



JANUARY 2007

Meetings and Events

ITB Virtual Workshop
Telephone Conference
January 24-25, 2007

19th Open Grid Forum (OGF19)
Chapel Hill, North Carolina
January 29 - February 2, 2007

OSG Consortium All Hands Meeting
San Diego, California
March 5-8, 2007

HPDC 2007: IEEE International Symposium on High Performance Distributed Computing (HPDC)
Monterey Bay, California
June 27-29, 2007

[View Full Calendar](#)

Citing OSG in Publications

Whenever you make use of Open Science Grid resources, services or tools, we would be grateful to have you acknowledge this use in any resulting publications. We suggest the following or similar text:

"This research was done using resources provided by the Open Science Grid, which is supported by the National Science Foundation and the U.S. Department of Energy's Office of Science."

ITB Virtual Workshop

We are pleased to announce a ITB Virtual Workshop January 24-25, 2007 to work on the next deployment and testing cycle of the Integration Test Bed. The goal of this workshop is to install ITB 0.5.2 on a number of sites in advance of the deployment of the next production release of the the OSG 0.6 software stack. The results will be a testbed that will enable further service validations (especially information and accounting services) and application validation by VOs. This workshop will be similar to the one held in March, 2006. We need volunteers to help with all phases of validation and documentation, in addition to checking the installation at your site.

Spotlight on an OSG Contributor

Open Science Grid News

[OSG HOME](#) | [SUBMIT NEWS](#) | [SUBSCRIBE](#) | [ARCHIVE](#) | [ABOUT OSG NEWS](#)

From the Executive Director

Welcome back after the holidays. I'll start by responding to your requests to provide information from the Executive Team and Executive Board. The Executive Team meets weekly to discuss the status of OSG and any issues that arise. Our decisions, arrived at by consensus as far as possible, are mainly technical and therefore affect you directly. The agenda and minutes are publicly posted. In the last few weeks we have decided to recommend deployment of the CEMON-based collection of information and to ask every site to install OSG Accounting, as it will be required in the next release of OSG.



The Executive Board, a bigger body with representation from projects and stakeholders, meets every six weeks or so. The EB also includes External Projects, non-OSG-funded software and service development projects, who report at our meetings.

I encourage you to look at the Year 1 Project Plan and the Security Plan that we've posted. Also peruse the document database to find other recent documents.

One of our primary goals is to get you, the users, successfully finding OSG sites that will run your applications. VORS tells if a site is up or down, but there is so much more to using a particular resource: data access, data flow, Message Passing Interface (MPI) and grid accessible storage, to name but a few. We have some ideas on how to give you a better handle on this, for example new sensors to send data to the GOC, sample customizable applications, and pilot jobs to monitor a resource. We welcome your ideas, too.

EGEE is using VDT 1.4.0 and the APAC Grid project is testing VDT. VDT 1.6.1 is released and being validated. One of the first challenges of the new year is getting the OSG 0.6.0 release out. For those of you who've asked, we're not ready to call it version 1.0. We know there is more to do to make the OSG infrastructure user friendly and as robust as local facilities are today. We don't want to imply that we have all the answers in place today.

The OSG Communications team is in the process of assessing this newsletter with the aim of enhancing its usefulness to the overall consortium. Do you read it? If so, which parts? What would you like to see in it? If you haven't already, please take a moment to send your thoughts on the newsletter to osg-contact@fnal.gov or call Anne Heavey at 630-840-8039. This newsletter is for you!

~ Ruth Pordes

OSG Joins Effort to Advance Information Systems

The NSF's Office of International Science and Engineering is sponsoring an OSG activity to advance interoperability with our Scandinavian peers. Towards this end, OSG is joining efforts with EGEE and NorduGrid to understand and define the needs for interoperability, scalability and robustness of our information services. OSG is also moving forward to extend the current information services (which provide grid users with access to information about grid resources) to include discovery services (which identify available grid services). University of



Participants in the Service Discovery in Grids Workshop in Copenhagen. (Click for larger image)



Wayne Betts has been with the STAR collaboration since 1994, in both physics and computing roles, and at Brookhaven National Lab for the past five years. He is currently a senior applications analyst, primarily providing systems administration and support for STAR's experiment control systems.

A year and a half ago, Wayne "dove into STAR's grid efforts by installing the ITB software stack on a gatekeeper with little idea what I was getting into." It soon occurred to him that a grid is much like a heterogeneous, multi-CPU computer. This gave him a better conceptual grasp of grid architecture with which to move ahead. He now maintains the STAR-BNL compute element's OSG configuration, assists other STAR site admins and helps STAR's "grid-aware" users to understand the infrastructure and components. Users and site admins, confronted with an error or failure, go to Wayne to help pinpoint the responsible component. He troubleshoots to determine if the cause lies in the software stack, an application failure or bug, user error, environment configuration, or some combination.

Currently, STAR's primary clusters are located at BNL and NERSC. STAR has only recently started "dipping its toes" into remote sites. Wayne and his colleagues have already used lessons learned from this experience to change practices at their own sites and to improve the STAR Unified Meta Scheduler's (SUMS) "grid-awareness."

Supported By



Office of Science/
U.S. DOE

South Florida is participating in these efforts.

Several OSG members attended the [Service Discovery in Grids workshop](#) held December 12-14, 2006 in Copenhagen: Adriana Iamnitchi, Lydia Prieto (USF), Laura Pearlman (ISI), Steve Timm (Fermilab), and Shaowen Wang (UI). At the workshop, we compared and contrasted past and current solutions, including MDS4, BDII and others, and found more similarities than expected. The information provider in most, for example, contains information about each site-level interface in an index (e.g., a catalog). On the other hand, grid level caching of fault tolerance and scalable performance information is typically handled differently by each. We walked through a number of common use cases in order to determine the kinds of information that users need to find, and to estimate the frequency of users' queries. This helped us produce a road map for moving ahead.

Beyond agreeing to work towards a common schema and interface definition, we analyzed bootstrapping, discussed approaches to scalability and fault tolerance for these systems, and developed a list of research topics. (Read our [full workshop report](#). OSG's next task in this area will be defining the Service Discovery interface at the upcoming Open Grid Forum (OGF19).

~ Adriana Iamnitchi & Lydia Prieto (University of South Florida)

OSG Generic Information Provider (OSG-GIP)

The Open Science Grid Generic Information Provider (OSG-GIP) has been developed by the GROW (Grid Research and education group @ IoWa) Collaboration. Originally developed by Laurence Field at CERN, GIP is a grid monitoring tool that aggregates static and dynamic resource information for use with LDAP-based information systems. The GIP produces information consistent with the GLUE Schema. The GROW collaboration customized and extended the GIP to enable LCG-OSG interoperability and OSG resource selection.

OSG-GIP components (See figure at right) can be summarized as follows:

1. Core GIP components - collect and publish monitoring information.
2. Configuration files - capture static information about the OSG resource.
3. Dynamic plugins - provide up-to-date information about resource usage.
4. Dynamic providers - allow flexibility of adding new information sources.

The OSG-GIP provides the following three main capabilities:

Automatic configuration:

GIP configuration requires extensive understanding of the GLUE schema, and is therefore difficult. The configuration tool in OSG-GIP gets GIP working out-of-the-box using existing OSG configuration information (e.g. `osg-attributes.conf`) so that minimal administrator intervention is required.

Monitoring:

Dynamic plugin and provider scripts capture and report the state of various OSG resources.

Validation:

An OSG-GIP validation tool, also known as [GlueChecker](#), pulls information from the OSG level BDII (previously from site level GRIS) and performs sanity checks. ([See the User-interface for GlueChecker.](#))

~ Anand Padmanabhan & Shaowen Wang (University of Iowa)



