

Open Science Grid Annual Report 2008–2009

The Open Science Grid Consortium

NSF Grant 0621704



Open Science Grid

Miron Livny	University of Wisconsin	Principal Investigator
Ruth Pordes	Fermilab	Co-PI, Executive Director
Kent Blackburn	Caltech	Co-PI, Council co-Chair
Paul Avery	University of Florida	Co-PI, Council co-Chair

Table of Contents

1. Introduction to Open Science Grid	3
1.1. Virtual Organizations	3
1.2. Software Platform.....	3
1.3. Common Services and Support	4
1.4. OSG Today (June 2009).....	4
2. Participants:	5
2.1. People	5
2.2. Partner Organizations	8
2.3. Participants: Other Collaborators	9
3. Activities and Findings:.....	9
3.1. Research and Education Activities.....	9
3.2. Findings	11
3.3. Training and Development	39
3.4. Outreach Activities.....	41
4. Publications and Products.....	42
4.1. Journal publications.....	42
4.2. Book(s) and/or other one time publication	42
4.3. Other specific products.....	42
4.4. Internet dissemination	42
5. Contributions.....	43
5.1. Contributions within Discipline	43
5.2. Contributions to Other Disciplines.....	43
5.3. Contributions to Education and Human Resources.....	44
5.4. Contribution to Resources for Science and Technology	44
5.5. Contributions Beyond Science and Engineering	45
6. Special Requirements	45
6.1. Objectives and Scope	45
6.2. Special Reporting Requirements	46

Notes on Fastlane instructions – these are all in Italics:

Graphics, Equations, Fonts

Unfortunately, current Web technology does not allow for the text formatting (bold, italics, fonts, superscripts, subscripts, etc.) nor for graphics, special equation formats, or the like. If pasted in from other software applications, they will be lost in transfer to our database. We hope that the technology will soon catch up in this respect. In the meantime our system does allow you to attach one PDF file with graphics, equations or both (no text please, other than labels or legends why this restriction?). You may refer to the graphics or equations in that file from any text entry in this system.

1. Introduction to Open Science Grid

The Open Science Grid (OSG) enables collaborative science by providing a national cyber-infrastructure of distributed computing and storage resources. The goal of the OSG is to transform processing and data intensive science through a cross-domain, self-managed, nationally distributed cyber-infrastructure that brings together campus and community resources. This system is designed to meet the needs of Virtual Organizations (VOs) of scientists at all scales. OSG is jointly funded by the Department of Energy and the National Science Foundation to build, operate, maintain, and evolve a facility that will meet the current and future needs of large scale scientific computing. To meet these goals, OSG provides common services and support, a software platform, and a set of operational principles that organizes users and resources into Virtual Organizations.

1.1. Virtual Organizations

Virtual Organizations (VOs) are at the heart of OSG principles and its model for operation. VOs are a collection of researchers who join together to accomplish their goals; typically they share the same mission, but that is not a requirement for establishing an OSG VO. A VO joins OSG to share their resources, computing and storage with the other OSG VOs and to be able to access the resources provided by other OSG VOs as well as share data and resources with international computer grids (i.e. EGEE). The resources owned by a VO are often geographically distributed; a set of co-located resources is referred to as a site and thus a VO may own a number of sites. Thus there are two key aspects of VOs: 1) the user community within a VOs that submits jobs into the OSG; and 2) the set of computing and storage resources that are owned by a VO and connected to the OSG. In some cases, VOs do not bring resources to OSG and are only users of available resources on OSG.

A key principle in OSG is the autonomy of VOs that allows them to develop an operational model that best meets their science needs; this autonomy applies both to their user community and sites. OSG requires each VO to establish certain roles (i.e. VO manager, VO admin, VO Security Contact) and agree to a set of policies (e.g. Acceptable User Policy) which allow operation of the OSG as a secure and efficient grid. VOs administer, manage, and support their own user community. In addition, many VOs provide common software infrastructure designed to meet the specific needs of their users. VOs as providers of resources also have great autonomy in building and operating their sites. Sites use the OSG software stack to provide the “middleware layers” that make their sites ready for connection to the OSG. Sites set policies on how their resources will be used by their own users and other VOs; the only requirement is that sites support at least one other VO but the site controls the conditions under which that resource is available. However, OSG does not tightly restrict what hardware or operating system software a VO may supply or what software it may use to access OSG or provide resources on OSG: they are autonomous and are allowed to make such choices as long as they meet the basic requirements. This autonomy allows a VO to build its computing resource to meet its specific needs and makes it more likely that a VO will choose to join OSG because it doesn’t have to compromise its own needs to do so.

1.2. Software Platform

The primary goal of the OSG software effort is to build, integrate, test, distribute, and support a set of common software for OSG administrators and users. OSG strives to provide a software

stack that is easy to install and configure even though it depends on a large variety of complex software.

The key to making the OSG infrastructure work is a common package of software provided and supported by OSG called the OSG Virtual Data Toolkit (VDT). The VDT includes Condor and Globus technologies with additional modules for security, storage and data management, workflow and other higher level services, as well administrative software for testing, accounting and monitoring. The needs of the domain and computer scientists, together with the needs of the administrators of the resources, services and VOs, drive the contents and schedule of releases of the VDT. The OSG middleware allows the VOs to build an operational environment that is customized to their needs.

The OSG supports a heterogeneous set of operating systems and versions and provides software that publishes what is available on each resource. This allows the users and/or applications to dispatch work to those resources that are able to execute it. Also, through installation of the VDT, users and administrators operate in a well-defined environment and set of available services.

1.3. Common Services and Support

To enable the work of the VOs, the OSG provides direct staff support and operates a set of services. These functions are available to all VOs in OSG and provide a foundation for the specific environments built, operated, and supported by each VO; these include:

- Information, accounting, and monitoring services that are required by the VOs; and forwarding of this information to external stakeholders on behalf of certain VOs,
- Reliability and availability monitoring used by the experiments to determine the availability of sites and to monitor overall quality,
- Security monitoring, incident response, notification and mitigation,
- Operational support including centralized ticket handling,
- Collaboration with network projects (e.g. ESNet, Internet2 and NLR) for the integration and monitoring of the underlying network fabric which is essential to the movement of petascale data,
- Site coordination and technical support for VOs to assure effective utilization of grid connected resources,
- End-to-end support for simulation, production, analysis and focused data challenges to enable the science communities accomplish their goals.

These centralized functions build centers of excellence that provide expert support for the VOs while leveraging the cost efficiencies of shared common functions.

1.4. OSG Today (June 2009)

OSG provides an infrastructure that supports a broad scope of scientific research activities, including the major physics collaborations, nanoscience, biological sciences, applied mathematics, engineering, and computer science. OSG does not own any computing or storage resources, but instead they are all contributed by the members of the OSG Consortium and are

used both by the owning VO and other VOs; recent trends show that about 20-30% of the resources are used on an opportunistic basis by VOs that do not own them.

With about 80 sites (see Figure 1) and 30 VOs, the usage of OSG continues to grow; the usage varies depending on the needs of the stakeholders. During stable normal operations, OSG provides approximately 600,000 CPU wall clock hours a day with peaks occasionally exceeding 900,000 CPU wall clock hours a day; approximately 100,000 to 200,000 opportunistic wall clock hours are available on a daily basis for resource sharing.

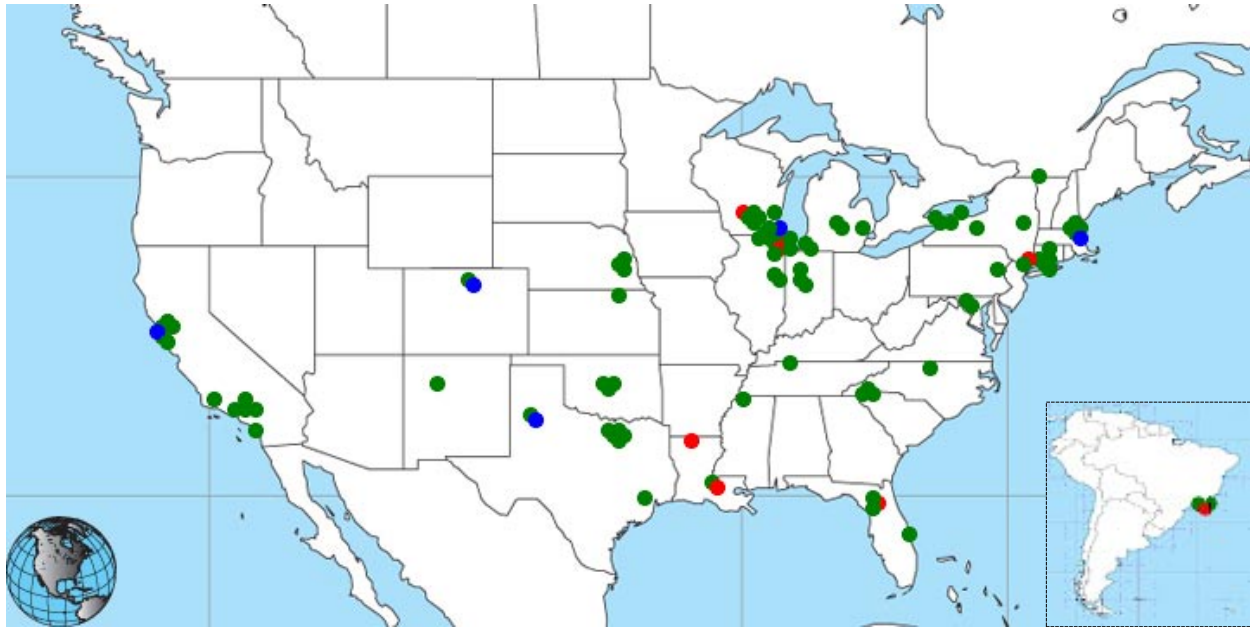


Figure 1: Sites in the OSG Facility

2. Participants:

2.1. People

What people have worked on the project (please note inside the project, a distinction should be made between paid and unpaid effort).

Name	Description	Paid?	160 Hours	Institution
OSG PIs				
Paul Avery	Co-PI & Council Co-Chair	No	Yes	UFlorida
Kent Blackburn	Co-PI & Council Co-Chair	Yes	Yes	Caltech
Miron Livny	Co-PI & Facility Coordinator	Yes	Yes	UWisconsin
Ruth Pordes	Co-PI & Executive Director	Yes	Yes	Fermilab
PIs and Area Coordinators				
Mine Altunay	Security Officer	Yes	Yes	Fermilab
Alina Bejan	Education Co-Coordinator	Yes	Yes	UChicago
Alan Blatecky	Co-PI	No	No	RENCI

Brian Bockelman	Metrics Coordinator	Yes	Yes	UNebraska
Eric Boyd	PI	No	No	Internet2
Rich Carlson	Internet2 Extensions Coordinator	No	No	Internet2
Jeremy Dodd	Co-PI	No	No	Columbia
Dan Fraser	Production Coordinator	Yes	Yes	UChicago
Robert Gardner	Co-PI & Integration Coordinator	Yes	Yes	UChicago
Sebastien Goasguen	PI & Campus Grids Coordinator	Yes	Yes	Clemson
Howard Gordon	Co-PI	No	No	BNL
Anne Heavey	iSGTW Editor	Yes	Yes	Fermilab
Matt Crawford	Storage Extensions Coordinator	Yes	Yes	Fermilab
Tanya Levshina	Storage Software Coordinator	Yes	Yes	Fermilab
Fred Luehring	Co-PI	No	No	Indiana
Scott McCaulay	Co-PI	No	No	Indiana
John McGee	Co-PI & Engagement Coordinator	No	Yes	RENCI
Doug Olson	Co-PI	Yes	Yes	LBNL
Maxim Potekhin	Extensions-WMS Coordinator	Yes	Yes	BNL
Robert Quick	Operations Coordinator	Yes	Yes	Indiana
Abhishek Rana	VOs Group Coordinator	Yes	Yes	UCSD
Alain Roy	Software Coordinator	Yes	Yes	UWisconsin
David Ritchie	Communications Coordinator	No	Yes	Fermilab
Chander Sehgal	Project Manager	Yes	Yes	Fermilab
Igor Sfiligoi	Extensions Scalability Coordinator	Yes	Yes	UCSD
Piotr Sliz	PI	No	No	Harvard
David Swanson	PI	No	No	UNebraska
Todd Tannenbaum	Condor Coordinator	Yes	Yes	UWisconsin
John Towns	Co-PI	No	No	UIUC
Mike Tuts	Co-PI	No	No	Columbia
Shaowen Wang	PI	No	Yes	UIUC
Torre Wenaus	Co-PI & Extensions Co-Coordinator	No	Yes	BNL
Michael Wilde	Co-PI	Yes	Yes	UChicago
Frank Wuerthwein	PI & Extensions Co-Coordinator	No	Yes	UCSD
Technical Staff				
Linton Abraham	Staff	Yes	Yes	Clemson
Warren Andrews	Staff	Yes	Yes	UCSD
Charles Bacon	Staff	Yes	Yes	UChicago
Andrew Baranovski	Staff	Yes	Yes	Fermilab
James Basney	Staff	Yes	Yes	UIUC
Chris Bizon	Staff	No	Yes	RENCI
Jose Caballero	Staff	Yes	Yes	BNL
Tim Cartwright	Staff	Yes	Yes	UWisconsin
Keith Chadwick	Staff	Yes	Yes	Fermilab
Barnett Chiu	Staff	No	No	BNL
Elizabeth Chism	Staff	Yes	Yes	Indiana
Ben Clifford	Staff	Yes	Yes	UChicago
Toni Coarasa	Staff	Yes	Yes	UCSD
Simon Connell	Staff	No	No	Columbia
Ron Cudzewicz	Staff	Yes	No	Fermilab

Britta Daudert	Staff	Yes	Yes	Caltech
Peter Doherty	Staff	Yes	Yes	Harvard
Ben Eisenbraun	Staff	No	No	Harvard
Robert Engel	Staff	Yes	Yes	Caltech
Michael Ernst	Staff	No	No	BNL
Jamie Frey	Staff	Yes	Yes	UWisconsin
Arvind Gopu	Staff	No	Yes	Indiana
Chris Green	Staff	Yes	Yes	Fermilab
Kyle Gross	Staff	Yes	Yes	Indiana
Soichi Hayashi	Staff	Yes	Yes	Indiana
Ted Hesselroth	Staff	Yes	Yes	Fermilab
John Hover	Staff	Yes	No	BNL
Keith Jackson	Staff	Yes	Yes	LBNL
Scot Kronenfeld	Staff	Yes	Yes	UWisconsin
Tom Lee	Staff	No	Yes	Indiana
Ian Levesque	Staff	No	No	Harvard
Marco Mambelli	Staff	Yes	Yes	UChicago
Doru Marcusiu	Staff	No	No	UIUC
Terrence Martin	Staff	Yes	Yes	UCSD
Jay Packard	Staff	Yes	No	BNL
Sanjay Padhi	Staff	Yes	Yes	UCSD
Anand Padmanabhan	Staff	Yes	Yes	UIUC
Christopher Pipes	Staff	Yes	Yes	Indiana
Jeff Porter	Staff	Yes	Yes	LBNL
Craig Prescott	Staff	No	No	UFlorida
Mats Rynge	Staff	No	Yes	RENCI
Iwona Sakrejda	Staff	Yes	Yes	LBNL
Aashish Sharma	Staff	Yes	Yes	UIUC
Neha Sharma	Staff	Yes	Yes	Fermilab
Tim Silvers	Staff	Yes	Yes	Indiana
Alex Sim	Staff	Yes	Yes	LBNL
Ian Stokes-Rees	Staff	No	Yes	Harvard
Marcia Teckenbrock	Staff	Yes	Yes	Fermilab
Greg Thain	Staff	Yes	Yes	UWisconsin
Suchandra Thapa	Staff	Yes	Yes	UChicago
Aaron Thor	Staff	Yes	Yes	BNL
Von Welch	Staff	Yes	No	UIUC
James Weichel	Staff	Yes	Yes	UFlorida
Amelia Williamson	Staff	Yes	No	UFlorida

2.2. Partner Organizations

Here you let NSF know about partner organizations outside your own institution – academic institutions, other nonprofits, industrial or commercial firms, state or local governments, schools or school systems, or whatever – that have been involved with your project. Partner organizations may provide financial or in-kind support, supply facilities or equipment, collaborate in the research, exchange personnel, or otherwise contribute. The screens will lead you through the obvious possibilities, but will also give you an opportunity to identify out-of-the-ordinary partnership arrangements and to describe any arrangement in a little more detail.

Partner Organizations – Why?

NSF cannot achieve its ambitious goals for the science and technology base of our country with its own resources alone. So we place strong emphasis on working in partnership with other public and private organizations engaged in science, engineering, and education and on encouraging partnerships among such organizations. We also seek partnerships across national boundaries, working with comparable organizations in other countries wherever mutually beneficial.

So we need to gauge and report our performance in promoting partnerships. We need to know about the partnerships in which our awardees have engaged and to what extent they have been effective.

We use a pre-established list of organizations to ensure consistency and to avoid both lost information and double counting where the same organization is identified by different names.

The members of the Council and List of Project Organizations

1. Boston University
2. Brookhaven National Laboratory
3. California Institute of Technology
4. Clemson University
5. Columbia University
6. Cornell University
7. Distributed Organization for Scientific and Academic Research (DOSAR)
8. Fermi National Accelerator Laboratory
9. Harvard University (medical school)
10. Indiana University
11. Information Sciences Institute/University of South California
12. Lawrence Berkeley National Laboratory
13. Purdue University
14. Renaissance Computing Institute
15. Stanford Linear Accelerator Center (SLAC)
16. University of California San Diego
17. University of Chicago
18. University of Florida
19. University of Illinois Urbana Champaign/NCSA
20. University of Nebraska – Lincoln
21. University of Wisconsin, Madison

2.3. Participants: Other Collaborators

You might let NSF know about any significant:

** collaborations with scientists, engineers, educators, or others within your own institution – especially interdepartmental or interdisciplinary collaborations;*

** non-formal collaborations or contacts with scientists, engineers, educators, or others outside your institution; and*

** non-formal collaborations or contacts with scientists, engineers, educators, or others outside the United States.*

The OSG relies on external project collaborations to develop the software to be included in the VDT and deployed on OSG. Collaborations are in progress with: Community Driven Improvement of Globus Software (CDIGS), SciDAC-2 Center for Enabling Distributed Petascale Science (CEDPS), Condor, dCache collaboration, Data Intensive Science University Network (DISUN), Energy Sciences Network (ESNet), Internet2, National LambdaRail (NLR), BNL/FNAL Joint Authorization project, LIGO Physics at the Information Frontier, Fermilab Gratia Accounting, SDM project at LBNL (BeStMan), SLAC Xrootd, Pegasus at ISI, U.S. LHC software and computing.

OSG also has close working arrangements with “Satellite” projects, defined as independent projects contributing to the OSG roadmap, with collaboration at the leadership level. Current Satellite projects include:

- “Embedded Immersive Engagement for Cyberinfrastructure”, (CI-Team, OCI funded, NSF 0753335)
- Structural Biology Grid: based from Harvard Medical School; 114 partner labs – Piotr Sliz, Ian Stokes-Rees (MCB funded)
- VOSS: “Delegating Organizational Work to Virtual Organization Technologies: Beyond the Communications Paradigm” (OCI funded, NSF 0838383)
- CILogon: “Secure Access to National-Scale CyberInfrastructure” (OCI funded, NSF 0850557)

3. Activities and Findings:

3.1. Research and Education Activities

OSG provides an infrastructure that supports a broad scope of scientific research activities, including the major physics collaborations, nanoscience, biological sciences, applied mathematics, engineering, computer science and, through the engagement program, other non-physics research disciplines. The distributed facility is quite heavily used, as described below and in the attached document showing usage charts.

OSG continued to provide a laboratory for research activities that deploy and extend advanced distributed computing technologies in the following areas:

- Integration of the new LIGO Data Grid security infrastructure, based on Kerberos identity and Shibboleth/Groupier authorization, with the existing PKI authorization infrastructure, across the LIGO Data Grid (LDG) and OSG.

- Support of inter-grid gateways which transport information, accounting, service availability information between OSG and European Grids supporting the LHC Experiments (EGEE/WLCG).
- Research on the operation of a scalable heterogeneous cyber-infrastructure in order to improve its effectiveness and throughput. As part of this research we have developed a comprehensive “availability” probe and reporting infrastructure to allow site and grid administrators to quantitatively measure and assess the robustness and availability of the resources and services.
- Scalability and robustness enhancements to Condor technologies. For example, extensions to Condor to support Pilot job submissions have been developed, significantly increasing the job throughput possible on each Grid site.
- Deployment and scaling in the production use of “pilot-job” workload management system – ATLAS PanDA and CMS glideinWMS. These developments were crucial to the experiments meeting their analysis job throughput targets.
- Scalability and robustness enhancements to Globus grid technologies. For example, comprehensive testing of the Globus Web-Service Gram which has resulted in significant coding changes to meet the scaling needs of OSG applications
- Development of an at-scale test stand that provides hardening and regression testing for the many SRM V2.2 compliant releases of the dCache, BeStMan, and Xrootd storage software.
- Integration of BOINC-based applications (LIGO’s Einstein@home) submitted through grid interfaces.
- Further development of a hierarchy of matchmaking services (OSG MM), ReSS or REsource Selection Services that collect information from more than 60 OSG sites and provide a VO based matchmaking service that can be tailored to particular application needs.
- Investigations and testing of policy and scheduling algorithms to support “opportunistic” use and backfill of resources that are not otherwise being used by their owners, using information services such as GLUE, matchmaking and workflow engines including Pegasus and Swift.
- Comprehensive job accounting across 76 OSG sites, publishing summaries for each VO and Site, and providing a per-job information finding utility for security forensic investigations.

The key components of OSG’s education program are:

- Organization and participation in more than 6 grid schools and workshops, including invited workshops at the PASI meeting in Costa Rica and the first US eHealthGrid conference, and co-sponsorship of the International Grid Summer School in Hungary as well as the Online International Grid Winter School which was totally electronically based.
- Active participation in more than 5 “Campus Infrastructure Days (CI Days) events. CI Days is an outreach activity in collaboration with Educause, Internet2, TeraGrid and the MSI institutions. Each event brings together local faculty, educators and IT personnel to learn about their combined needs and to facilitate local planning and activities to meet the cyber-infrastructure needs of the communities.

- Invited participation in the TeraGrid Supercomputing 08 education workshop, participation in the Grace Hopper Conference GHC08 October 1-4, Colorado and Applications of HPC, Grids, and Parallel Computing to Science Education Aug 15, 2008, U of Oklahoma
- Support for student computer science research projects from the University of Chicago, performing FMRI analysis and molecular docking, as well as evaluating the performance and usability of the OSG infrastructure.

3.2. Findings

- Scientists and researchers can successfully use a heterogeneous computing infrastructure with job throughputs of more than 25,000 CPU days per day (an increase of an average of 5,000 CPU days per day over the last six months), dynamically shared by up to ten different research groups, and with job-related data placement needs of the order of Terabytes.
- Initial use of opportunistic storage in conjunction with opportunistic processing provides value and can significantly increase the effectiveness of job throughput and performance.
- Federating the local identity/authorization attributes with the OSG authorization infrastructure is possible. We know there are multiple local identity/authorization implementations and it is useful to have an exemplar of how to integrate with at least one.
- The effort and testing required for inter-grid bridges involves significant costs, both in the initial stages and in continuous testing and upgrading. Ensuring correct, robust end-to-end reporting of information across such bridges remains fragile and human effort intensive.
- Availability and reliability testing, accounting information and their interpretation are proving their worth in maintaining the attention of the site administrators and VO managers. This information is not yet complete. Validation of the information is also incomplete, needs additional attention, and can be effort intensive.
- The scalability and robustness of the infrastructure has reached the performance needed for initial LHC data taking, but not yet reached the scales needed by the LHC when it reaches stable operations. The goals for the commissioning phase in FY09 have been met and are only now being sustained over sufficiently long periods.
- The job “pull” architecture does indeed give better performance and management than the “push” architecture.
- Automated site selection capabilities are proving their worth when used. However they are inadequately deployed. They are also embryonic in the capabilities needed – especially when faced with the plethora of errors and faults that are encountered on a loosely coupled set of independent computing and storage resources used by a heterogeneous mix of applications with greatly varying I/O, CPU and data requirements.
- Analysis of accounting and monitoring information is a key need which requires dedicated and experienced effort.
- Transitioning students from the classroom to be users is possible but continues as a challenge, partially limited by the effort OSG can dedicate to this activity.
- Many communities are facing the same challenges as OSG in educating new entrants to get over the threshold of understanding and benefiting from distributed computing.

3.2.1. Findings enabled by the Distributed Infrastructure: Science Deliverables

Physical Sciences:

CMS: US-CMS relies on Open Science Grid for critical computing infrastructure, operations, and security services. These contributions have allowed US-CMS to focus experiment resources on being prepared for analysis and data processing, by saving effort in areas provided by OSG. OSG provides a common set of computing infrastructure on top of which CMS, with development effort from the US, has been able to build a reliable processing and analysis framework that runs on the Tier-1 facility at Fermilab, the project supported Tier-2 university computing centers, and opportunistic Tier-3 centers at universities. There are currently 18 Tier-3 centers registered with the CMS computing grid in the US which provide additional simulation and analysis resources to the US community. In addition to common interfaces, OSG has provided the packaging, configuration, and support of the storage services. Since the beginning of OSG the operations of storage at the Tier-2 centers have improved steadily in reliability and performance. OSG is playing a crucial role here for CMS in that it operates a clearinghouse and point of contact between the sites that deploy and operate this technology and the developers. In addition, OSG fills in gaps left open by the developers in areas of integration, testing, and tools to ease operations. The stability of the computing infrastructure has not only benefitted CMS. CMS' use of resources has been very much cyclical so far, thus allowing for significant use of the resources by other scientific communities. OSG is an important partner in Education and Outreach, and in maximizing the impact of the investment in computing resources for CMS and other scientific communities.

In addition to computing infrastructure OSG plays an important role in US-CMS operations and security. OSG has been crucial to ensure US interests are addressed in the WLCG. The US is a large fraction of the collaboration both in terms of participants and capacity, but a small fraction of the sites that make-up WLCG. OSG is able to provide a common infrastructure for operations including support tickets, accounting, availability monitoring, interoperability and documentation. As CMS has entered the operations phase, the need for sustainable security models and regular accounting of available and used resources has become more important. The common accounting and security infrastructure and the personnel provided by OSG is a significant service to the experiment.

ATLAS: US ATLAS continues to depend crucially on the OSG infrastructure. All our facilities deploy the OSG software stack as the base upon which we install the ATLAS software system. The OSG has been helpful in improving usability of the grid as seen by US ATLAS production and analysis, and mitigating problems with grid middleware. Examples include

- GRAM dependency in CondorG submission of pilots, limiting the scalability of PanDA pilot submission on the grid. The OSG WMS program has developed a 'pilot factory' to work around this by doing site-local pilot submission without every pilot seeing the gatekeeper and GRAM.
- gLExec for analysis user tracing and identity management, now deployed for production by FNAL/CMS and planned for EGEE deployment soon. US ATLAS will benefit from its addition to the OSG software stack, and has benefitted from OSG WMS support in integrating gLExec with PanDA.

- OSG-standard site configuration, providing a ‘known’ environment on OSG WNs. This has lessened the application-level work of establishing homogeneity.
- Tools for resource discovery. We use OSG tools to gather the information on resource availability, health, and access rights that is required to fully utilize the resources available.
- Supported storage systems and their SRM v2.2 interfaces, including dCache (3 Tier-2 sites) and BeStMan-Xrootd (2 Tier-2 sites). In addition, we anticipate BeStMan-Xrootd systems to become adopted by several Tier-3 facilities in the coming year, and so will rely on the continued packaging, testing, and support provided by the OSG Storage teams.
- Software components that have allowed interoperability with European ATLAS sites, including selected components from the gLite middleware stack including LCG client utilities (for file movement, supporting space tokens as required by ATLAS), and file catalogs (server and client).
- We anticipate adoption of Internet2 monitoring tools such as perfSonar and NDT within the VDT, which will provide another support point for network troubleshooting as regards both Tier-2 and Tier-3 facilities.

We greatly benefit from OSG's Gratia accounting services, as well as the information services and probes that provide OSG usage and site information to the application layer and to the WLCG for review of compliance with MOU agreements. We rely on the VDT and OSG packaging, installation, and configuration processes that lead to a well-documented and easily deployable OSG software stack, and OSG's integration testbed and validation processes that accompany incorporation of new services into the VDT. US ATLAS and ATLAS operations increasingly make use of the OSG trouble ticketing system (which distributes tickets originating from OSG and EGEE to the US ATLAS RT tracking system) and the OSG OIM system which communicates downtimes of US ATLAS resources to WLCG and International ATLAS. We also benefit from and rely on the infrastructure maintenance aspects of the OSG such as the GOC that keep the virtual US ATLAS computing facility and the OSG facility as a whole operational.

The US-developed PanDA distributed production and analysis system based on just-in-time (pilot based) workflow management is in use ATLAS-wide for production and analysis, and is (since 2006) a part of the OSG's workload management effort as well. Both ATLAS and OSG have benefited from this activity. The OSG WMS effort has been the principal driver for improving the security of the PanDA system, in particular its pilot job system, bringing it into compliance with security policies within the OSG and WLCG, in particular the requirement that gLExec be used for user analysis jobs to assign the job's identity to that of the analysis user. The OSG WMS effort also continues to deepen the integration of PanDA with the Condor job management system, which lies at the foundation of PanDA's pilot submission infrastructure. For the OSG, PanDA has been deployed as a tool and service available for general OSG use. A team of biologists uses PanDA and OSG facilities for protein folding simulation studies (using the CHARMM simulation code) underpinning a recent research paper, and additional users are trying out PanDA. We are increasing PanDA's offerings to the OSG community with a present focus on offering VOs simple data handling tools that allow them to integrate their data into a Panda-managed workflow. Reciprocally the OSG WMS effort will continue to be the principal source for PanDA security enhancements, further integration with middleware and particularly Condor, and scalability/stress testing of current components and new middleware integration.

LIGO: The Einstein@Home data analysis application that searches for gravitational radiation from spinning neutron stars using data from the Laser Interferometer Gravitational Wave Observatory (LIGO) detectors was identified over a year ago as an excellent LIGO application for migration onto the Open Science Grid (OSG). This is due to the fact that this particular search is virtually unbounded in the scientific merit achieved by additional computing resources. The original deployment in spring of 2008 was based on the WS-Gram interface which had limited availability on the OSG. Late in 2008, the Einstein@Home grid application effort began to rework the application to support the Globus Toolkit 2 Gram interface supported on all OSG sites. Beginning in February of 2009, the new application was deployed on the Open Science Grid. Several modifications to the code ensued to address stability, reliability and performance. By May of 2009, the code was running reliably in production on close to 20 sites across the OSG that support job submission from the LIGO Virtual Organization.

The Einstein@Home application is now averaging roughly 6,000 CPU hours per day on the OSG (see Figure 2). In terms of scientific contributions to the search for spinning neutron stars, this accounts for approximately 160,000 Einstein@Home Credits per day (a “Credit” is defined as a unit of data analysis by the Einstein@Home team; on average the OSG contributes slightly more than 1 Credits per CPU hour.) with a peak performance of 210,000 credits seen in a single day. The total contributions to the Einstein@Home search from the OSG is now ranked 30th in the world based on all credits since November 2008 and is on a daily bases in the top ten contributors, averaging 9th place in the world at this time. In the future, LIGO plans to reengineer the job submission side of the Einstein@Home to utilize Condor-G instead of raw GRAM job submissions to improve the loading and reduce overhead seen on OSG gatekeepers. This should allow more reliable job submission and provide further improvements in efficiency.

In the past year, LIGO has also begun to investigate ways to migrate the data analysis workflows searching for gravitational radiation from binary black holes and neutron stars onto the Open Science Grid for production scale utilization. The binary inspiral data analyses typically involve working with tens of terabytes of data in a single workflow. Collaborating with the Pegasus Workflow Planner developers at USC-ISI, LIGO has identified changes to both Pegasus and to the binary inspiral workflow codes to more efficiently utilize the OSG where data must be moved from LIGO archives to storage resources near the worker nodes on OSG sites. One area of particular focus has been on the understanding and integration of Storage Resource Management (SRM) technologies used in OSG Storage Element (SE) sites to house the vast amounts of data used by the binary inspiral workflows so that worker nodes running the binary inspiral codes can effectively access the data. To date this has involved standing up a SRM Storage Element on the LIGO Caltech OSG integration testbed site. This site has 120 CPU cores with approximately 30 terabytes of storage currently configured under SRM. The SE is using BeStMan and Hadoop for the distributed file system shared among the worker nodes. This effort is just beginning and will require further integration into Pegasus for the workflow planning to begin to evaluate the nuances of migration onto the OSG production grid. How to properly advertise OSG SE configuration information to most efficiently utilize the combination of storage and computation necessary to carry out the binary inspiral gravitation radiation searches is also an active area for this research.

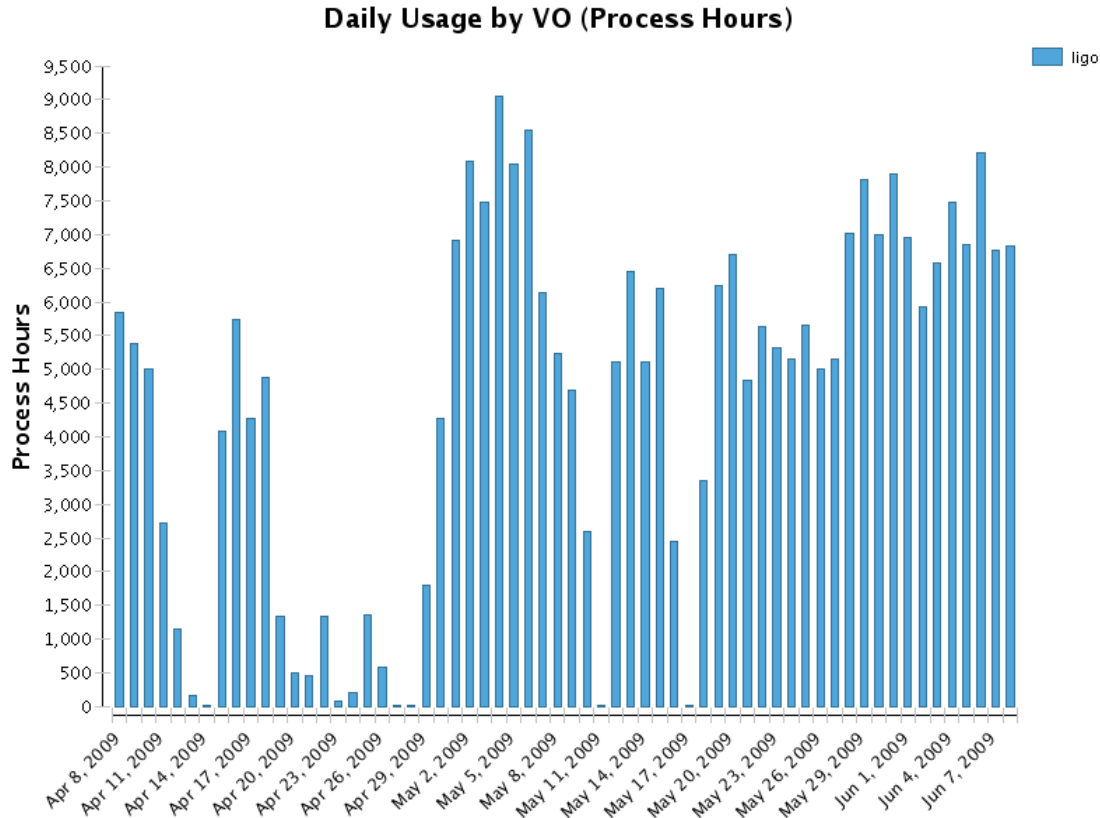


Figure 2: OSG Usage by LIGO's Einstein@Home application for the two month period covering both the month before full deployment of the new code and the first month of running at production levels with the new code using the GRAM 2 job submission interface.

LIGO has also been working closely with the OSG to evaluate the implications of its requirements on authentication and authorization within its own LIGO Data Grid and how these requirements map onto the security model of the OSG and the Department of Energy Grids Certificate Authority policies. This has involved close collaboration between the LIGO Scientific Collaboration's Auth Project and the OSG security team.

D0 at Tevatron: The D0 experiment continues to rely heavily on OSG infrastructure and resources in order to achieve the computing demands of the experiment. The D0 experiment has successfully used OSG resources for many years and plans on continuing with this very successful relationship into the foreseeable future.

All D0 Monte Carlo simulation is generated at remote sites, with OSG continuing to be a major contributor. During the past year, OSG sites simulated 330 million events for D0, approximately 1/3 of all production. An extensive study was undertaken in 2008 to understand and increase production efficiencies, which varied significantly site to site. It was determined that sites that did not have local storage elements had lower job efficiencies than those that did. D0 thereupon requested OSG to have relevant sites implement local storage elements and worked with Fermilab Computing Division to improve the infrastructure on the experiment's side. The resulting improvements greatly increased the job efficiency of Monte Carlo production.

Over the past year, the average number of Monte Carlo events produced per week by OSG has nearly doubled. In September 2008, D0 had its first 10 million events produced in a week by

OSG. In recent months 10 million events/week is becoming the standard and a new record of 13 million events/week was set in May 2009. Much of this increase is due to improved efficiency, increased resources, (D0 used 24 sites in the past year and uses 21 regularly), automated job submission, use of resource selection services and expeditious use of opportunistic computing. D0 plans to continue to work with OSG and Fermilab computing to continue to improve the efficiency of Monte Carlo production on OSG sites.

The primary processing of D0 data continues to be run using OSG infrastructure. One of the very important goals of the experiment is to have the primary processing of data keep up with the rate of data collection. It is critical that the processing of data keep up in order for the experiment to quickly find any problems in the data and to keep the experiment from having a backlog of data. D0 is able to keep up with the primary processing of data by reconstructing nearly 6 million events/day. Over the past year D0 has reconstructed over 2 billion events on OSG facilities.

OSG resources have allowed D0 to meet its computing requirements in both Monte Carlo production and in data processing. This has directly contributed to D0's 40 published papers during the past year.

CDF at Tevatron: The CDF experiment continues to use OSG infrastructure and resources in order to provide the collaboration with enough Monte Carlo data to keep a high level of physics results. CDF, in collaboration with OSG, aims to improve the infrastructural tools in the next years to increase the Grid resources usage.

During last six months CDF has been operating the pilot-based Workload Management System (glideinWMS) as the submission method to remote OSG sites. This system went into production three months ago on the CDF North American Grid (NAmGrid) portal. Figure 3 shows the number of running jobs on NAmGrid and demonstrates that there has been steady usage of the facilities, while Figure 4, a plot of the queued requests, shows that there is large demand. The emphasis or recent work has been to validate sites for reliable usage of Monte Carlo generation and to develop metrics to demonstrate smooth operations. One impediment to smooth operation has been the rate at which jobs are lost and re-started by the batch system. It should be noted that there were a significant number of restarts until week 21, after which the rate tailed down significantly. At that point, it was noticed that most re-starts occurred at specific sites, which were subsequently removed from NamGrid. Those sites and any new site will be tested and certified in integration using Monte Carlo jobs that have previously been run in production. We are also adding more monitoring to the CDF middleware to allow faster identification of problem sites or individual worker nodes. Issues of data transfer and the applicability of opportunistic storage is being studied as part of the effort to understand issues affecting reliability.




Figure 3: Running CDF jobs on NAMGrid




Figure 4: Waiting CDF jobs on NAMGrid, showing large demand

A legacy glide-in infrastructure developed by the experiment is still running on the portal to on-site OSG resources (CDFGrid). Plots of the running jobs and queued requests are shown in Figure 5 and Figure 6. Among the major issues we encountered in achieving smooth and efficient operations was a serious unscheduled downtime in April. Subsequent analysis found the direct cause to be incorrect parameters set on disk systems serving the OSG gatekeeper software stack and data output areas. No OSG software was implicated in the root cause analysis. There were also losses of job slots due to attempts to turn on opportunistic usage. The proper way to handle this is still being investigated. Instabilities in Condor software caused job loss at various times. Recent Condor upgrades have led to steadier running on CDFGrid. Finally, Job re-starts on CDFGrid cause problems in data handling and job handling synchronization. A separate effort is under way to identify the causes for these re-starts and to provide recovery tools.




Figure 5: Running CDF jobs on CDFGrid




Figure 6: Waiting CDF jobs on CDFGrid

CDF recently conducted a review of the CDF middleware and usage of Condor and OSG. While there were no major issues, a number of cleanup projects have been identified that will add to the long-term stability and maintainability of the software. These projects are now being executed. The use of glideinWMS in CDFGrid is planned. Integration testing is completed; deployment awaits the end of the summer conference season.

Thanks to OSG resources and infrastructure CDF has been able to publish another 50 physics papers during this year including 4 discoveries in the last six months.

Nuclear physics: The STAR experiment has continued the use of data movement capabilities between its established Tier-1 and Tier-2 centers and between BNL and LBNL (Tier-1), Wayne State University and NPI/ASCR in Prague (two fully functional Tier-2 centers). A new center,

the *Korea Institute of Science and Technology* Information (KISTI) has joined the STAR collaboration as a full partnering facility and resource provider in 2008 and activities surrounding the exploitation of this new potential have taken a large part of STAR's activity in the 2008/2009 period.

The RHIC run 2009 had been projected to bring to STAR a fully integrated new data acquisition system with data throughput capabilities going from 100 MB/sec reached in 2004 to 1000 MB/sec. This is the second time in the experiment's lifetime STAR computing has to cope with an order of magnitude growth in data rates. Hence, a threshold in STAR's Physics program was reached where leveraging all resources across all available sites has become essential to success. Since the resources at KISTI have the potential to absorb up to 20% of the needed cycles for one pass data production in early 2009, efforts were focused on bringing the *average* data transfer throughput from BNL to KISTI to 1 Gb/sec. It was projected (Section 3.2 of the STAR computing resource planning, "The STAR Computing resource plan", STAR Notes CSN0474, <http://drupal.star.bnl.gov/STAR/starnotes/public/csn0474>) that such a rate would sustain the need up to 2010 after which a maximum of 1.5 Gb/sec would cover the currently projected Physics program up to 2015. Thanks to the help from ESNet, Kreonet and collaborators at both end institutions this performance was reached (see <http://www.bnl.gov/rhic/news/011309/story2.asp>, "From BNL to KISTI: Establishing High Performance Data Transfer From the US to Asia" and ¹ <http://www.lbl.gov/cs/Archive/news042409c.html>, "ESnet Connects STAR to Asian Collaborators"). At this time baseline Grid tools are used and the OSG software stack has not yet been deployed. STAR plans to include a fully automated job processing capability and return of data results using BeStMan/SRM (Berkeley's implementation of SRM server).

Encouraged by the progress on the network tuning for the BNL/KISTI path and driven by the expected data flood from Run-9, the computing team is re-addressing all of its network data transfer capabilities, especially between BNL and NERSC and between BNL and MIT. MIT has been a silent Tier-2, a site providing resources for local scientist's research and R&D work but has not been providing resources to the collaboration as a whole. MIT has been active since the work made on Mac/X-Grid reported in 2006, a well-spent effort which has evolved in leveraging additional standard Linux-based resources. Data samples are routinely transferred between BNL and MIT. The BNL/STAR gatekeepers have all been upgraded and all data transfer services are being re-tuned based on the new topology. Initially planned for the end of 2008, the strengthening of the transfers to/from well established sites was a delayed milestone (6 months) to the benefit of the BNL/KISTI data transfer.

At Prague / Bulovka, data transfers are also handled using a BeStMan SRM client but in interoperability mode with a *Disk Pool Manager* (DPM) SRM door. Xrootd remains the low-human cost middleware of choice for STAR and its Tier-2 center storage aggregation strategy but sites such as Prague typically rest on components such as DPM, already deployed within the context of other grid projects. Data rates between BNL and Prague, reaching 300 Mb/sec at the moment, are sufficient to sustain the local needs. Local data access in Prague rests on the use of the *STAR Unified Meta-Scheduler* (SUMS) offering users a common interface for job submission. STAR's approach provides a transparent submission interface to both Grid and non-Grid resources and SUMS remains at the heart of STAR's strategy to migrate an entire class of jobs to Grid resources. Analysis of data sets now entirely relies on access to Scalla/Xrootd data aggregation at BNL (since 2006) and DPM/rfio access at Prague (2007/2008). Users make extensive use of SUMS abstraction to seamlessly launch jobs on the respective farms; the same

job description works on both farms. STAR has plans to utilize the Prague resources for opportunistic Monte-Carlo event processing by mid to end of 2009.

A research activity involving STAR and the computer science department at Prague has been initiated to improve the data management program and network tuning. We will study and test a multi-site data transfer paradigm, coordinating movement of datasets to and from multiple locations (sources) in an optimal manner, using a planner taking into account the performance of the network and site. This project relies on the knowledge of file locations at each site and a known network data transfer speed as initial parameters (as data is moved, speed can be re-assessed so the system is a self-learning component). The project has already shown impressive gains over a standard peer-to-peer approach for data transfer. Although this activity has so far impacted OSG in a minimal way, we will use the OSG infrastructure to test our implementation and prototyping at the end of summer 2009. To this end, we paid close attention to protocols and concepts used in Caltech's Fast Data Transfer (FDT) tool as its streaming approach has non trivial consequence and impact on TCP protocol shortcomings.

STAR has continued to use and consolidate the BeStMan/SRM implementation and has engaged in active discussions, steering and integration of the messaging format from the Center for Enabling Distributed Petascale Science's (CEDPS) Troubleshooting team, in particular targeting use of BeStMan client/server troubleshooting for faster error and performance anomaly detection and recovery. At the time of this report, tests and a base implementation are underway to pass BeStMan based messages using syslog-ng. Several problems have already been found, leading to better and more robust implementations. We believe we would have a case study within months and able to determine if this course of action represents a path forward to distributed message passing. STAR has finished developing its own job tracking and accounting system, a simple approach based on adding tags at each stage of the workflow and collecting the information via recorded database entries and log parsing. The work was presented at the CHEP 2009 conference (*Workflow generator and tracking at the rescue of distributed processing. Automating the handling of STAR's Grid production*, Contribution ID 475, CHEP 2009, <http://indico.cern.ch/contributionDisplay.py?contribId=475&confId=35523>). The STAR SBIR Tech-X/UCM project, aimed to provide a fully integrated *User Centric Monitoring* (UCM) toolkit, has reached its end-of-funding cycle. The project is being absorbed by STAR personnel who aim to deliver a workable monitoring scheme at application level. The library has been used in nightly and regression testing to help further development (mainly scalability, security and integration into Grid context). The knowledge and a working infrastructure based on syslog-ng may very well provide a simple mechanism for merging UCM with CEDPS vision.

STAR grid data processing and job handling operations have continued their progression toward a full Grid-based operation relying on the OSG software stack and the OSG Operation Center issue tracker. The STAR operation support team has been efficiently addressing issues and stability. Overall the grid infrastructure stability seems to have increased. To date, STAR has however mainly achieved simulated data production on Grid resources. Since reaching a milestone in 2007, it has become routine to utilize non-STAR dedicated resources from the OSG for the Monte-Carlo event generation pass and to run the full response simulator chain (requiring the whole STAR framework installed) on STAR's dedicated resources. On the other hand, the relative proportion of processing contributions using non-STAR dedicated resources has been marginal (and mainly on the FermiGrid resources in 2007). This disparity is explained by the fact that the complete STAR software stack and environment, which is difficult to impossible to

recreate on arbitrary grid resources, is necessary for full event reconstruction processing and hence, access to generic and opportunistic resources are simply impractical and not matching the realities and needs of running experiments in Physics production mode. In addition, STAR's science simply cannot suffer the risk of heterogeneous or non-reproducible results due to subtle library or operating system dependencies and the overall workforce involved to ensure seamless results on all platforms exceeds our operational funding profile. Hence, STAR has been a strong advocate for moving toward a model relying on the use of Virtual Machine (see contribution at the OSG booth @ CHEP 2007) and have since closely work, to the extent possible, with the CEDPS Virtualization activity, seeking the benefits of truly opportunistic use of resources by creating a complete pre-packaged environment (with a validated software stack) in which jobs will run. Such approach would allow STAR to run any one of its job workflow (event generation, simulated data reconstruction, embedding, real event reconstruction and even user analysis) while respecting STAR's policies of reproducibility implemented as complete software stack validation. The technology has huge potential in allowing (beyond a means of reaching non-dedicated sites) software provisioning of Tier-2 centers with minimal workforce to maintain the software stack hence, maximizing the return to investment of Grid technologies. The multitude of combinations and the fast dynamic of changes (OS upgrade and patches) make the reach of the diverse resources available on the OSG, workforce constraining and economically un-viable.

This activity reached a world-premiere milestone when STAR made use of the Amazon/EC2 resources, using Nimbus Workspace service to carry part of its simulation production and handle a late request. These activities were written up in iSGTW (*Clouds make way for STAR to shine*, <http://www.isgtw.org/?pid=1001735>, Newsweek (*Number Crunching Made Easy - Cloud computing is making high-end computing readily available to researchers in rich and poor nations alike* <http://www.newsweek.com/id/195734>), SearchCloudComputing (*Nimbus cloud project saves brainiacs' bacon* http://searchcloudcomputing.techtarget.com/news/article/0,289142,sid201_gci1357548,00.html) and HPCWire (*Nimbus and Cloud Computing Meet STAR Production Demands* <http://www.hpcwire.com/offthewire/Nimbus-and-Cloud-Computing-Meet-STAR-Production-Demands-42354742.html?page=1>). This was the very first time cloud computing had been used in the HENP field for scientific production work with full confidence in the results. The results were presented during a plenary talk at CHEP 2009 conference where others presented "tests" rather than actual use (Belle Monte-Carlo testing was most interesting as well). We believe this represents a breakthrough and have since, actively engaged in discussions with the OSG management for the inclusion of such technology into the program of work (present or future) of the Open Science Grid project.

All STAR physics publications acknowledge the resources provided by the OSG.

MINOS: Over the last three years, computing for MINOS data analysis has greatly expanded to use more of the OSG resources available at Fermilab. The scale of computing has increased from about 50 traditional batch slots to typical user jobs running on over 1,000 cores, with a strong desire to expand to about 5,000 cores (over the past 12 months they have used 3.1M hours on OSG from 1.16M submitted jobs). This computing resource, combined with 90 TBytes of dedicated BlueArc (NFS mounted) file storage, has allowed MINOS to move ahead with traditional and advanced analysis techniques, such as Neural Network, Nearest Neighbor, and Event Library methods. These computing resources are critical as the experiment moves beyond the early, somewhat simpler Charged Current physics, to more challenging Neutral Current, $\nu+e$ and other analyses which push the limits of the detector. We use a few hundred cores of

offsite computing at collaborating universities for occasional Monte Carlo generation. MINOS is also starting to use TeraGrid resources at TACC, hoping to greatly speed up their latest processing pass.

Astrophysics: The Dark Energy Survey (DES) used approximately 20,000 hours of OSG resources in 2008, with DES simulation activities ramping up in the latter part of the year. The most recent DES simulation produced 3.34 Terabytes of simulated imaging data, which were used for testing the DES data management data processing pipelines as part of the so-called Data Challenge 4. These simulations consisted of 2,600 mock science images of the sky, along with another 740 calibration images, each 1 GB in size. Each image corresponds to a single job on OSG and simulates the sky covered in a single 3-square-degree pointing of the DES camera. The processed simulated data are also being actively used by the DES science working groups for development and testing of their science analysis codes. DES expects to roughly double its usage of OSG resources over the following 12 months.

Structural Biology: During the past year SBGrid-RCN (Structural Biology Research Coordination Network) has become actively involved with OSG in several activities. In 2008 they integrated two computing clusters at Harvard Medical School with OSG. The initial configuration successfully supported isolated chunks of computations, but more work had to be performed to establish a sustainable grid infrastructure. In particular, although their grid resources were accessible for internal job submissions, some critical system probes were failing, and therefore SBGrid was inaccessible to external sites.

Within the last 12 months, in phase II of the project, they have fine-tuned the setup and currently operate within stringent, predefined site metrics. All elements of the computational grid are preconfigured with the latest software from the OSG Virtual Data Toolkit. In addition, they also created a storage element and incorporated a 114-CPU Mac-Intel cluster with OSG. Their computational portal connects to internal RCN resources, allowing SBGrid to accommodate computations submitted from Northeastern University. They also have the ability to redirect computations to the OSG Cloud. External sites can also utilize SBGrid resources.

In order to facilitate phase II of integration in September of 2008 SBGrid-RCN established a joint RCN-OSG taskforce. The aim of this initiative was twofold: a) to rapidly resolve remaining configuration issues and b) facilitate refinement of existing OSG documentations and procedures. The task force was deemed successful, with all technical issues resolved by November. The task force was closed in December 2008.

In phase II of the project SBGRID-RCN successfully utilized extensive external resources for structural biology computations. Most jobs have been submitted to the UCSD, Wisconsin, and Fermilab. On January 27th 2009 RCN reported a peak utilization of 6,000 hours/day/site.

The RCN has contributed in several ways to OSG operations. Ian Stokes-Rees has worked diligently to ensure that throughout the integration RCN provides a continuous feedback to OSG, and that it works with OSG to improve existing procedures, documentation and Virtual Data Toolkit software. Piotr Sliz (PI of SBGrid) was elected to the OSG Council in March 2009.




Figure 7: Utilization of remote Open Science Grid sites by SBGrid in November, December and January. Peak utilization of 6,000 CPU hours was reported on January 26th 2009.

SBGrid-RCN has been a leading participant in the newly established Biomed HPC Collaborative. The initiative aims to coordinate efforts of High Performance Biomedical Computing groups from Boston area (participants include Beth Israel Deaconess Medical Center, Boston University, Brown University, Dana Farber Cancer Institute, Harvard and several affiliated schools, Northeastern University, Partners Healthcare, The Broad Institute, Tufts University, University of Massachusetts, University of Connecticut Health Center and Wyss Institute for Biologically Inspired Engineering). SBGrid RCN has been providing guidance on Open Science Grid integration, and in collaboration with the OSG has seeded a supporting initiative to interlink existing biomedical resources in the Boston area.

Multi-Disciplinary Sciences: The Engagement team has worked directly with researchers in the areas of: biochemistry (Xu), molecular replacement (PRAGMA), molecular simulation (Schultz), genetics (Wilhelmsen), information retrieval (Blake), economics, mathematical finance (Buttimer), computer science (Feng), industrial engineering (Kurz), and weather modeling (Etherton).

The computational biology team led by Jinbo Xu of the Toyota Technological Institute at Chicago uses the OSG for production simulations on an ongoing basis. Their protein prediction software, RAPTOR, is likely to be one of the top three such programs worldwide.

A chemist from the NYSGrid VO using several thousand CPU hours a day sustained as part of the modeling of virial coefficients of water. During the past six months a collaborative task force between the Structural Biology Grid (computation group at Harvard) and OSG has resulted in porting of their applications to run across multiple sites on the OSG. They are planning to publish science based on production runs over the past few months.

Computer Science Research: A collaboration between OSG extensions program, the Condor project, US ATLAS and US CMS is using the OSG to test new workload and job management

scenarios which provide “just-in-time” scheduling across the OSG sites using “glide-in” methods to schedule a pilot job locally at a site which then requests user jobs for execution as and when resources are available. This includes use of the “GLExec” component, which the pilot jobs use to provide the site with the identity of the end user of a scheduled executable.

3.2.2. Findings of the Distributed Infrastructure: The OSG Facility

OSG Facility: The facility provides the platform that enables production by the science stakeholders; this includes operational capabilities, security, software, integration, and engagement capabilities and support. In the last year, we have increased focus on providing “production” level capabilities that the OSG VOs can rely on for their computing work and get timely support when needed. Maintaining a production facility means paying particular attention to detail and effectively prioritizing the needs of our stakeholders while constantly improving the infrastructure; this is facilitated by the addition of a Production Coordinator (Dan Fraser) to the OSG staff who provides focus specifically on these issues. Other improvements to the platform this year included: (1) attention to software technology that will improve incremental software delivery to sites to minimize disruption of production activities; (2) the addition of new probes into the RSV infrastructure for reporting site capability and availability; (3) a redesign of the ticketing infrastructure that makes it easier to submit and manage tickets; (4) support for new storage technologies such as BeStMan and Xrootd based on stakeholder needs; and (5) new tools needed by ATLAS and CMS for data management.

The stakeholders continue to ramp up their use of OSG, and the ATLAS and CMS VOs are ready for the restart of LHC data taking and being ready to run the anticipated heavy workloads.




Figure 8: OSG facility usage vs. time broken down by VO

In the last year, the usage of OSG resources by VOs has roughly doubled from 2,000,000 hours per week to over 4,000,000 hours per week, sustained; additional detail is provided in attachment 1 entitled “Production on the OSG.” OSG provides an infrastructure that supports a broad scope

of scientific research activities, including the major physics collaborations, nanoscience, biological sciences, applied mathematics, engineering, and computer science. Most of the current usage continues to be in the area of physics but non-physics use of OSG is a growth area with current usage of 195,000 hours per week (averaged over the year) spread over 13 VOs.

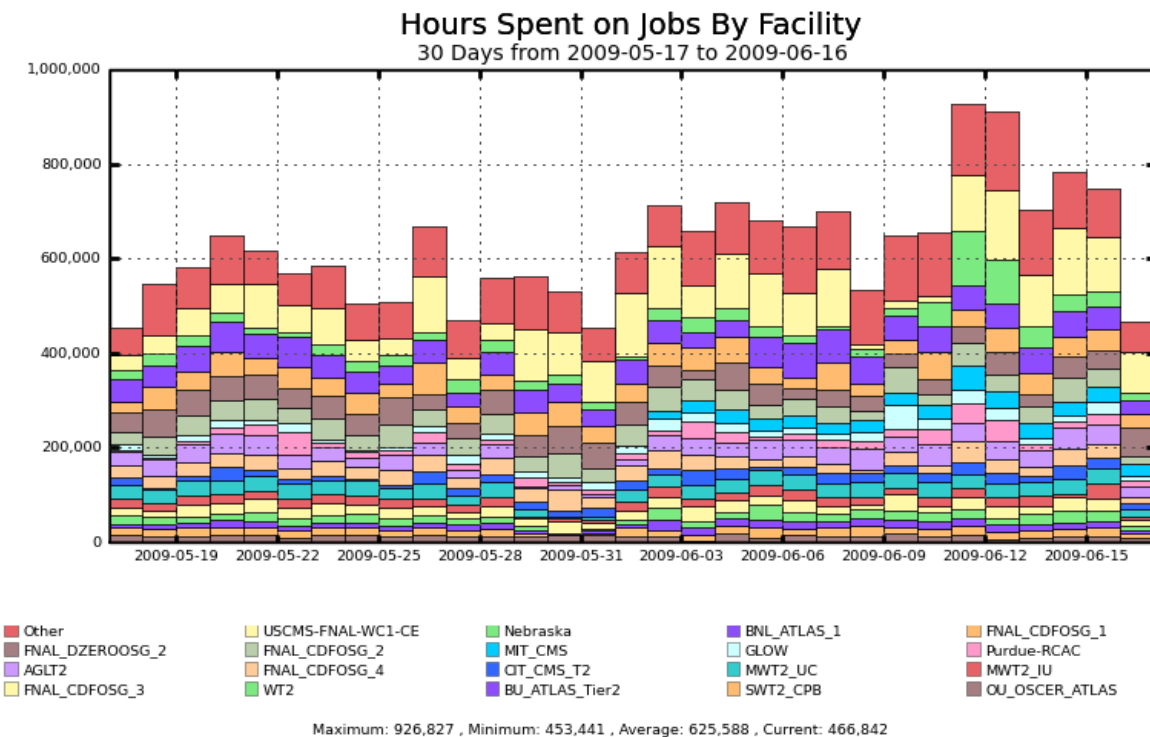


Figure 9: OSG facility usage vs. time broken down by Site. (Other represents the summation of all other “smaller” sites)

With about 80 sites, the production provided on OSG resources continues to grow; the usage varies depending on the needs of the stakeholders. During stable normal operations, OSG provides approximately 600,000 CPU wall clock hours a day with peaks occasionally exceeding 900,000 CPU wall clock hours a day; approximately 150,000 opportunistic wall clock hours are available on a daily basis for resource sharing.

Middleware/Software: To enable a stable and reliable production platform, the middleware/software effort has increased focus on support and capabilities that improve administration, upgrades, and support. Between June 2008 and June 2009, OSG’s software efforts focused on supporting OSG 1.0 and developing OSG 1.2.

As in all major software distributions, significant effort must be given to ongoing support. OSG 1.0 was released in August 2008 and was extensively documented in last year’s annual report. Subsequently, there have been 20 incremental updates to OSG 1.0, which demonstrates that OSG 1.0 was successful in one of our main goals: being able to support incremental updates to the software stack, something that has been traditionally challenging in the OSG software stack.

While most of the software updates to OSG 1.0 were “standard” updates featuring numerous bug fixes, security fixes, and occasional minor feature upgrades, three updates are worthy of deeper discussion. As background, the OSG software stack is based on the VDT grid software distribution. The VDT is grid-agnostic and used by several grid projects including OSG,

TeraGrid, and WLCG. The OSG software stack is the VDT with the addition of OSG-specific configuration.

- 1) VDT 1.10.1i was released in September 2008, and it changed how we ship certificate authority (CA) certificates to users. Instead of the CA certificates coming from a software provider (i.e. the VDT team), they are supplied by the OSG security team. As of early 2009, the VDT team still provides a “convenience” installation of CA certificates that is simply the IGTF-certified CAs, but the OSG security team is responsible for building the CA distribution used by most OGS sites, thus correctly placing responsibility with security experts. In addition, VDT users (most likely from other grids) can now easily provide their own CA distributions as appropriate.
- 2) VDT 1.10.1q was released in December 2008 and represents the culmination of significant efforts of the storage sub-team of the VDT. This release added support for new types of storage elements based on BeStMan (which provides an SRM interface) and Xrootd (which provides a distributed file system). While we continue to support dCache, new storage technologies are a major new focus for OSG and this has required substantial effort to develop our ability to support them. It is important for smaller OSG sites that wish to deploy an SE because it is simpler to install, configure, and maintain than dCache, perhaps at the cost of some scalability and performance. Support for BeStMan with Xrootd was requested by the ATLAS experiment, but is likely to be of interest to other OSG users as well.
- 3) VDT 1.10.1v was a significant new update that stressed our ability to supply a major incremental upgrade without requiring complete re-installations. To do this, we supplied a new update program that assists site administrators with the updating process and ensures that it is done correctly. This updater will be used for all future updates provided by the VDT. The update provided a new version of Globus, an update to our authorization infrastructure, and an update to our information infrastructure. It underwent significant testing both internally and by VOs in our integration testbed.

In the last several months, we have been hard at work at creating OSG 1.2. As much as OSG 1.0 has improved our ability to provide software updates without requiring a fresh installation, there were several imperfections in our ability to do so. The LHC data taking will be restarted at the end of September 2009, and it is imperative that we are able to provide software updates smoothly so that LHC sites can upgrade during data taking. Therefore we have developed a new version of the VDT (2.0.0) that will be the basis for OSG 1.2. As of early June 2009, a pre-release of OSG 1.2 is in testing by the OSG integration testbed, and we expect it to be ready for deployment by the beginning of August 2009, in time for sites to be able to install before the LHC data taking restart.

OSG 1.2 contains very few software upgrades, but has focused instead on improvements to packaging. Because of this, we expect testing to go fairly smoothly. That said, there have been some software upgrades to meet the needs of OSG stakeholders, such as upgrades to MyProxy (for ATLAS) and new network diagnostic tools (requested by ATLAS, but useful to most OSG sites).

In the fall of 2008, we added the Software Tools Group (STG), which watches over the small amount of software development being done in OSG. Although we strongly prefer not to develop software, there are some needs that are not met by sourcing software from external

providers; in these cases, the STG, led by Alain Roy and Mine Altunay, watches over the requirements, development, and release of this software.

A few other notable software developments:

- In November 2008, we held a meeting with external software providers, to improve our communication and processes between OSG and software providers.
- In the spring of 2009, we developed a testbed for improved testing of BeStMan and Xrootd.
- We are preparing for an OSG Storage Forum to be held at the end of June 2009 that will bring together OSG site administrators and storage experts.

The VDT continues to be used by external collaborators. EGEE/WLCG uses portions of VDT (particularly Condor, Globus, UberFTP, and MyProxy). The VDT team maintains close contact with EGEE/WLCG due to the OSG Software Coordinator's (Alain Roy's) weekly attendance at the EGEE Engineering Management Team's phone call. TeraGrid and OSG continue to maintain a base level of interoperability by sharing a code base for Globus, which is a release of Globus, patched for OSG and TeraGrid's needs.

Operations: Operations provides a central point for operational support for the Open Science Grid. The Grid Operation Center (GOC) performs real time monitoring of OSG resources, supports users, developers and system administrators, maintains critical information services, provides incident response, and acts as a communication hub. The primary goals of the OSG Operations group are: supporting and strengthening the autonomous OSG resources, building operational relationships with peering grids, providing reliable grid infrastructure services, ensuring timely action and tracking of operational issues, and quick response to security incidents. In the last year, the GOC continued to provide the OSG with a reliable facility infrastructure while at the same time improving services to offer more robust tools to the stakeholders of the OSG.

The GOC continued to provide and improve numerous stable services for the OSG. The OSG Information Management (OIM) database that provides the definitive source of information about OSG entities at the person, resource, support agency, or virtual organization level was updated to allow new data to be provided to OSG stakeholders, as well as cleaning up the database backend and enhancing the aesthetics. The services have been used to provide operations automation, simplifying and reducing some time consuming administrative tasks as well as providing automated reporting to the WLCG. Operations automation allowed us to be prepared to better handle the needs of the stakeholders during the LHC data-taking. The Resource and Service Validation (RSV) monitoring tool is going through a second round of updates improving stability and allowing new security and administrator use functionality. Redundant BDII (Berkeley Database Information Index) servers, requested by US CMS, are now in place in Bloomington and Indianapolis, allowing us to provide a BDII data survivability with load-balancing and failover. MyOSG is an information consolidating tool and is being deployed, allowing customizable "dashboards" to be created by OSG users and administrators based on their own specific needs. MyOSG allows administrative, monitoring, information, validation and accounting services to be displayed at a single address. A public interface to view trouble tickets that the GOC is working is now available. This interface allows issues to be tracked and updated by user and it also allows GOC personnel to use OIM meta-data to route tickets much

more quickly, reducing the amount of time needed to look up contact information of resources and support agencies. Several other hardware and service upgrades have taken place:

- The TWiki environment used for collaborative documentation was updated with new functionality and with security fixes.
- The BDII was updated to improve performance.
- The power and networking infrastructure in the racks holding the servers providing the OSG services was enhanced.
- A migration to a virtual machine environment for many services is being undertaken to allow flexibility in providing high availability services.

OSG Operations is currently preparing to support the LHC start-up, in addition to focusing on service reliability and operations automation. We are actively preparing for the stress of the LHC start-up on services by testing, by putting proper failover and load-balancing mechanisms in place, and by implementing administrative ticketing automation. Service reliability for GOC services has always been high and we will begin gathering metrics that can show the reliability of these services exceed the requirements of Service Level Agreements (SLAs) that will be agreed to with the OSG stakeholders. The first SLA was written and agreed to for the CMS use of the BDII; a list of needed SLAs has been documented. Operations automation is important to permit the GOC work to be scalable into the future and we will conduct more research into the best ways to allow process automation and problem alerts that will allow us to keep up with the growth of OSG.

Integration and Site Coordination: The mission of the OSG integration activity is to improve the quality of grid software releases deployed on the OSG and enable greater success by the sites in achieving effective production.

In the last year, the Integration effort delivered high quality software packages to our stakeholders resulting in smooth implementation of the OSG 1.0 and its update to OSG 1.0.1 releases; several process innovations were key to these results. During the release transition to OSG 1.0, several iterations of the Validation Test Bed (VTB) were made using a 3-site test bed which permitted quick testing of pre-release VDT updates, functional tests, and install and configuration scripts. The ITB was deployed on 12 sites providing compute elements and four sites providing storage elements (dCache and BeStMan packages implementing SRM v1.1 and v2.2 protocols); 36 validation processes were defined across these compute and storage resources in readiness for the production release. Pre-deployment validation of applications from 12 VOs were coordinated with the OSG VOs support group. Other accomplishments include both dCache and SRM-BeStMan storage element testing on the ITB; delivery of a new site configuration tool; and testing of an Xrootd distributed storage system as delivered by the OSG Storage group.

The OSG Release Documentation continues to receive significant edits from the community of OSG participants. The collection of wiki-based documents capture processes for install, configure, and validation methods as used throughout the integration and deployment processes. These documents were updated and received review input from all corners of the OSG community (33 members participated for the OSG 1.0 release) resulting in a higher quality output. A new initiative has been launched to align site administrator's documentation with other groups in OSG to promote re-use and consistency.

The community of resource providers comprising the OSG Facility is diverse in terms of the scale of computing resources in operation, research mission, organizational affiliation, and technical expertise, leading to a wide range of operational performance. The Sites Coordination activity held two face-to-face workshops (a dedicated meeting at SLAC, a second co-located with the OSG All Hands meeting at the LIGO observatory). Both of these were hands-on covering several technical areas for both new and advanced OSG administrators.

Virtual Organizations Group: A key objective in OSG is to facilitate, enable, and sustain Science communities to *produce* Science using the OSG Facility. To accomplish this goal, the Virtual Organizations Group (VO Group) directly interfaces with each VO to address requirements, feedback, issues, and roadmaps for production-scale operations of the “at-large” (i.e. all VOs except ATLAS, CMS, LIGO which are directly supported by the OSG Executive Team) Science communities.

The focus is to: (a) improve efficiency and utilization of OSG Facility; (b) provide an avenue for operational, organizational, and scientific discussions with each at-large stakeholder; (c) facilitate broad stakeholder participation in the OSG software engineering lifecycle; (d) enable tactical methods for sustenance of communities that have a newly formed VO; and (e) provide a channel for the OSG Storage group to work directly with all stakeholders, and thus strengthen the data-grid capabilities of OSG. Some of the major work items in the last year were:

- Feedback from most of the science communities to the OSG team was completed to improve planning for their needs. Input was gathered from 17 at-large VOs covering: scope of use; VO mission; average and peak utilization of OSG; resource provisioning to OSG; and plans, needs, milestones. This information was reported to the OSG Council on behalf of ALICE, CDF, CompBioGrid, D0, DES, DOSAR, Fermilab VO, GEANT4, GPN, GRASE, GROW, GUGrid, IceCube, MARIACHI, nanoHUB, NYSGrid, and SBGrid.
- Pre-release Science Validation on the Integration Testbed (ITB) was completed for OSG Release1.0, and its incremental updates. In partnership with OSG Integration, a rigorous OSG process has been designed and is regularly executed prior to each software release to assure quality. Each participating Science stakeholder tests their own use scenarios, suggesting changes, and signaling an official approval of each major OSG release. In ITB 0.9.1 validation, 12 VOs participated, 7 VOs ran real Science applications, 6 VOs participated in storage validation, of which, 4 VOs conducted introductory validation of opportunistic storage. In terms of process execution, this was a coalition of 36+ experts, 20+ from VO communities. After careful validation and feedback, official ‘green flags’ toward OSG 1.0 were given by ATLAS, CDF, CIGI, CMS, DES, DOSAR, Dzero, Engagement, Fermilab VO, LIGO, nanoHUB, SBGrid, and SDSS. Subsequently as part of ITB 0.9.2, a smaller-scale cycle was organized for the incremental Release1.0.1.
- Joint Taskforces were executed for ALICE, D0, nanoHUB, and SBGrid. Via joint staffing and planning between OSG and the collaborations, we addressed wide-ranging technical and process items that enabled production use of OSG by the VOs. During the last year: (1) the ALICE-OSG Taskforce integrated LHC AliEn grid paradigm to startup ALICE production on OSG, using the current scale of ALICE resources in the US. (2) the D0-OSG Taskforce led to a significant improvement in D0’s procedures, D0’s grid infrastructure, and in the overall D0 monte-carlo event production on OSG. In part due to this work, D0 has continued to reach new levels of Monte Carlo production; in May 2009, D0 reached a new peak of 13

million events per week, (3) the SBGrid-OSG Taskforce worked closely together to enable SBGrid resource infrastructure and to evolve design and implementation of the SBGrid Molecular Replacement science application, (4) the nanoHUB-OSG Taskforce successfully made gradual improvements in one another's infrastructure to increase nanoHUB production volume and job efficiency across OSG, and (5) the Geant4-OSG Task Force, currently active, is working to enable Geant4's Regression Testing production runs on the OSG Facility.

- Production-scale Opportunistic Storage provisioning and usage was initiated on OSG. In partnership with the OSG Storage group, a technical model was designed and enabled on select SRM storage sites of CMS and ATLAS, followed by its sustained active usage by D0.
- The Annual OSG Users meeting was organized at BNL in June 2008, with emphasis on VO security and policy.

The VO Group continues to provide bidirectional channels between Science communities and all facets of the OSG, to assure that the *needs and expectations of Science* communities are understood, absorbed, and translated into work activities and decisions in OSG.

Engagement: A major priority of Open Science Grid is helping new science communities benefit from the infrastructure we are putting in place by working closely with these communities over periods of several months. The Engagement activity brings the power of the OSG infrastructure to scientists and educators beyond high-energy physics and uses the experiences gained from working with new communities to drive requirements for the natural evolution of OSG. To meet these goals, engagement helps in: providing an understanding of how to use the distributed infrastructure; adapting applications to run effectively on OSG sites; engaging the deployment of community owned distributed infrastructures; working with the OSG Facility to ensure the needs of the new community are met; providing common tools and services in support of the engagement communities; and working directly with and in support of the new end users with the goal to have them transition to be full contributing members of the OSG. These goals and methods remain the same as they have been in previous years.

During this program year, the Engagement team has successfully worked with the following researchers who are in full production use of the Open Science Grid, including: Steffen Bass (+3), theoretical physics, Duke University; Anton Betten, mathematics, Colorado State; Jinbo Xu (+1), protein structure prediction, Toyota Technological Institute; Vishagan Ratnaswamy, mechanical engineering, New Jersey Institute of Technology; Abishek Patrap (+2), systems biology, Institute for Systems Biology; Damian Alvarez Paggi, molecular simulation, Universidad de Buenos Aires; Eric Delwart, metagenomics, UCSF; Tai Boon Tan, molecular simulation, SUNY Buffalo; Blair Bethwaite (+1), PRAGMA. Additionally, we have worked closely with the following researchers who we expect will soon become production users: Cynthia Hays, WRF, University of Nebraska-Lincoln; Weitao Wang (+2), computational chemistry, Duke University; Kelly Fallon, The Genome Center at Washington University. Figure 10 shows the diversity and level of activity among Engagement users for the previous year, and Figure 11 shows the distribution by OSG facility of the roughly 3 million CPU hours that Engagement users have consumed during that same time frame.

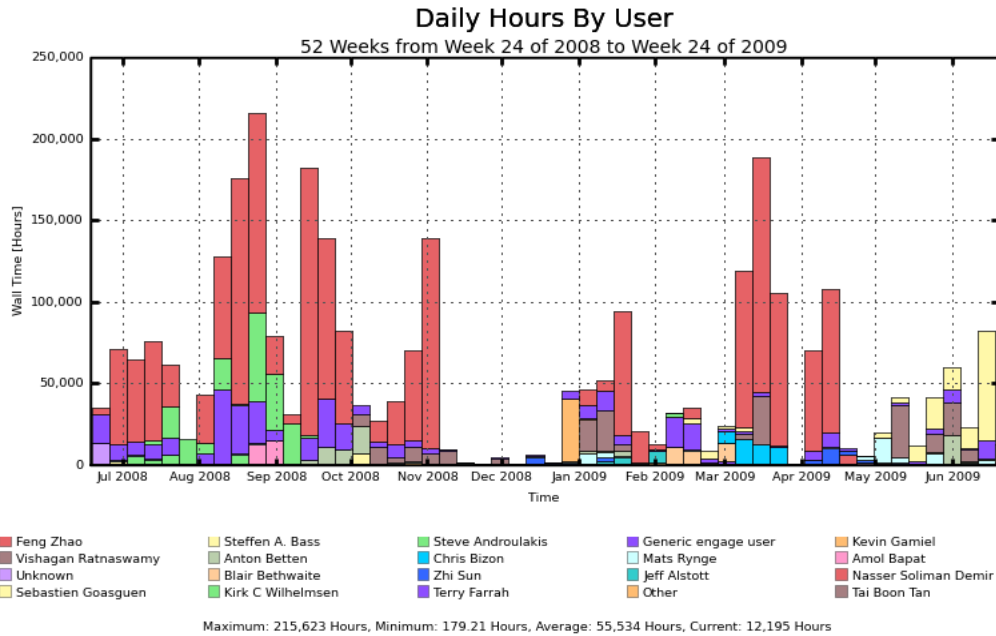


Figure 10: Engage user activity for one year

In addition to developing the new production users, the Engagement Team has added a compute element from RENCI which is providing on order of 4k cpu hours per day to the Engagement VO as well as other VOs such as LIGO and nanoHUB. We have assisting the SB-Grid engagement effort (Peter Doherty), initiated discussions with two research teams regarding MPI jobs (Cactus, SCEC), and have begun exporting the Engagement methodology to the separately funded activities of RENCI's TeraGrid Science Gateway program as described in a TG'09 paper.

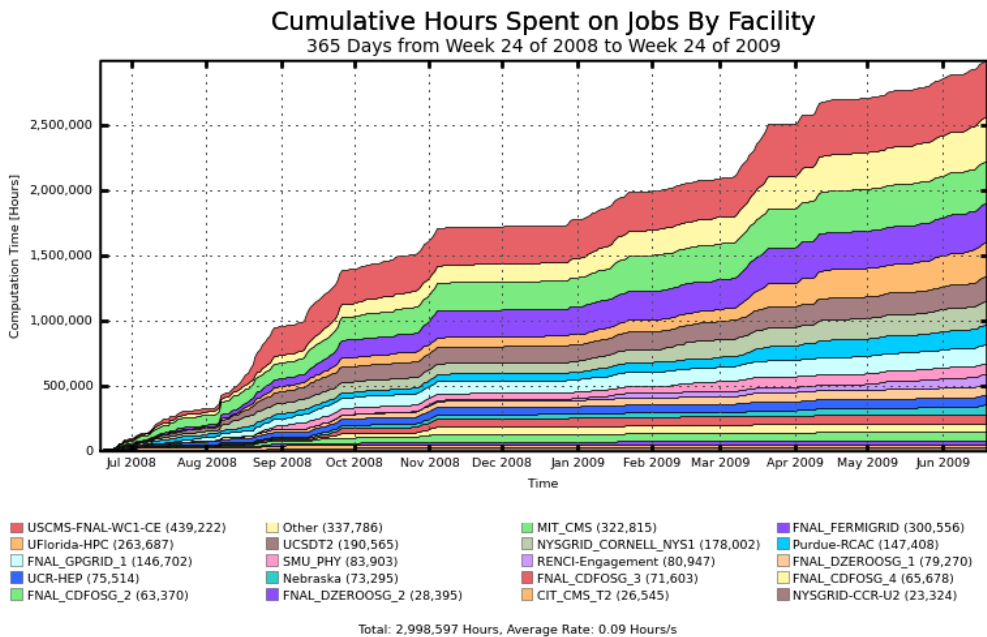


Figure 11: CPU hours by facility for Engage Users

Campus Grids: The Campus Grids team's goal is to include most US universities in the national cyberinfrastructure. By helping universities understand the value of campus grids and resource sharing through the OSG national framework, this initiative aims at democratizing cyberinfrastructures by providing all resources to users and doing so in a collaborative manner.

In the last year, the campus grids team worked closely with the OSG education team to provide training and outreach opportunities to new campuses. A workshop for new site administrators was organized at Clemson and was attended by representatives from four new campuses: University of Alabama, University of South Carolina, Florida International Miami and from the Saint Louis Genome Center. These four sites are currently in active deployment or planning to join OSG and contribute resources. A second workshop was held at RENCi, and included representatives from Duke, UNC-Chapel Hill, and NCSU among others. UNC-CH has recently launched its TarHeel Grid, and Duke University is actively implementing a campus grid being incubated by a partnership between academic computing and the physics department. Overall 24 sites have been in various levels of contacts with the campus grids team and seven are now in active deployment. Finally, a collaborative activity was started with SURAGrid to study interoperability of the grids; and an international collaboration with the White Rose Grid from the UK National Grid Service in Leeds as well as the UK Interest Group on Campus Grid was initiated this period.

As a follow-up to work from the prior year, the Clemson Windows grid contributed significant compute cycles to Einstein@home run by LIGO and reached the rank of the 4th highest contributor world-wide; work is currently in-progress to report this under the OSG production metrics via the Gratia accounting system. Finally as a teaching activity, five Clemson students (of which two are funded by OSG and two are funded by CI-TEAM) will compete in the TeraGrid 2009 parallel computing competition. One of them, an African-American female, will attend the ISSGC 2009 in France due to OSG financial support.

Security: The Security team continued its multi-faceted approach to successfully meeting the primary goal of maintaining operational security, developing security policies, acquiring, or developing necessary security tools and software, and disseminating security knowledge and awareness.

We continued our efforts to improve the OSG security program. The security officer asked for an external peer evaluation of the OSG security program which was conducted by the EGEE and TeraGrid security officers and a senior security person from University of Wisconsin. The committee produced a report of their conclusions and suggestions, according to which we are currently adjusting our current work. This peer evaluation proved so successful that we have received requests from EGEE and TeraGrid officers to join an evaluation of their security programs.

As part of our operational tests, we conducted a security drill over OSG Tier-1 sites in cooperation with the EGEE security team. Both teams conducted the same drill scenario and evaluated the sites on similar grading schemes. The drill was useful for both the OSG sites and the security team. We obtained a detailed knowledge of EGEE's operational guidelines and were able to measure sites' performance across different middleware and security procedures. Our sites performed excellently and both the OSG Tier-1 sites scored the maximum allowed 100 points plus bonus points.

While preparing for the LHC date taking restart, we evaluated our infrastructure against potential security related disruptions. We conducted a risk assessment of OSG infrastructure and devoted a 2-day meeting with OSG executive team to discuss our findings. For identified threats, we have already started contingency planning and we are currently implementing the plan. We also identified key providers for OSG infrastructure and included them in our contingency planning. DOEGrids CA is the first of such providers, and we worked with them to produce the contingency plan. This process was very useful and both sides gained new insights over potential risks and we plan to work with other key providers subsequently.

We continued the day-to-day work on incident response and monitoring of the OSG infrastructure. During the last year, we had a single incident that required considerable effort. The incident was tangential to the grid middleware; it exploited vulnerabilities found in the ssh protocol and affected some OSG sites. However, we did not have any grid incident as a result of the ssh incident and our actions were preventive in nature. A positive consequence of the incident is that TeraGrid, OSG, and EGEE security officers have started a joint incident sharing community.

We continued to examine the day-to-day behavior of OSG members to gain a better understanding of their needs. Although in its early stages, this effort already helped us uncover a few potential problems with our infrastructure. We realized that sites are reluctant in updating their security configurations to OSG-suggested values. An examination revealed that an administrative tool that can merge the updated values without overwriting site-specific variables can solve the problem. We developed this tool and plan to release it in OSG 2.0. To provide better monitoring, the security team wrote a suite of security probes which allow a site administrator to compare local security configurations against OSG-suggested values. We plan to include the probes in OSG 2.0 release.

Finally, another joint middleware project, Authorization Interoperability project, has been completed. The project aimed to achieve message-level interoperability between EGEE and OSG's security components. The software has been released and tested thoroughly.

As part of our efforts in security policies and procedures, we continued the OSG representation in Joint Security Policy Group (JSPG), which is the main body for preparing and suggesting security policies to WLCG. We have provided regular feedback to US sites regarding policy changes and their potential impacts over the sites. In addition, we have completed the Software Vulnerability Procedure, Privacy Policy, and Certificate Authorities Policy, received approvals and started enforcing these policies and procedures.

Metrics and Measurements: The metrics and measurements activity aims to give the OSG, VOs, and external agencies a view of the OSG's progress. The metrics for FY08 were completed and published. The formal reporting activity in this area was altered slightly in FY09 to focus on receiving input from all areas of the OSG. This effort, called "internal metrics" has involved all OSG Area Coordinators and produced an outline of each area's goals; we are now in the process of taking measurements for these metrics. Metrics and measurements also provide OSG Management with ad-hoc reports on an as-needed basis.

In addition to reports, there are continuous and monthly activities supported by the metrics and measurements activity. The area maintains a repository of historical data from the OSG information services; this allows the OSG to track the number of active running jobs throughout the year, as well as track the number of cores deployed and the types of CPU on the distributed

facility. This is a continuation of a Year 2 effort, and we now have a large set of historical data on batch system activities on the grid. In this year, we started the roll-out of transfer accounting to many WLCG sites; while this effort is not complete, we believe a large percentage of OSG transfers are now being accounted for. We are in the process of trying to integrate this information – usually only accessible by experts – into GOC displays, making it available to more users.

The Metrics and Measurements area continues to be involved with the coordination of WLCG-related reporting efforts. It sends monthly reports to the JOT that highlights MOU sites' monthly availability and activity. It produces monthly data for the metric thumbnails and other graphs available on the OSG homepage. Current projects include coordinating the reporting of WLCG Installed Capacity data, development of storage reporting, and rolling out transfer accounting. The accounting extensions effort requires close coordination with the Gratia project at FNAL, and includes effort contributed by ATLAS.

The Metrics and Measurements area will also begin investigating how to incorporate the network performance monitoring data into its existing reporting activities. Several VOs are actively deploying perfSONAR based servers that provide site-to-site monitoring of the network infrastructure, and we would like to further encourage network measurements on the OSG.

Based on input from the stakeholders, OSG assessed the various areas of contribution and documented the value delivered by OSG. The goal of this activity was to develop and estimate the benefit and cost effectiveness, and thus provide a basis for discussion of the value of the Open Science Grid (OSG).

Findings of the Distributed Infrastructure: Extending Science Applications

In addition to operating a facility, the OSG includes a program of work that extends the support of Science Applications both in terms of the complexity as well as the scale of the applications that can be effectively run on the infrastructure. We solicit input from the scientific user community both as it concerns operational experience with the deployed infrastructure, as well as extensions to the functionality of that infrastructure. We identify limitations, and address those with our stakeholders in the science community. In the last year of work, the high level focus has been threefold: (1) improve the usability and scalability, as well as our understanding thereof; (2) establish and operate a workload management system for OSG operated VOs; and (3) establish the capability to use storage in an opportunistic fashion at sites on OSG.

In the present year, we made a change in the way we track the needs of our primary stakeholders: ATLAS, CMS, and LIGO. We established the notion of a “senior account manager” for each of the three. That person then met on a quarterly basis with senior management of the stakeholders to go over their needs. Additionally for ATLAS and CMS, we started documenting, and revising their feedback in form of a prioritized “wishlist” with deliverable dates. This quarterly updated list then informed much of the work in the extensions program.

Scalability, Reliability, and Usability: As the scale of the hardware that is accessible via the OSG increases, we need to continuously assure that the performance of the middleware is adequate to meet the demands. There were four major goals in this area for the last year and they were met via a close collaboration between developers, user communities, and OSG.

- At the job submission client level, the goal is 20,000 jobs running simultaneously and 200,000 jobs run per day from a single client installation, and achieving in excess of 95%

success rate while doing so. The job submission client goals were met in collaboration with CMS, CDF, Condor, and DISUN, using glideinWMS. This was done via a mix of controlled environment and large scale challenge operations across the entirety of the WLCG. For the controlled environment tests, we developed an “overlay grid” for large scale testing on top of the production infrastructure. This test infrastructure provides in excess of 20,000 batch slots across a handful of OSG sites. An initial large-scale challenge operation was done in the context of CCRC08, the main LHC computing challenge in May 2008. Here we submitted a typical CMS application to 40 sites distributed worldwide, at a scale of up to 4000 simultaneously running jobs. Condor scalability limitations across large latency networks were discovered and this led to substantial redesign and reimplementations of core Condor components, and subsequent successful scalability testing with a CDF client installation in Italy submitting to the CMS server test infrastructure on OSG. Testing with this testbed exceeded the scalability goal of 20,000 running simultaneously and 200,000 jobs per day across the Atlantic. This paved the way for production operations to start in CMS across the 7 Tier-1 centers. Data Analysis Operations across the roughly 50 Tier-2 and Tier-3 centers available worldwide today is more challenging, as expected, due to the much more heterogeneous level of support at those centers. During STEP09 data analysis at the level of 6000 to 10000 jobs was sustained for a two week period via a single glideinWMS instance serving close to 50 sites worldwide. The goal of this exercise was to determine whether or not the totality of the pledged resources on the global CMS grid could be utilized. As of this writing, STEP09 is still ongoing and OSG participates with expertise and effort in this exercise.

- At the storage scheduling level, the present goal was to have 2Hz file handling rates. We expected that to be sufficient given that it translates into more than 10Gbps for Gbyte file sizes. An SRM scalability of 5Hz was achieved in collaboration with the dCache developers, and demonstrated at the CMS Tier-1 center at FNAL. While we clearly exceeded our goals, we also realized that significant further increases are needed in order to cope with the increasing scale of operations by the large LHC VOs Atlas and CMS. The driver here is stage-out of files produced during data analysis. The large VOs find that the dominant source of failure in data analysis is the stage-out of the results, followed by read-access problems as second most likely failure. This has led to a resetting of the goal to a much more aggressive 50Hz srmls and srmcp. In addition, storage reliability is receiving much more attention now, given its measured impact on job success rate. In a way, the impact of jobs on storage has become more and more of a visible issue in part because of the large improvements in the submission tools, monitoring, and error accounting within the last two years. The improvements in submission tools coupled with the increased scale of resources available is driving up the load on storage. The improvements in monitoring and error accounting are allowing us to fully identify the sources of errors. OSG is very actively engaged in understanding the issues involved, working with both the major stakeholders as well as partner grids.
- At the functionality level, this year’s goal was to roll out the capability of opportunistic space use. The roll-out of opportunistic storage was exercised on the OSG ITB preceding OSG v1.0. It has since been deployed at several dCache sites on OSG, and successfully used by D0 for the production operations on OSG. CDF is presently in the testing stage for adapting

opportunistic storage into their production operations on OSG. However, a lot of work is left to do before opportunistic storage is an easy to use and widely available capability of OSG.

- OSG has successfully transitioned to a “bridge model” with regard to WLCG for its information, accounting, and availability assessment systems. This implies that there are aggregation points at the OSG GOC via which all of these systems propagate information about the entirety of OSG to WLCG. For the information system this implies a single point of failure, the BDII at the OSG GOC. If this service fails then all resources on OSG disappear from view. ATLAS and CMS have chosen different ways of dealing with this. While ATLAS maintains its own “cached” copy of the information inside Panda, CMS depends on the WLCG information system. To understand the impact of the CMS choice, OSG has done scalability testing of the BDII. We find that the service is reliable up to a query rate of 10 Hz. The OSG GOC is deploying monitoring of the query rate of the production BDII in response to this finding. The goal is for the GOC to monitor this rate in order to understand the operational risk implied by this single point of failure.

In addition, we have worked on a number of lower priority objectives:

- On WS-GRAM scalability and reliability in collaboration with LIGO, DISUN, CDIGS/Globus, and OSG. As Globus is transitioning to GRAM5, we are committed to work with them on large scale testing of GRAM5 as it evolves.
- On testing of a Condor client interface to the CREAM compute element (CE) in support of ATLAS and CMS. CREAM is a webservices based CE developed by INFN in EGEE. WLCG sites in Europe and Asia are considering replacing their Globus GRAM with CREAM on some yet to be determined timescale. Both ATLAS and CMS require Condor client interfaces to talk to CREAM. We have successfully completed the first two of three phases of testing and we are now waiting for CREAM deployment on the production infrastructure on EGEE to allow for large scale testing.
- On testing the Condor client interface to the ARC compute element (CE) in support of ATLAS and CMS; ARC is the middleware deployed on NorduGrid. The situation here is similar to CREAM, except that ARC CEs are already deployed on the NorduGrid production sites.
- In the area of usability, an “operations toolkit” for dCache was started. The intent was to provide a “clearing house” of operations tools that have been developed at experienced dCache installations, and derive from that experience a set of tools for all dCache installations supported by OSG. This is significantly decreasing the cost of operations and has lowered the threshold of entry. Site administrators from both the US and Europe have uploaded tools, and the first two releases were derived from that. These releases have been downloaded by a number of sites, and are in regular use across the US, as well as some European sites.
- Work has started on putting together a set of procedures that would allow us to automate scalability and robustness tests of a Compute Element. The intent is to be able to quickly “certify” the performance characteristics of new middleware, a new site, or deployment on new hardware. Once we have such procedures, we can then offer this as a service to our resource providers so that they can assess the performance of their deployed or soon to be deployed infrastructure.

Workload Management System: The primary goal of the OSG Workload Management System (WMS) effort is to build, integrate, test and support operation of a flexible set of software tools and services for efficient and secure distribution of workload among OSG sites. There are currently two suites of software utilized for that purpose within OSG: Panda and glideinWMS, both drawing heavily on Condor software.

The Panda system continued as a supported WMS service for the Open Science Grid, and a crucial infrastructure element of ATLAS experiment at LHC. We completed the migration of Panda software to Oracle database backend, which enjoys strong support from major OSG stakeholders and allows us to host an instance of the Panda server at CERN where ATLAS is located, creating efficiencies in support and operations areas.

To foster wider adoption of Panda in the OSG user community, we created a prototype of data service that will make easier its utilization by individual users, by providing a Web-based user interface for uploading and management of input and output data, and a secure backend that allows Panda pilot jobs to both download and transmit data as required by the Panda workflow. No additional software is required on users' desktop PCs or lab computers, and this will be helpful for smaller research groups, who may lack manpower to support the full software stack.

Progress was made with the glideinWMS system approaching the project goal of pilot-based large-scale workload management. Version 2.0 has been released and is capable of servicing multiple virtual organizations with a single deployment. The FermiGrid facility has expressed interest in putting this in service for VOs based at Fermilab. Experiments such as CMS, CDF and MINOS are currently using glideinWMS in their production activities. Discussions are underway with new potential adopters including DZero and CompBioGrid. We also continued the maintenance of the gLExec (user ID management software), a collaborative effort with EGEE, as a project responsibility.

In the area of WMS security enhancements, we completed integration of gLExec into Panda. It is also actively used in glideinWMS. In addition to giving the system more flexibility from security and authorization standpoint, this also allows us to maintain a high level of interoperability of the OSG workload management software with our WLCG collaborators in Europe, by following a common set of policies and using compatible tools, thus enabling both Panda and glideinWMS to operate transparently in both domains. An important part of this activity was integration test of a new suite of user authorization and management software (SCAS) developed by WLCG, which involved testing upgrades of gLExec and its interaction with site infrastructure.

Work continued on the development of Grid User Management System (GUMS) for OSG. This is an identity mapping service which allows sites operate on the Grid while relying on traditional methods of user authentication, such as UNIX accounts or Kerberos. Based on our experience with GUMS in production since 2004, a number of new features have been added which enhance its usefulness for OSG.

This program of work continues to be important for the science community and OSG for several reasons. First, having a reliable WMS is a crucial requirement for a science project involving large scale distributed computing which processes vast amounts of data. A few of the OSG key stakeholders, in particular LHC experiments ATLAS and CMS, fall squarely in that category, and the Workload Management Systems developed and maintained by OSG serve as a key enabling factor for these communities. Second, drawing new entrants to OSG will provide benefit of access to opportunistic resources to organizations that otherwise wouldn't be able to

achieve their research goals. As more improvements are made to the system, Panda will be in a position to serve a wider spectrum of science disciplines.

Storage Extensions: The Storage Extensions area contributes to the enhancement of software used in Storage Elements (SEs) on the Open Science Grid, and software used to discover, reserve, and access those Storage Elements. This includes additional features needed by users and sites, as well as improvements to the robustness and ease-of-use of middleware components. This year we have worked on the architecture, requirements and design of a framework and tools for supporting the opportunistic use of Storage Elements on the Open Science Grid.

The storage discovery prototype tools built this year and in the hands of testers include client and server sides of a command-line discovery tool that searches for matches to user requirements among XML descriptions of resources. The XML is created, when necessary, by translation from LDIF formatted data accessed through LDAP. Role-based access control is provided through gPlazma in order to match behavior of SEs, which also use gPlazma.

We are coordinating with and informing storage activities outside of OSG. These include efforts to improve and regularize logging from the various SRM implementations and the consolidated monitoring of the MCAS project.

Internet2 Joint Activities: Internet2 partnered with OSG to develop and test a suite of tools and services that would make it easier for OSG sites to support their widely distributed user community. A second goal is to leverage the work within OSG to create scalable solutions that will benefit the entire Internet2 membership.

Identifying and resolving performance problems continues to be a major challenge for OSG site administrators. A complication in resolving these problems is that lower than expected performance can be caused by problems in the network infrastructure, the host configuration, or the application behavior. Advanced tools that can quickly isolate which problem(s) exist will go a long way toward improving the grid user experience and making grids more useful to more scientific communities.

In the past year, Internet2 has worked with the OSG software developers to incorporate several advanced network diagnostic tools into the VDT package. These client programs interact with perfSONAR-based servers, described below, to allow on-demand testing of poorly performing sites. By enabling OSG site administrators and end users to test any individual compute or storage element in the OSG environment, we can reduce the time it takes to begin the network troubleshooting process. It will also allow site administrators or users to quickly determine if a performance problem due to the network, a host configuration issue, or an application behavior.

In addition to deploying client tools via the VDT, Internet2 staff, working with partner networks in the US and internationally, have created a simple live-CD distribution mechanism for the server side of these tools. This bootable CD allows an OSG site-admin to quickly stand up a perfSONAR-based server to support the OSG users. These perfSONAR boxes automatically register their existence in a global database, making it easy to find new servers as they become available. Internet2 staff also identified an affordable 1U rack mountable computer that can be used to run this server software. OSG site administrators can now order this standard hardware, ensuring that they can quickly get started with a known good operating environment.

These servers provide two important functions for the OSG site-admin. First they provide an end point for the client tools deployed via the VDT package. OSG users and site-admins can run on-

demand tests to begin troubleshooting performance problems. The second function they perform is to host regularly scheduled tests between peer sites. This allows a site to continuously monitor the network performance between itself and the peer sites of interest. The USATLAS community has begun monitoring throughput between the Tier1 and Tier2 sites. Finally, on-demand testing and regular monitoring can be performed to both peer sites and the Internet2 or ESNNet backbone network using either the client tools, or the perfSONAR servers. Internet2 will continue to interact with the OSG admin community to learn ways to improve this distribution mechanism.

Another major task for Internet2 is to provide training on the installation and use of these tools and services. In the past year Internet2 has participated in several OSG site-admin workshops, the annual OSG all-hands meeting, and interacts directly with the LHC community to determine how the tools are being used and what improvements are required. Internet2 has provided hands-on training in the use of the client tools, including the command syntax and interpreting the test results. Internet2 has also provided training in the setup and configuration of the perfSONAR server, allowing site-admins to quickly bring up their server. Finally, Internet2 staff has participated in several troubleshooting exercises. This effort includes running tests, interpreting the test results and guiding the OSG site-admin through the troubleshooting process.

National Lambda Rail Activities: National LambdaRail (NLR) continues to provide network connectivity and services in support of the OSG community. Recently NLR was selected to provide ultra high-performance, fiber-optic circuits as part of the network infrastructure to support the Large Hadron Collider (LHC) in the U.S. NLR will provide two 10 Gb/s circuits between Chicago and New York, enabling LHC data access and exchange by the U.S. Tier-1 facilities. These two links are scheduled to be installed and operational by mid-July 2009.

NLR's Layer3 (PacketNet) and Layer2 (FrameNet) services continue to provide basic network connectivity between OSG sites. User can now configure their own Dynamic VLAN System (DVS) on the NLR Sherpa tool. This tool allows users to provision, modify, enable, and disable dedicated or non-dedicated VLANs on FrameNet in real time, without requiring intervention from the NLR NOC. Based on user input, NLR has added additional features to the Sherpa tools, including the ability to schedule dedicated and non-dedicated VLANs for specific time periods. A Sherpa client module is available that allows a researcher to programmatically interact with the Dynamic VLAN system from a remote host.

NLR's PacketNet infrastructure continues to provide good end-to-end connectivity for OSG sites. The default bandwidth for NLR connections is 10-Gigabit Ethernet (GE), thus enabling near 10-GE data streams end-to-end for sites with internal and regional 10-GE infrastructure. The NLRView infrastructure provides a number of test points to help troubleshoot end-to-end performance issues on networks and end systems. Two of the tools provided that allow end users to troubleshoot and measure performance are NPAD and NDT. These tests run on the NLRView Performance Test PC located in the NLR PacketNet Pops. Each of these PCs has 10-GE NICs directly connected to the network.

3.3. Training and Development

Training and outreach to campus organizations, and the development of the next generation of computational scientists is a core part of the OSG program. The OSG Education and Training program brings domain scientists and computer scientists together to provide a rich training

ground for the engagement of students, faculty and researchers in learning the OSG infrastructure, applying it to their discipline and contributing to its development.

During the last year, OSG has sponsored and conducted numerous training events for students and faculty and participated in additional invited workshops. The result was that grid computing training and education reached about 200 new professionals, of which about 30% were from under-represented groups and about 20% were women. Four major OSG sponsored “Grid School” training events were organized and delivered in the past 12 months: 1) Midwest Grid School, Chicago, September 2008 (3 days); 2) Site administrator workshop, Clemson University, March 2009 (1 day); 3) New Mexico Grid School, April 2009, and 4) North Carolina Grid School, April 2009 (2 days). In line with the overall OSG outreach goals, we have reached minority serving and under-resourced institutions by providing training for students and faculty in the state of New Mexico (participants from UNM campuses and Navajo College). Following each of these workshops, feedback was gathered in the form of online surveys, recorded and analyzed, and follow-on research and engagement assistance opportunities were planned and are being provided (including access to the hands-on curricula and to ongoing support offered by the OSG EOT staff).

In addition to these major training events, OSG staff conducted numerous smaller training and outreach workshops (for example, GHC08 and SC08 in Oct and Nov respectively) and is a co-organizer of the International Summer School on Grid Computing (ISSGC09) school that will be held in July 2009 in France with attendance from many countries. The OSG Education team coordinated the selection of students who participated will participate in ISSGC09, and arranges sponsorship for US-based students to attend this workshop. In addition, OSG staff provides direct contributions to the International Grid School by attending, presenting, and being involved in lab exercise development and student engagement. Another aspect of the EGEE-OSG collaboration at the education level involves the participation in the International Winter School on Grid Computing (IWSGC09), an online training event spanning 6 weeks in March-April 2009.

The content of the training material has been enhanced over the past year to include an updated module targeting new OSG Grid site administrators. Work continues to enhance the on-line delivery of user training and we are seeing an increased number of students and faculty that are signing up for the self-paced online course (about 50 active participants). In addition to individuals, we have made the course accessible as support material for graduate courses on grid computing at universities around the country, such as Rochester Institute of Technology, University of Missouri at St. Louis and Clemson University.

OSG collaborates with Educause, Internet2 and TeraGrid to sponsor day-long workshops local to university campus-wide CyberInfrastructure (CI) Days. These workshops bring expertise to the campuses to foster research and teaching faculty development, IT facility planning, and CIO awareness and dialog. Ongoing dialog and coordination between the EOT programs of TeraGrid, OSG and the Supercomputing conference education programs take place frequently during the year.

In the area of international outreach, we were active in Africa and South America. In South Africa, OSG staff conducted two grid training schools within the past 12 months: one workshop at the Witwatersrand University in Johannesburg in July 2008 and another workshop at the University of Johannesburg in April 2009. This second workshop was followed by a series of

encounters with domain scientists with the goal of providing them with guidance in using cluster and grid computing techniques for advancing their research. In South America, a week-long event was co-organized with a team in Brazil to present the Second Brazilian LHC Computing workshop in December 2008, Sao Paolo, Brazil. OSG provided the staff for teaching a 3-day course in grid computing at Universidad de Chile in Santiago, Chile. The Chilean team has solicited our help in building the Chilean National Cyberinfrastructure and we are currently in discussions trying to understand their specifics. We have ongoing discussions regarding possible ways of supporting the implementation of their regional and national cyberinfrastructure in Colombia; OSG has trained a few staff members as grid site administrators via US-based events and is ready to offer further advice and expertise in the practical aspects of their projects.

3.4. Outreach Activities

3.4.1. U.S. Outreach

We present a selection of the presentations and book publications from OSG in the past year:

- Joint EGEE and OSG Workshop at the High Performance and Distributed Computing (HPDC 2009): “Workshop on Monitoring , Logging and Accounting, (MLA) in Production Grids. <http://indico.fnal.gov/conferenceDisplay.py?confId=2335>
- Presentation at BIO-IT WORLD CONFERENCE & EXPO 2009, Ramping Up Your Computational Science to a Global Scale on the Open Science Grid <http://www.bio-itworldexpo.com/>
- Presentation to International Committee for Future Accelerators (ICFA) <http://www-conf.slac.stanford.edu/icfa2008/Livny103008.pdf>
- Contributions to the DOE Grass Roots Cyber Security R&D Town Hall white paper.
- Book contribution for “Models and Patterns in Production Grids” being coordinated by TeraGrid.
- <http://osg-docdb.opensciencegrid.org/0008/000800/001/OSG-production-V2.pdf>
- Workshop at Grace Hopper conference. <http://gracehopper.org/2008/assets/GHC2008-Program.pdf>

3.4.2. International Outreach

- Co-sponsorship of the International Summer School on Grid Computing in France (<http://www.issgc.org/>). OSG is sponsoring 10 students to attend the 2 week workshop, provided a key-note speaker and 3 teachers for lectures and hands-on exercises. We are also presenting OSG engagement to the students before the school, and following up with individuals following the school. (We thank NSF for additional funds through Louisiana State University for this program)
- Continued co-editorship of the highly successful International Science Grid This Week newsletter, www.isgtw.org. OSG is very appreciative that DOE and NSF have been able to supply funds matching the European effort starting in January 2009. A new full time editor has been hired and will start in July 2009. Future work will include increased collaboration with TeraGrid.

- Presentations at the online International Winter School on Grid Computing
<http://www.iceage-eu.org/iwsgc09/index.cfm>

4. Publications and Products

4.1. Journal publications

These are listed in detail in attachment 2 entitled “OSG VO Publications.”

4.2. Book(s) and/or other one time publication

“New Science on the Open Science Grid”, Ruth Pordes et al. Published in J.Phys.Conf.Ser.125:012070,2008.

“The CMS experiment at the CERN LHC”, by CMS Collaboration (R. Adolphi et al.), 361pp. Published in JINST 3:S08004, 2008.

“The Open Science Grid status and architecture”, Ruth Pordes et al. Published in J.Phys.Conf.Ser.119:052028,2008.

4.3. Other specific products

4.3.1. Teaching aids

OSG developed web based training materials for Grid Schools continue to be used by Gregor von Laszewski, Associate Professor in Computer Science at RIT and Computer Science departments in South America.

4.3.2. Technical Know-How

OSG is developing an experienced and expert workforce in the operational, management and technical aspects of high throughput production quality distributed infrastructures. This experience includes the use, diagnosis, security and support of distributed computing technologies including Condor, Globus, X509 based security infrastructure, data movement and storage, and other technologies included in the Virtual Data Toolkit.

4.4. Internet dissemination

OSG co-sponsors the weekly newsletter International Science Grid This Week: <http://www.isgtw.org/>. The other major partner in this newsletter is the Enabling Grids for EScience (EGEE) project in Europe. Additional contributions, as well as a member of the editorial board, come from the TeraGrid. The newsletter has been very well received, having just published 130 issues with subscribers totaling approximately 4,800. It covers the global spectrum of science, and projects that support science, using distributed computing.

OSG research highlights each describe a science result from the project:
http://www.opensciencegrid.org/About/What_We%27re_Doing/Research_Highlights

The results published in the last year are accessible via the following links:

- [Clouds make way for STAR to shine](#) (April 2009)
- [Single and Playing Hard to Get](#) (March 2009)
- [Protein Structure: Taking It to the Bank](#) (December 2008)
- [Opportunistic Storage Increases Grid Job Success Rate](#) (October 2008)

- [Simulating Starry Images-Preparing for the Dark Energy Survey](#) (July 2008)

And these are also provided in attachment 3 entitled “OSG Research Highlights.”

OSG has a comprehensive web site and information repository: <http://www.opensciencegrid.org>.

5. Contributions

Describe the unique contributions and specific products of your project, include major accomplishments and innovations and the success of your project. Contributions within the Discipline: How have your findings, techniques you developed or extended, or other products from your project contributed to the principal disciplinary field(s) of the project?

5.1. Contributions within Discipline

Contributions within Discipline – What?

Having summarized project activities and principal findings in one earlier section, and having listed publications and other specific products in another, here say how all those fit into and contribute to the base of knowledge, theory, and research and pedagogical methods in the principal disciplinary field(s) of the project.

Please begin with a summary that an intelligent lay audience can understand (Scientific-American style). Then, if needed and appropriate, elaborate technically for those more knowledgeable in your field(s). How you define your field or discipline matters less to NSF than that you cover (here or under the next category – “Contributions to Other Disciplines”) all contributions your work has made to science and engineering knowledge and technique. Make the most reasonable distinction you can. In general, by “field” or “discipline” we have in mind what corresponds with a single academic department or a single disciplinary NSF division rather than a subfield corresponding with an NSF program – physics rather than nuclear physics, mechanical engineering rather than tribology, and so forth. If you know the coverage of a corresponding NSF disciplinary division, we would welcome your using that coverage as a guide.

Contributions within Discipline – Why?

A primary function of NSF support for research and education – along with training of people – is to help build a base of knowledge, theory, and technique in the relevant fields. That base will be drawn on many times and far into the future, often in ways that cannot be specifically predicted, to meet the needs of the nation and of people. Most NSF-supported research and education projects should be producing contributions to the base of knowledge and technique in the immediately relevant field(s).

The OSG has delivered to the science of the physics collaborations who are the major stakeholders and helped to refine and advance the capabilities of distributed computing technologies.

5.2. Contributions to Other Disciplines

During the last 12 months OSG has added contributions to:

Protein Modeling and structure Prediction: Researchers at the Toyota Institute in Chicgao routinely use the OSG for running the protein threading software RAPTOR for protein structure prediction. This software is now number 2 world-wide in the Critical Assessment of Techniques for Protein Structure Prediction - a biennial experiment sponsored by NIH which represents the Olympic Games of the protein structure prediction community. Science papers based on the results are in publication review.

Structural Biology: through support of applications hosted by, and partnership with, the Structure Biology Grid project at Harvard Medial School.

Molecular Dynamics: modeling gases through work at the University of Buffalo.

Computer Science: through production processing support as part of algorithm development for genetic link analysis at the Technion; Through informing and helping the scalability and performance enhancements of BeStMan, Condor, Globus, dCache, and Pegasus software.

Mathematics: through modeling which is leading to a published paper (under review) at Colorado State University.

Campus Shared Cyberinfrastructure: through work at the University of North Carolina, University of Nebraska, Duke University and others.

5.3. Contributions to Education and Human Resources

Contributions to Human Resources Development– What?

Describe how your project has contributed to human resource development in science, engineering, and technology by:

- * providing opportunities for research and teaching in science and engineering areas;*
- * improving the performance, skills, or attitudes of members of underrepresented groups that will improve their access to or retention in research and teaching careers;*
- * developing and disseminating new educational materials or providing scholarships; or*
- * providing exposure to science and technology for pre-college teachers, young people, and other non-scientist members of the public.*

Contributions to Human Resources Development – Why?

A major aim of NSF programs is to contribute to the human-resource base for science and technology, including the base of understanding among those who are not themselves scientists or engineers. A core NSF strategy is to encourage integration of research and education. NSF needs to know and be able to describe how the work we support actually furthers that aim and that strategy. Moreover, contributions of this sort are important in the evaluation of results from your project when we and reviewers are considering a new proposal.

See Section 3.3 Training and Development.

5.4. Contribution to Resources for Science and Technology

Contributions to Resources for Research and Education – What?

To the extent you have not already done so in describing project activities and products, please identify ways, if any, in which the project has contributed to resources for research and education used beyond your own group and immediate colleagues, by creating or upgrading:

- * physical resources such as facilities, laboratories, instruments, or the like;*
- * institutional resources for research and education (such as establishment or sustenance of societies or organizations); or*
- * information resources, electronic means for accessing such resources or for scientific communication, or the like.*

Contributions to Resources for Research and Education – Why?

Physical, institutional, and information resources are important parts of the science and technology base that NSF seeks to sustain and build. Where particular projects build or sustain those resources for a broader community of scientists, engineers, technologists, and educators, that is a significant outcome which should be counted among the results that have come from federal support of science and engineering research and education. And you should get credit for those results.

Some NSF projects serve this purpose in a direct and primary way and so might report the outputs in earlier sections. Many NSF projects do not serve it at all, and are not expected to. But many serve it in ways ancillary to their primary purposes and activities. This is the place to report such contributions.

The OSG infrastructure currently provides access to the following resources. It must be remembered that OSG does not own any resources. They are all contributed by the members of the OSG Consortium, and are used both locally and by the owning Virtual Organization. Only a percentage that varies between 10 and 30% are in general available for use by the OSG.

Number of processing resources on the production infrastructure	89
Number of Grid interfaced data storage resources on the production	34
Number of Campus Infrastructures interfaced to the OSG	4
Number of National Grids interoperating with the OSG	2
Number of processing resources on the Integration infrastructure	21
Number of Grid interfaced data storage resources on the integration infrastructure	6
Number of Cores accessible to the OSG infrastructure	~49,000
Size of Tape storage accessible to the OSG infrastructure	~10 Petabytes @ LHC Tier 1s
Size of Disk storage accessible to the OSG infrastructure	~10 Petabytes
CPU Wall Clock usage of the OSG infrastructure	Average of 25,000 CPU days/day during May 2009

5.4.1. The OSG Virtual Data Toolkit

The OSG Virtual Data Toolkit (VDT) provides the underlying packaging and distribution of the OSG software stack. VDT continues to be the packaging and distribution vehicle for Condor, Globus, MyProxy, and common components of the OSG and EGEE software. VDT packaged components are also used by EGEE, the LIGO Data Grid, the Australian Partnership for Advanced Computing, GridUNESP the Sao Paulo state grid, and the UK national grid, and the underlying middleware versions are shared between OSG and TeraGrid.

Much of the work for the OSG Virtual Data Toolkit (VDT) has been focused on the needs of the OSG stakeholders and the OSG software release (described previously in Section 2.2.2). The VDT continues to be used by external collaborators. EGEE/WLCG uses portions of VDT (particularly Condor, Globus, UberFTP, and MyProxy). The VDT team maintains close contact with EGEE/WLCG and TeraGrid and OSG continue to maintain a base level of interoperability by sharing a code base for Globus, which is a release of Globus, patched for OSG and TeraGrid's needs. The Earth Science Grid (ESG) has investigated adoption of the VDT, and several discussions have been had with them about it.

5.5. Contributions Beyond Science and Engineering

None

6. Special Requirements

Provide evidence that you have complied with any special award terms. (These usually pertain to Cooperative Agreements).

6.1. Objectives and Scope

A brief summary of the work to be performed during the next year of support if changed from the original proposal.

No change.

6.2. Special Reporting Requirements

OSG has put in place processes and activities that meet the terms of the Cooperative Agreement and Management Plan:

- The Joint Oversight Team meets periodically, as scheduled by DOE and NSF, via phone to hear about OSG progress, status, and concerns. Follow-up items are reviewed and addressed by OSG, as needed.
- Two intermediate progress reports were submitted to NSF in February and June of 2007.
- The Science Advisory Group (SAG) met in June 2007. The OSG Executive Board has addressed feedback from the Advisory Group. Another meeting of the SAG is planned for late 2009.
- In February 2008, a DOE annual report was submitted.
- In July 2008, an annual report was submitted to NSF.
- In December 2008, a DOE annual report was submitted.

As requested by DOE and NSF, OSG staff provides pro-active support in workshops and collaborative efforts to help define, improve, and evolve the US national cyberinfrastructure.