



Document Name	OSG Project Status, February 2007
Authors	OSG Executive Board

Summary .....	1
Deliverables and Major Milestones Met.....	2
Progress Reports.....	2
OSG Facility .....	2
Education and Outreach .....	4
Extensions .....	4
Office of International Science and Engineering Activities .....	4
Communications and Administration.....	4
“SciDac Expectations” Deliverables.....	5
Project Management .....	5
Web Site.....	5
Project Reporting .....	6
Coordination and Communication Activities.....	6
Institutional PI reports.....	6
DOE Laboratories (to accompany FWPs): .....	6
Brookhaven National Laboratory: .....	6
FermiLab .....	7
LBNL .....	8
Stanford Linear Accelerator Laboratory.....	8
Universities: .....	9
Boston University: .....	9
Caltech .....	9
Columbia University.....	9
Cornell University.....	9
USC Information Sciences Institute.....	10
RENCI .....	10
University of California San Diego .....	10
University of Chicago.....	11
University of Iowa .....	11
University of Wisconsin .....	12
Indiana University.....	12
University of Florida.....	13
Expectations for SciDAC projects 9/13/2006 .....	14

## Summary

The OSG Project has organized for the next phase of the program of work and the new structure needed for a funded project with planned deliverables, milestones and metrics. The OSG Consortium and Council have continued to be engaged and supportive, endorsing all actions requested to date and providing an ongoing forum for contributing the resources used by, the applications using and engaging new and old partners and organizations.

The OSG Operational Performance has met the stated needs of the stakeholders, but the availability of sites across the facility to a broad set of users other than the “owners”, and the ease of use of the sites once they are available, are major focuses of necessary work.

### **Deliverables and Major Milestones Met**

Ready mechanisms to interface Condor-based local pools to OSG infrastructure	10/02/06
Testing and Validation Frameworks	10/02/06
GOC Risk Analysis Report	10/30/06
Gather User requirements for OSG 0.6.0	10/31/06
Initial test release of SRM/dCache for installation on OSG sites	11/15/06
Document & deploy the improved process	11/28/06
Interoperability with EGEE ticket handling system achieved	12/01/06
<b>Baseline OSG First Year Plan</b>	<b>12/01/06</b>
<b>Evaluate common OSG and EGEE Site Functional Tests</b>	<b>12/05/06</b>
Develop monitoring for OSG Authentication Service (GUMS)	12/06/06
Documented plan for Panda/Condor integration phase 1	12/14/06
<b>Validate the SRM/dCache prototype deployment candidate</b>	<b>12/15/06</b>
Specify transfer metrics for viewing on the first OSG transfer aggregation prototype.	12/15/06
<b>Release VDT for OSG 0.6.0</b>	<b>12/19/06</b>
<b>Release Security Plan</b>	<b>01/02/07</b>
Accept & process 15 identity services	01/02/07
Complete integration to preliminary VDT release candidate	01/02/07
Internal Review	01/02/07
Extend VO Management Service (VOMS) monitoring	01/05/07
Provide facility documentation	01/23/07
Demonstrate capacity to handle 50 tickets a week	01/30/07
<b>Sustained operations of LIGO workflow at UCSD at the level of 25 jobs for one week.</b>	<b>02/01/07</b>

LIGO's first milestone on OSG involved running a LIGO workflow for 7 days on at least 25 batch slots at all times at the UCSD tier-2 site, thus being a joint project for DISUN and OSG of LIGO. To meet this milestone, a number of hurdles had to be overcome, most notably the clustering of jobs on a batch slot. This exercised a new functionality within Pegasus, the USC Information Sciences Institute provided Middleware LIGO is using. LIGO exceeded the milestone by close to a factor 4, and a couple weeks ahead of schedule. The graph (generated using Monalisa) shows the execution of LIGO jobs at OSG sites in January. Days 11-18 show a continuous amount of work being done at UCSD, with many peaks of 100 jobs.

### **Progress Reports**

#### **OSG Facility**

During the first six months of the project, the OSG Facility has:

Operated the OSG infrastructure which was used by 20 VOs, of which ATLAS, CDF, CMS, sustained production running of several hundred to a thousands of jobs, and GADU, GLOW, Nanohub had periods of production running. D0 tested 10 sites on the infrastructure for a reprocessing run of ~1500 CPUs for several months that is starting in January.

ATLAS	Met the US ATLAS production running throughput needs. PANDA managing
-------	--

	production on more than 10 sites, with total CPU usage of more than 100K CPU hours per day.
CDF	Made significant contributions to CDF remote Monte Carlo production running across 8 sites on the OSG. 3 are “opportunistic” i.e. not otherwise used compute cycles.
GADU	Provided 2000 CPU days in 30 days opportunistic use.
GLOW	Provided 6000 CPU days in 60 days across 15 sites of opportunistic use.
CMS	Met the CMS CSA06 data challenge deliverables and milestones. Job success rate ~90%.
DZero	Testing of 10 sites for reprocessing run to start in February 2007.

Released four new versions of the VDT—including a production release, which will be used by the next release of the EGEE as well as OSG software—and two minor releases of VDT and the OSG software cache to fix security issues in the software. These new releases offered features to enable easier updating of the software, the new OSG accounting service and site “proxy caching” for the LHC experiments, upgraded versions of >10 middleware packages, added new supported platforms, and fixed numerous bugs. DCache was also integrated into the VDT.

Established the Validation testbed, which is a small well controlled set of sites, which can provide very quick turnaround to new VDT development releases. The Integration Testbed of ~20 sites is currently validating the OSG 0.6.0 release. The Validation testbed’s work resulted in a significantly more reliable release before the Integration Testbed began its work. The details are well [documented on the Twiki](#).

Responded to four Security Incident alerts. One of these resulted in a patch to the OSG software release. Delivered the OSG Risk Assessment (restricted) and [Security Plan](#) (OSG document 389). These were additionally requested as input to an Open Grid Forum security committee. Wrote Security Agreement for hosting of a core OSG service – in this case to the Tilted Planet company who host the OSG web portal. Attended the DOE Cybersecurity conference and submitted a [blueprint for OSG Security](#).

Established the Troubleshooting team, defined their procedures and protocols and accepted three troubleshooting activities, with many details [documented on the Twiki](#):

- STAR job efficiency on STAR OSG sites. This is proving a good exemplar of the need for such activities – the application and OSG teams working closely to test the underlying services and gradually work up to understand the full end to end system. This work is resulting in configuration changes at NERSC – the site selected for the focused tests, and the storage software (SRM/DRM).
- Web service Gram configuration and use by OSG applications on the Integration Test bed.
- Condor configuration and use on sites without shared files systems.

Identified potential new users for OSG, established relationships with them and started to provide assistance in getting their applications operational on the OSG. The Engage VO has been established.

During this period the weekly Facility meetings have been established and a smooth transition happened between Operations Coordinators due to the departure of Leigh Grundhoefer.

## **Education and Outreach**

These activities have been in preparations for the hands on Grid workshops and participation in the I2U2 project. The material from the previous schools is well on the way to being reworked to be reusable also as introductory user material for new OSG users. The first grid workshop is a collaboration with University of Illinois Chicago, Northwestern University and the University of Chicago and will be at UIC on March 23-24<sup>th</sup>.

We are preparing to post and hire a full-time OSG EOT coordinator. We are participating in the Program Committee for the next International Grid Summer School and contributed to the program.

## **Extensions**

OSG Storage Services: Extensions were made to the SRM/dCache for packaging, documentation and support on OSG, authorization, and the dCache billing/accounting and testing system. The OSG deployment infrastructure is integrated back into the main line of the dCache code.

OSG Workload Management: This is designed to be a “just in time” system, where pilot jobs are sent to the available execution sites and the decision of which job to run on which site is taken when the site is available to receive it. The goal is to release this new service into VDT/OSG over the next six months. Both PANDA and the CMS-Gfactory are prototypes of this service. Planning meetings between OSG, ATLAS, CMS and Condor have established the roadmap and the OSG extensions and integration of the existing services have begun.

Collaboration with the Center for Enabling Distributed Petascale Science: OSG members regularly attend CEPDS meetings. The Extensions for Storage, Troubleshooting and “edge services” activities of the OSG are working with their counterparts in CEDPS to ensure coherence in the deliverables and approachers. The OSG Integration TestBed is offered as an early testbed for CEDPS deliverables.

## **Office of International Science and Engineering Activities**

We helped organize and then attended the “Information and Discovery Services Workshop” at the Nordic DataGrid Facility (NDGF) to work on the next phase of interoperability to ensure transparent access across NDGF and OSG for data movement and job scheduling. The following junior faculty from OSG, with their students, are participating in this activity:

- Adriana Iamnitchi, Assistant Professor, Computer Science and Engineering, University of South Florida;
- Shaowen Wang, Adjunct Assistant Professor, Department of Geography, Ph. D. The University of Iowa;

The following junior researchers from the US DOE Laboratories and the Globus project are also participating in the activity: Jenny Schopf, Laura Pearlman, Globus; Gabriele Garzoglio, Steve Timm, Fermilab.

## **Communications and Administration**

During the past six months the OSG staff has:

- Written and signed 75% of the Statements of Work (SOWs) to disburse funds to the institutions. Completed the baselined WBS and change control process for OSG. (This is our first attempt – expect revisions). This involved detailed discussions and agreements as to the individual site responsibilities.
- Established processes for the Resources Managers and Finance Board and held first Finance Board meeting. Held two Executive Board meetings, and two OSG Blueprint meetings.

- Had regular phone discussions between OSG and TeraGrid management and are participating in planning and organization for the first Campus Cyberinfrastructure Days.
- Hired the OSG Communicator (Anne Heavey) and iSGTW Editor (Daniel Venton).

We held a planning meeting with the Nuclear Physics Stakeholders in OSG and developed a joint member between the Experiment and OSG management defining the next steps in our roadmap and collaboration ([OSG 520](#)) Mechanisms for communication and the actual collaboration with the External Technical Projects on which OSG depends for software and application development have been established. The following External Projects are members of the OSG Executive Board:

- Community Driven Improvement of Globus Software (CDIGS) /distributed systems technologies.
- Center for Enabling Distributed Petascale Science (CEDPS).
- Condor.
- dCache.
- Data Intensive Science University Network (DISUN).
- DRM.
- Laser Interferometer Gravitational-Wave Observatory (LIGO) Physics and the Information Frontier.
- OSG Accounting.
- OSG Privilege/Authorization.
- Security for Open Science.
- UltraLight/advanced networks.
- U.S. LHC (US representation to the WLCG).

## **“SciDac Expectations” Deliverables**

---

### ***Project Management***

The OSG Project Plan (2006-12) is [OSG Document 514](#).

The OSG Management Plan is [OSG Document 314](#).

The signed Statements of work include the following text about “Source code” developed by the project *“To comply with clauses defined in the SciDAC grant solicitation documents, any software produced as a result of the work done using SciDAC funds will be placed under open source licensing agreement.”* The VDT includes licence information from each software product showing the open source nature of the contributions.

OSG Metrics document (in draft form) published as OSG Document 541.

### ***Web Site***

The OSG communications and administration staff have responsibility for the web presence for OSG:

[Web portal](#) (<http://www.opensciencegrid.org>) for overview and communication, project overview, research plan, publications, presentations, interactions, progress reports, [OSG Twiki](#) based collaborative documentation area where all activity and technical information.

[Managed Document Repository](#) for the Project and Consortiums reference documentation including security, agreements and policies. The document librarian is Marcia Teckenbrock [marcia@fnal.gov](mailto:marcia@fnal.gov)

David Skinner from the SciDAC outreach center has agreed to attend the upcoming OSG All Hands meeting in March to discuss with the Communications group what information we can provide and how.

***Project Reporting.***

Overview of OSG in 6 slides: [OSG Document 506](#).

This note (OSG 540) is the first OSG Project progress report.

The December OSG monthly report and this note provides the first information about project accomplishments.

***Coordination and Communication Activities***

OSG PI and Executive Director attended the February SciDAC-2 kickoff meeting.

OSG Security and Policy Officers attended the [DOE Open Cybersecurity workshop](#) and submitted a blueprint for [OSG Security](#).

Plans for an article (possibly joint with CEDPS and ESG) in the SciDAC review.

Monthly OSG newsletter.

**Institutional PI reports**

---

**DOE Laboratories (to accompany FWPs):**

***Brookhaven National Laboratory:***

OSG Identity Mapping Services (GUMS):

GUMS V1.2 has been extended, is now feature complete, and is now being tested/polished. New features include requests for extended security access; improved configuration management and monitoring; refactoring of the code for support and maintenance; additional granularity for the access mappings;

Progress in the Workload Management area (third column indicates where support for the work came from):

Sep 1 2006	First planning meeting for OSG just-in-time workload management extensions (at FNAL); define initial program (Glide-in factory for Condor-based site-local pilot submission with centralized grid-wide pilot management (factory submission to sites via the grid)	
Sep 2006	Defined and initiated generic VO-neutral Panda development (TestPilot subsystem)	ATLAS
Sep 2006	First dedicated manpower on BNL OSG extensions program (UT Arlington student, 50%)	OSG, ATLAS
Sep/Oct 2006	Established Condor testbed at BNL for Panda/Condor pilot factory development and integration	ATLAS, OSG
Oct 2006	Demonstrated Panda pilot operation and job processing on generic OSG, LCG sites (no ATLAS specificity)	ATLAS
Oct 2006	Demonstrate ATLAS analysis (pathena) operation on generic OSG, LCG sites: this is the only milestone defined for this period that is not yet completed. It is ~60% complete.	ATLAS
Oct 2006	Provided framework for Panda-based execution of non-ATLAS VO workload on generic OSG	ATLAS
Nov 2006	Demonstrated Panda-based execution of non-ATLAS VO	ATLAS,

	workload (CHARMM) on generic OSG	CHARMM
Nov/Dec 2006	Drafted Panda/Condor first-phase integration plan, with glide-in based pilot factory as initial target	OSG, ATLAS, Condor
Dec 2006	Second planning meeting for OSG workload management extensions (at UT Arlington); agreed on first-phase integration plan for workload management bringing together Condor, ATLAS, CMS developments	ATLAS, CMS, OSG, Condor
Dec 2006	Reached operating scale of 250 Panda queues at almost 200 SEs (gatekeepers) across OSG and LCG	ATLAS
Dec 2006	Established second Panda instance (CERN) to develop and test multi-instance operation	ATLAS
Oct-Dec 2006	Investigated and tested various Condor components for use in pilot factory and Condor integration	OSG, ATLAS
Jan 2007	Resolved Panda DB performance issues, improving scalability of current Panda configuration (single server/DB instance) by a factor of ~5-10	ATLAS
Jan 2007	Adapted generic VO-neutral Panda subsystem (TestPilot) to support ATLAS production	ATLAS
Feb 2007	Proceeding with BNL OSG extensions hires, after hold due to budget issues	

### ***FermiLab***

#### VDT:

The first packaging, documentation and prototype releases of the dCache based storage service have been delivered to VDT. The installation has been tested on four OSG CMS sites. It will be deployed on the Integration Test Bed over the next month.

#### Security:

The Security Plan has been delivered. Operational Security is in process. Policies developed jointly by the OSG, EGEE and WLCG (through the Joint Security Policy Group) are in progress.

#### Extensions and User Support

Extensions has focussed on those needed to include SRM/dCache into VDT have been done. This has been in close collaboration with DESY as part of the dCache collaboration. Testing and deployment of SRM V2.2 in test mode for the WLCG, LHC Experiments, and shortly the OSG has been done.

Support has been provided for the GADU, Engage, D0 and NanoHub activities on OSG. A “VO” based test program has been prototyped to feed into the core site testing frameworks – VORS, GridEX – to help in this support.

#### Communications and Administration

The OSG Communicator has been hired (Anne Heavey) and the communications team has provided 3 OSG 1 pagers, 1 poster, the OSG news and science highlights. The OSG projec associate has been hired (Chander Sehgal) and has taken over the budget and planning processes.

## **LBL**

Progress relative to Milestones and Deliverables for the Year1 OSG-LBNL SOW. There is one Agency Reportable Milestone in year 1:

*Feb. 27, 2007 SRM Tester V2.2 for OSG 0.6, and technical evaluation of OSG 0.6*

Status: The SRM Tester V2.2 is being used to test SRM-based storage elements on OSG. The technical evaluation of the OSG 0.6 software release is in progress and LBNL is participating at the level of 1 FTE.

Other deliverables:

- *Setup, install and maintain and ITB reference site (VTB) as an on-going member of the integration testbed.* A node has been set up and used for installations of the appropriate VDT releases for the current VTB and ITB activities. A 3-node VTB cluster is in progress of being deployed.
- *Technical evaluation of the OSG 0.6 software release* Documented installation issues with the VDT 1.6.1 release. Continued testing of the Gratia-SGE probe, part of the accounting infrastructure.
- *Establish SRM test procedures for the OSG 0.6 software release, SRM-Tester V2.2:* SRM Tester V2.2 in use, testing SRM/CASTOR (CERN, RAL), SRM/dCache (FNAL), SRM/DPM (CERN), SRM/StoRM (INFN), SRM/BeSRM (a.k.a. DRM, LBNL), SRM/LStore (VU). OSG SE Monitoring server node is being set up.
- *Establish and maintain SE monitoring web site for OSG SEs.* SRM Tester results are posted to twiki web site.
- *Integrate all OSG VOs into the OSG RA procedures for handling grid certificate requests.* All regular VOs with any significant number of certificates have RA agents identified and operational. All remaining VOs have Sponsors identified so that OSG operations agents can process requests.
- *Establish improved functionality of the DOEGrids CA and OSG RA procedures for LIGO, as part of the OSG-LIGO PIF collaborative program.* In progress – LIGO has determined precise requirements needs and the implementation plan based on those requirements will be developed in February.

## **Stanford Linear Accelerator Laboratory**

The effort is focused on Policy and collaboration with our peers – to ensure a consistent and common security policy infrastructure across many federated grids for LHC and other collaborations.

Middleware Security Group (MSWG); Joint EGEE, OSG security middleware technical group.	In the next day, we primarily spent it discussion the various authorization models in use between EGEE and OSG, and discussing future directions while maintaining interoperability.
Joint Security Policy Group (OSG policy group in collaboration with EGEE and WLCG)	Drafting top-level security policy document and Grid Site Agreement to present to the OSG Executive Board and Council for approval and implementation. To date OSG has not agreed to the policies as worded.
Policy Management Association (PMA) meetings	Contribute to The Americas Grid Policy Management Association (TAGPMA); International Grid Task Force (IGTF) group of all PMAs, European Union (EU) gridPMA.

Requirements and standards groups	<a href="#">Cyber Security R&amp;D Workshop: Open Grid Forum 19</a> : day on Federated Identity. Internet2 and Federated Identity (Shibboleth) requirements meeting.
-----------------------------------	--

## **Universities:**

### ***Boston University:***

Released Pacman 3.19 in response to changes requested by the VDT. Support has been provided as needed. The SOW is in progress.

### ***Caltech***

Caltech has either completed or is on track for all its milestones to date. The Resource co-Manager has been active in all activities of the Executive Team and has reviewed and coordinated the SOW process. The group is still working through the Caltech Office of Sponsored Research to gain approval of Caltech Statement of Work and Budget, due to changes in the charges for overheads being proposed.

- Contributed to the organization and implementation of the new OSG Validation Test Bed (VTB).
- Using the new VTB, installed, tested, evaluated candidate OSG 0.6.0 software stack.
- Provided leadership, coordination and technical expertise in both the VTB and Integration Test Bed (ITB).
- Work hand in hand with Pegasus developers to support new workflows capable of using locally cached LIGO data sets.
- Using LIGO Binary Inspiral workflow, met requirement to run LIGO analysis pipeline at a sustained level of 25 job slot or more for one week at University of California San Diego OSG site.
- Working to extend LIGO workflow technology to support Storage Resource Manager (SRM) on the OSG.
- Working with Pegasus developers to extend functional support of worker nodes lacking shared file systems.

### ***Columbia University***

OSG funds will provide partial support for the Site Administrator/Developer, and Wits has now secured additional funding to make this a full-time position. Potential candidates are being identified, with the goal of filling the position by March 31, 2007. OSG Education, Outreach and Training staff, will provide technical expertise and a visit once the cluster is established.

### ***Cornell University***

Cornell has been contributing to the Storage Service extensions for accounting/billing in dCache:

- Devised and implemented a prototype web service to capture the billing data from all the CMS dCache servers at a central site.
- Based on web-service experience, implemented and tested a prototype dCache billing data probe to register data with Gratia (meets 1/15/07 milestone)
- Transfer metrics specifications (12/15 milestone) has been transferred to Gratia project – billing probe will be updated by Cornell to meet Gratia specifications
- Ongoing review of dCache billing code: several bugs and necessary enhancements for proper, efficient billing probe operation reported (2nd deliverable).

### ***USC Information Sciences Institute***

The Pegasus team at USC Information Sciences Institute (ISI) has been working closely with LIGO on enabling the running of the large-scale LIGO analysis HIPE on the OSG resources. To this end ISI: 1) updated the Pegasus resource information provider to point to the new VORS OSG information system, 2) provided the capability to stage-in executables on demand to remote sites, and 3) enhanced the Pegasus partitioning capabilities to enable scalable execution of LIGO workflows. Finally ISI conducted research on the dynamic clean up of completed data dependencies. As a result of the joint LIGO-Pegasus efforts, we were able to achieve and surpass the first milestone: of running a constant of 25 jobs at the OSG UCSD cluster for a week.

### ***RENCI***

During this reporting period, we have focused on identifying and qualifying key potential new users for OSG, establishing relationships with these users, ramping up our capabilities to provide assistance in getting their applications operational on the OSG, and providing feedback to the consortium regarding our experiences. We have deployed a new OSG VO for the purposes of the Engagement activities (the Engage VO), and have succeeded in getting this new VO accepted by an initial set of resource providers. We have three active relationships with users who are new to OSG, representing the science domains of Biology (protein folding), atmospheric modeling, and materials science. Each of these engagement activities are in various stages of running codes on OSG, though the Biology application is nearly ready for their first large run, on the order of tens of thousands of jobs. We continue to recruit new high profile users in diverse communities, and log our experiences in bringing new applications and users to this infrastructure with an eye towards enhancing OSG's ability to attract and retain new user communities.

### ***University of California San Diego***

UCSD efforts on OSG are focused in the wbs area 1.3, Science Applications and Extensions. Focus in the last 6 month was particularly on wbs items 1.3.1.1 (User group coordination), 1.3.2.1 (OSG CE performance), 1.3.2.2 (Data Storage and Access Management), and 1.3.2.4 (LIGO Support). LIGO's first milestone on OSG exceeded the milestone by close to a factor 4, and a couple weeks ahead of schedule. UCSD's role was largely advisory here.

In the Data Storage area, the work revolved around two topics, integrating OSG wide aggregation of dCache data transfer into the GRATIA accounting system, and validating the SRM/dCache deployment candidate for the OSG 0.6 release in VDT. Both are archetypical extensions activities, as it brings developers and OSG together to spec, validate, and integrate in a rapid prototype cycle. UCSD deployed prototypes, verified functionality, and feedback bugs, including performance problems. UCSD has worked closely with the primary developers on readying the products for OSG 0.6 prior to testing in the OSG integration testbed. For the GRATIA integration, we also sped the accounting information to be aggregated, and are presently involved in designing the views of the accounting presentation.

The effort in determining the OSG CE performance is just starting. This is another joint activity between DISUN and OSG in that it uses the DISUN production cluster in a parasitic way in order to implement a large-scale testbed for performance testing of the OSG compute element. UCSD contributes to the coordination of the application testing of the upcoming OSG 0.6.0 release candidate in order to make sure that the release does not negatively impact the applications communities.

Finally, the UCSD group is hosting the OSG Consortium All Hands meeting on March 5-7th, followed by a US LHC Tier-2 meeting. Chairing the program committee, and coordinating this meeting always takes effort.

### ***University of Chicago***

#### Facility

At UC we have been leading the OSG Integration effort (WBS 1.1.3) with direct oversight responsibilities of the new OSG personnel at LBNL, Caltech, and ½ FTE (a new hire, November 2006) at UC and overall coordination with VDT and site administrators contributing to the Integration Test bed (ITB). We organized a new Validation Test bed (VTB) to rapidly incorporate patches and overnight pre-release changes of the VDT specifically for GRATIA probe testing and development and the OSG compute element configuration process. This completed milestones for establishing VTB sites at three separate institutions (UC, CIT, LBNL) and two others that joined (FNAL, IU). We created a community code repository based on Subversion, a Pacman cache, and automated build scripts (code to cache) for VTB candidates; this infrastructure will be used by the GOC for ITB release packages. We developed automated VTB site validation scripts and a web reporting system. We setup a Xen-based set of VTB hosts for testing multiple services, emulating a small OSG “site” using UC ATLAS Tier3 resources. We are contributing to the testing and validation of WS-GRAM. We lead the weekly ITB meeting, the as-needed weekly VTB (or ITB-follow-up issues) meeting, and we coordinated a virtual workshop with about 20 OSG participants. The ITB 0.5.2 release, based on VDT 1.6.1 was deployed on 13 sites during January 24-25, 2007, which will be the basis of the OSG 0.6 production deployment. We have started work on the SRM/dCache SE testing on ITB 0.5.2, having completed the initial installation on our development ITB cluster. Additional work includes reporting project status in weekly Facilities meetings, frequently ET meetings, and coordination with other efforts in OSG such as metrics definition.

#### Education and Outreach

- Planning and preparations for the March 23-24<sup>th</sup> Grid training workshop at the University of Illinois at Chicago (UIC), co-sponsored by UIC, Northwestern and UChicago.
- Planning and preparation for Grid workshop sessions at Argentina’s Grid School (Ben Clifford will teach 3 days between March 7-14).
- Reviewed all the tutorial material from the 2006 Grid Workshop with technical write Forrest Christian so that we have a common base for working, and for splitting it up into easier-to-reuse pieces (motivated initially by the needs of the March workshops in Argentina and Chicago)
- Revised workflow workshop and presented the new module at I2U2 collaboration meeting at Adler in Chicago.
- At the Open Grid Forum sessions on education and training, establishing relationship with the UK e-Science education and training leaders (Malcolm Atkinson and David Fergusson)
- Described US progress in I2U2 and on Grid Workshops at the OGF Education sessions
- Preparing to post and hire a full-time OSG EOT coordinator.
- Provided management and architecture support for I2U2 project (weekly calls, developer support, problem resolution)
- Refined plans for S. Africa support with Jeremy Dodd of Columbia University and provided initial support contact for site deployment and Grid administration (Ben Clifford).

### ***University of Iowa***

The University of Iowa coordinates and contributes to the Troubleshooting efforts. Here are the problems that are completed over the past six months, and those in progress.

- Low efficiencies of job execution observed by STAR VO (the cause of the problem has been identified).
- Publishing and Mapping VO's on CE's (solved).
- Problem executing Condor Standard Universe jobs via the OSG interface. (solved)
- Consistent and effective mechanisms to get resource compliance with OSG policy (initial investigation)
- Jobs switch between idle/running in Condor queue holding CPUs (solved).
- Validating BDII/GIP information at site-level (solved).
- Advertising Incorrect Cluster Resource Information (initial investigation).
- Submitting GRAM jobs to a Tomcat container that doesn't have GRAM installed (solved).
- Submitting to GRAM from a client with clockskew (solved).
- Hundreds of thousands of bogus entries in PBS server logs (fix being tested).

Other activities have involved: (1) the joint efforts with the Nordu DataGrid Facility (NGDF), EGEE and Globus on information discovery infrastructure and the Glue Schema attribute definitions; (2) OSG Generic Information Provider development and support.

### ***University of Wisconsin***

The Facilities Coordinator established weekly meetings of the Team, oversees and monitors the daily usage statistics of the distributed facility and technical roadmap of the OSG, takes part in phone calls, face-to-face meetings and workshops.

The OSG PI worked with the OSG institutions on developing their Statement Of Work and establishing the subcontracts.

The UW team contributed to the troubleshooting effort and maintained the operation of the Grid Exerciser for the ITB.

Helped to define and implement the program for the next International Grid Summer School.

Contributed to the “just-in-time” scheduling effort.

The Software Coordinator manages and contributes to the OSG provisioning and VDT releases.

### **VDT**

- Meetings: Regular VDT Office Hours, OSG Operations, Integration, Security, and Executive Team Plus meetings.
- Moved the VDT mailing lists from ivdgl.org to opensciencegrid.org, and announced the move.
- Worked on VDT 1.5.0, 1.5.1, 1.5.2, 1.6.0 and 1.6.1 releases.
- Updated and added many software packages in the VDT, which are documented in the [1.5.1](#), [1.5.2](#), [1.6.0](#) and [1.6.1](#) release notes.
- Collaboration with TeraGrid to converge on a uniform middleware base.
- Added dCache for storage management to the VDT.
- Worked with the Validation and Integration testbeds to improve the VDT release in preparation for the OSG 0.6.0 release.

### ***Indiana University***

Indiana coordinates and provides Operations effort to the OSG. The central ticketing and operations centers are at IU. IU is responsible for ongoing operations and improvements in the

operational infrastructure for OSG – including extensions to and operation of the validation, monitoring and availability services. Milestones and deliverables met are:

- Trouble Ticket Exchange with EGEE (GGUS) - Accomplished and Service Validation - Complete - 12/01/06.
- Processing 15 Identity Requests through OSG RA - 29 Requests Approved Since Inception of OSG RA - 15 Accomplished 11/10/2006.
- Demonstrate capacity to handle 50 trouble tickets per week - ~29 Tickets Opened per Week, ~29 Closed per Week, with ~67 Tickets in Open status at any given time. (Number gathered between August and December 2006).
- Develop Version 2 of VO Resource Selector - Goal 03/07/2007.
- Publish Operations Plan - 06/01/2007.

### ***University of Florida***

The University of Florida is organizing their efforts in education and outreach for OSG and will participate in OSG's EOT effort, including the grid training workshops. Other responsibilities are to provide training in site and system administration for the OSG. The effort will be provided by students supervised by Craig Prescott of the High Performance Computing Group. Prescott is responsible for maintaining the HPC as a component of OSG.

Paul Avery is the US International Science Grid This Week editorial advisor and also contributing as one of the Resources Co-Managers

## **Expectations for SciDAC projects 9/13/2006**

---

All SciDAC projects are being funded as cooperative agreements. The Statement of Substantial Involvement that is required for establishing cooperative agreements, coupled with content of the solicitation, allows us to set project requirements beyond what is normally expected. This document outlines a common framework for managing these projects.

A number of requirements will be expected in common across all SciDAC projects. After the awards are announced, it is expected that each program manager will communicate these requirements to their PIs.

1. Project Management Plan. Each SciDAC project is expected to provide a Project Management Plan, a baseline against which project progress can be measured. This is the document that describes how and when a project's objectives are to be achieved by showing the major activities, milestones, deliverables, and resources required on the project. (Milestones must identify the person and institution accountable.) It should address all intellectual property and software issues as outlined below. Many large proposals included an outline for a management approach and presented milestones and deliverables. However, we are asking that a complete plan be submitted with the first sixty days of the award
  - a. Science applications must include a detailed description of the plan for partnerships that integrates advanced applied mathematics and computer science technologies with the domain-specific efforts.
  - b. Source code developed by enabling technology projects is expected to be fully and freely available for use and modification throughout the scientific computing community via a pre-approved open source process. Enabling technology projects should specify the type of open source license that will be used and the mechanisms, including web sites, workshops, and other community-based activities that will be used to disseminate information about software developed.
  - c. All projects must include a plan for supporting their software over the long term.
  - d. All projects must establish a baseline against which progress can be measured. For example, at the beginning of the project, the performance of a science application code should be established and documented relative to scalability, number of processors, etc.
2. Project Reviews. Projects should expect an annual review, the first occurring on or before eighteen months into the project. Cognizant program managers will establish how this review should be conducted, ranging from a site visit by the program manager to a structured review by a committee of peers. For example, projects that are staffed by teams established during SciDAC-1 should expect a different management approach than projects with newly established teams.
3. Web Site. Each project is expected to have a visible web presence.
  - a. Each project is expected to develop and maintain a web site. It should include a project overview, research plan, publications, presentations, interactions, progress reports, etc. Namely, it should provide a resource to the cognizant program manager that is maintained with current information so that information is accessible quickly without relying on direct contact with the project personnel. The institution and person responsible for keeping the web site current should be identified.
  - b. The SciDAC.org web site is being updated and all projects are expected to provide input for that site, as requested.

4. Project Reporting. Each project is expected to provide the following project information.
  - a. Within the first month, each project should provide an overview presentation on the project to the cognizant program manager (4-6 slides). It should include a problem statement/vision, why the effort is important, goals of the project, technical challenges, and timelines.
  - b. Project progress reports are to be provided on a TBD schedule—it may be quarterly, or semi-annually, or annually, as directed by the cognizant program manager.
  - c. Project Accomplishments (single concise paragraph of significant results) to support the budget process should be provided to the cognizant program manager, as requested, on a yearly basis—late spring.
5. **Coordination and Communication Activities.**
  - a Projects are expected to publish in the *SDA C Review*, as appropriate
  - b Projects are expected to participate in the annual SciDAC meeting
  - c. Projects may be asked to participate in workshops and other focused coordination activities, as appropriate.

9/13/2006