new workflows to utilize this service.*

c) Data movement across the DHTC fabric of >750 Petabytes/year.

*The data movement for Year 2 totaled 318PB. This is dominated by activity at the USLHC Tier-1 and Tier-2 sites; we believe this will increase significantly at the start of Run 2.*

*Year 3:*

a) >50% of resources support end-to-end capability to be schedulable as HTPC and simultaneously usable and available by HTC/single processor job.

*Delayed; the testing of HTPC was completed but the demand for multi-core jobs did not materialize as expected. At present about 6 sites are configured for HTPC.*

b) 50% of non-LHC workload supports remote I/O capabilities.

*In-progress, but not expected to achieve 50%. OSG VO workloads were supported by OSG Public storage (based on iRODS) which is deployed for use by researchers in vo=osg. HTTP caching is being deployed by Intensity Frontier experiments and new Campus researchers for read-only access to large datasets. Most of the non-LHC science drivers continue to have modest data needs that do not benefit dramatically from the remote I/O capabilities. The 50% goal is thus unlikely to be a relevant goal for the forseeable future.*

c) Data movement across the DHTC fabric of >750 Petabytes/year

*The data movement for (incomplete) Year 3 is currently down from Year 2. This number is dominated by activity at the USLHC Tier-1 and Tier-2 sites; we believe this will increase significantly at the start of Run 2.*

**Access to an expanded set of job and data services accessible via a single identity:**

*Year 1:*

a) 10% of non-LHC Users accessing OSG services using campus identities.

*Delayed; staff was redirected to implement the OSG PKI service due to the planned shut-down of the DOEGrids CA.*

b) Complete the architecture and design of the new set of ID management services.

*Delayed; staff was redirected to implement the OSG PKI service due to planned shut-down of the DOEGrids CA.*

*Year 2:*

a) 20% of non-LHC Users accessing OSG services using campus identities.

*With the deployment of OSG-Connect in August 2013, we have a good mechanism for using Campus identities to access OSG. At present the usage is still small but we expect this will grow as we identify more Campus researchers who can benefit from using DHTC for their science. With the traceability project, users can access the OSG Production Grid using only the identity recognized by the VO – often the campus identity.*

b) Deliver initial release of new set of ID management services.

*Work was redirected to remove the need for end-user certificates for VOs using glideinWMS to send jobs to sites. This project was an integral part of our long-term strategy of enabling secure access by middleware such as pilot job factories, and allowing them to authenticate users via means a community may have in place other than user certificates. To achieve this, we evaluated the user traceability and we found out that the glideinWMS system had sufficient capability to trace users.*

*Year 3:*

a) >30% of non-LHC Users accessing OSG services using campus identities.

*In-progress. All new campus researchers who access DHTC via OSG-Connect use their local campus identities.*

c) Transition new set of ID management services to production and operations.

*No longer relevant, as new IDM services not needed due to the success with Campus Grids using campus identities and other VOs using the traceability project to run at sites without the need for end-user certificates.*

**Improve the usability, expand the usage, lower barriers of adoption:**

*Year 1:*

a) Deliver report on integration of virtualized resources into the OSG fabric of services.

*Cancelled; staff was redirected to OASIS and HTCondor CE*

b) Prototype integration of one cloud resource into the production DHTC environment.

*OSG Connect allows access to cloud resources, including amazon, through the HTCondor and GlideinWMS mechanisms. Also, the LHC VOs are now using cloud resources routinely, including the cloud-provisioned resources at CERN, that are accessible through the job execution overlay infrastructure that OSG supports.*

c) 100% VDT packages available as RPMs.

*As of September 2012, the new OSG Software 3 distribution was complete and ready for* *deployment. All needed software packages were converted to RPMs, Enterprise Linux versions 5 and 6 (32- and 64-bit) were fully supported, and many packages were dropped from the OSG distribution and replaced by corresponding ones already in standard distributions.*

d) Production release of configuration management of RPM-packaged VDT software.

*In December 2012, the OSG Production team advised sites to upgrade to the new RPM software distribution and, in February 2013, over 60% of sites had upgraded at least their central service node.*

e) Deliver report on extending the job-level monitoring.

*Delayed; staff was redirected to support the SHA-2 transition and upgrade to Java. In particular, a great deal of effort went into updating the BeStMan and jGlobus software components to work in Java 7.*

*Year 2:*

a) Improve reliability of software distribution via duplication of software repositories.

*Completed as part of the 2012 switch to RPMs; OSG sites utilize mirrors in Indiana, Wisconsin, and Nebraska.*

b) Provide hosting of non-VDT software on behalf of OSG communities.

*Completed as part of the 2012 switch to RPMs; we have a new "contrib" repository where we can distribute contributed software on behalf of our communities without adding it to the OSG Software Stack. During 2013, this effort was extended to better allow OSG communities to utilize our build infrastructure.*

c) Deliver report on integration of advanced (100G, Terabit) networks.

*Delayed in order to put more effort into the network monitoring activity.*

d) Extend existing Dynamic Resource Allocation Services to include centralized policies that regulate allocation.

*Delayed. We have instead managed the available opportunistic cycles through fair-share policies within OSG VO at the Open Facility "on-ramp".*

e) Collaborate with U.S. LHC for initial deployment of simplified data services for non-

LHC.

*Delayed but currently in-progress; HTTP caching is being introduced to other experiments with the initial introduction to the Intensity Frontier experiments*

f) Integration of one cloud resource into the production DHTC environment.

*Cancelled as a stand-alone goal, but technology available through OSG Connect and the LHC production systems. For non-LHC users we have not yet identified cloud resources that are appropriate for OSG-style opportunistic use.*

g) Provide enhanced OSG-wide job monitoring in prototype.

*Delayed; staff was redirected to support SHA-2 transition and upgrade to Java 7. In particular, a great deal of effort went into updating the BeStMan and jGlobus software components to work in Java 7.*

*Year 3:*

a) Develop a trust flow diagram of VDT stack. Identify the effect of configuration parameters on the security of the software.

*Delayed; staff was redirected to the traceability project which provided ability for VOs to run at sites without end-user certificates.*

b) Provide OSG services to extend the dynamic resource allocation capability.

*Completed; using BOSCO, the campus grids activity is able to allocate resources across campus resources and HPC systems outside the OSG Production Grid. We will continue to extend these capabilities.*

c) All software available as source RPMs. Drop support for Pacman.

*Completed; the End of Life (EOL) date for Pacman was May 31, 2013.*

d) 20% of resources support simplified data management services for non-LHC VOs.

*In progress; all sites are now asked to support Squid caching to support HTTP data access.*

e) Full integration of multiple cloud resources into the production DHTC environment.

*In progress: Specific stakeholders (in particular, USATLAS and USCMS) have integrated cloud resources they own or purchase into their production DHTC environment. However,* *we have not had any of these resources register into the OSG Production Grid nor have we identified cloud resources that allow opportunistic usage.*

f) Transition enhanced OSG-wide job monitoring into operations.

*Delayed; staff was redirected to support SHA-2 transition and upgrade to Java 7 -- In particular, a great deal of effort went into updating the BeStMan and jGlobus software components to work in Java 7.*