

THE OPEN SCIENCE GRID

R. Pordes, Fermilab

for the Open Science Grid Consortium [1]

<http://www.opensciencegrid.org>

Abstract

The U.S. LHC Tier-1 and Tier-2 laboratories and universities are developing production Grids to support LHC applications running across a worldwide Grid computing system. Together with partners in computer science, physics grid projects and (active experiments, we will build a common national production grid infrastructure which is open in its architecture, implementation and use.

The Open Science Grid (OSG) model builds upon the successful approach of last year's joint Grid2003 project. The Grid3 shared infrastructure has for over eight months provided significant computational resources and throughput to a range of applications, including ATLAS and CMS data challenges, SDSS, LIGO, and biology analyses, and computer science demonstrators and experiments.

To move towards LHC-scale data management, access and analysis capabilities, we must increase the scale, services, and sustainability of the current infrastructure by an order of magnitude or more. Thus, we must achieve a significant upgrade in its functionalities and technologies.

The initial OSG partners will build upon a fully usable, sustainable and robust grid. Initial partners include the US LHC collaborations, DOE & NSF Laboratories and Universities & Trillium Grid projects. The approach is to federate with other application communities in the U.S. to build a shared infrastructure open to other sciences and capable of being modified and improved to respond to needs of other applications, including CDF, D0, BaBar, and RHIC experiments.

We describe the application-driven, engineered services of the OSG, short term plans and status, and the roadmap for a consortium, its partnerships and national focus.

THE OPEN SCIENCE GRID

The partners in the Open Science Grid Consortium have come together to build a sustainable national production Grid infrastructure in the United States that will be open to scientific collaborations. The Consortium will build on and evolve the existing Grid3 [2] common shared infrastructure together with the distributed computing facilities at the Laboratories and Universities, including the Fermilab SAMGrid [3] system.

The Open Science Grid Consortium will act to integrate, deploy, maintain and operate a shared common infrastructure to the benefit of all its users. To meet the long-term needs of the experimental physics community in the US, existing grid infrastructures must evolve by an order of magnitude or more in size, performance, and capabilities. The Consortium plans to evolve the grid infrastructure to meet both these needs and computational needs of other science partners.

Our vision is to make a significant step forward in cooperative development and Grid interoperation on a global scale. We plan to build a flexible real-time framework of services, using an engineered end-to-end approach and to accommodate multiple Grid systems and software suites, while ensuring a coherent operation of the whole system. The scope of the Grid infrastructure is to support science working environments ("gateways") for user groups ranging from individuals to teams of thousands of members, as well as a working environment for large-scale production Grid operations supporting 100s to 1000s of tasks with a wide range of demands for computing, data and network resources. The infrastructure will accommodate policies for resource usage and task priorities, while ensuring fair-sharing of network infrastructures and other resources.

The OSG Consortium is organized as a collaborative partnership with community goals & contributions, deliverables and self-organized governance: much like a physics experiment. The Consortium will involve all partners in the design, development and deployment to ensure that OSG will be able to function and evolve consistently while meeting the needs of individual members.

The US ATLAS [4] and US CMS [5] computing programs are contributing to the Consortium by presenting their storage and computing resources to OSG, developing and deploying their distributed data and processing services in a common environment, and running their applications on the shared infrastructure. Consortium participants come from DOE particle physics laboratory facilities, university and campus grids, physics experiments (including STAR [6], CDF [7], D0 [8], BaBar [9], LIGO [10], SDSS [11]), as well as computer science and bioinformatics. Grid technology providers (Condor [12], Globus [13]) are partners and also provide Grid middleware support through an evolution of the

Virtual Data Toolkit [14]. The US Trillium Grid projects (GriPhyN [15], iVDGL [16], PPDG [17]) are evolving their program of work to contribute to the Open Science Grid roadmap. Other scientific communities and projects are represented: e.g., ACDC at the University of Buffalo [18], the GLOW [19] campus Grid at the University of Wisconsin. OSG is partnering with other Grids, in particular the TeraGrid [20] to ensure effective use across the federated Grid infrastructures. The Consortium roadmap is endorsed by the High Energy Physics program offices in the DOE and NSF.

FACILITIES, INFRASTRUCTURE AND APPLICATIONS

The Open Science Grid (OSG) is the Grid under the governance of the Open Science Grid Consortium which is operated as a sustained production infrastructure for the benefit of the users. Other grids may operate under the governance of the Consortium, for example the Grid that is used for validating parts of the infrastructure before they are added to OSG. OSG interfaces to and/or includes participating facility, campus, and community grids, and interacts with grids external to the Consortium through federation and partnership. The OSG Virtual Organization (VO) is open to those Users and VOs that have contracts (agreements) with the Consortium.

The OSG program of work focuses on interfacing existing computing and storage facilities to the common infrastructure, adapting the applications to run on the shared infrastructure, and focuses on developing those services and operations that provide for a coherent, robust and operationally performant distributed system.

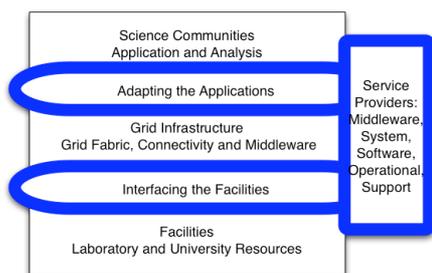


Figure 1: The Open Science Grid Infrastructure

Facilities accessible to OSG will implement common interfaces to their computational and storage resources and management services. They facilities will contribute to the definition of Grid-wide accounting and policy mechanisms, and support the system infrastructure for security, information and monitoring services.

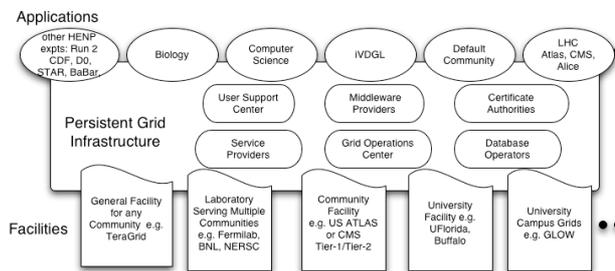


Figure 2: The Open Science Grid Participants

Due to their critical path reliance on OSG infrastructure, US ATLAS and US CMS play an important leadership role in the initial technology roadmap of the Consortium. The global nature of the LHC and other particle physics communities has led us to define a special partnership between OSG and our European peers, the EGEE [21] and LHC Computing Grid Projects [22].

TECHNICAL ROADMAP & BLUEPRINT

The OSG technical roadmap is being defined in a Blueprint, which will evolve to reflect the principles and goals of the Grid infrastructure architecture and design. This Blueprint is being developed with close attention to the architecture and designs of EGEE, LCG, and TeraGrid. The OSG Blueprint will not replicate the information in documents developed by those other groups, but attempt to identify areas of agreement and any potential divergence and gaps.

The keys to the OSG model, as presented in the Blueprint, are: to make the infrastructure both conceptually and in practice as simple and flexible as possible; to build from the bottom up a coherent, performant system which can accommodate the widest practical range of users of current Grid technologies, and to deploy an infrastructure to maximize the use of common in technologies.

The principles as defined in the Blueprint build upon the principles that underlined the success of the Grid2003 project in 2003-2004. Representative principles are:

- OSG will provide baseline services and a reference implementation; use of other services will be allowed.
- All services should support the ability to function and operate in the local environment when disconnected from the OSG environment.
- Users are not required to interact directly with resource providers.
- The requirements for participating in the OSG infrastructure should promote inclusive participation both horizontally (across a wide variety of scientific disciplines) and vertically (from small organizations like high schools to large ones like National Laboratories).

- The OSG architecture will follow the principles of symmetry and recursion.
- OSG infrastructure must always include a phased deployment, with the phase in production having a clear operations model adequate to the provision of production-quality service.

The following are representative Best Practices:

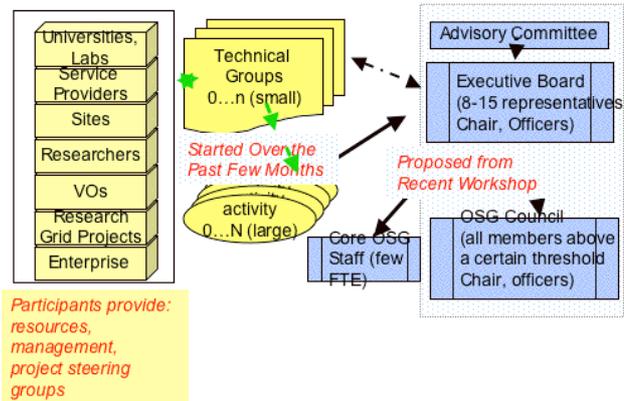
- OSG architecture is VO-based, meaning that most services are instantiated within the context of a VO.
- The infrastructure will support multiple versions of services and environments, and also support incremental upgrades
- OSG infrastructure should have minimal impact on a Site. Services that must run with super-user privileges will be minimized
- OSG infrastructure will support development and execution of applications in a local context, without an active connection to the distributed services.
- Services manage state and ensure their state is accurate and consistent.

The technical approach includes building on our experience with existing and emerging state of the art distributed system technologies and prototypes, not already present in simple Grids:

- To build a scalable (Web-)services fabric capable of supporting thousands of service-instances, integrated with Grid middleware and VO-specific software.
- To use a global view of resources (site facilities and networks) to ensure consistent use and support policies.
- To support a real-time monitoring and tracking system to ensure effective workflow and track resource usage and scheduling: task by task, VO by VO, and in real-time monitor the system state and state-change information.
- To define an architecture that supports complexity and global scalability with moderate effort, including: dynamic, automatic mutual discovery of services, automatic notification of clients, and services of subscribed “events”.

ORGANIZATION

The organization framework of the Consortium includes Technical Working Groups each with a general responsibility for a technical area such as security, storage services and support centers. The work and deliverables for OSG are organized through Activities which are sponsored and overseen by one or more of the technical groups.



The technical groups liaise and collaborate with their peer organizations in the U.S. and world-wide, as well as participate in relevant standards organizations. The Consortium governance is emerging, and will follow the model of an experiment collaboration as far as possible. The following technical groups have either started their work or are being defined:

- *Security* will propose the end-to-end security model and oversee the security infrastructure for OSG.
- *Storage and Data Management* will coordinate those activities related to data storage & management, as well as identify the user and VO requirements and technology solutions.
- *Monitoring, Accounting & Information Services* will address issues such as how resources on the OSG are advertised and accounted for, as well as how the OSG as a system is monitored.
- *Policy Services*, including authorization & access policies of Sites, VOs and Resources.
- *Support Centers*, which includes operations, will define the operational model. This includes the Virtual Data Toolkit support which provides the basic Grid middleware distribution.
- *Governance* will work through the details of participation, decision making, and oversight for the Consortium.
- *Education* will focus on developing education technologies for use on OSG and on training and other educational materials.

DEPLOYING THE OPEN SCIENCE GRID

The initial deployment of OSG is planned for spring 2005. An extended mix of applications is expected to run on OSG including analysis as well as simulation applications from US LHC, D0, CDF, STAR, BaBar, LIGO, SDSS as well as biology applications such as BLAST [23], SnB [24], and GADU [25]. The infrastructure will include an evolution of Grid3 in several areas as outlined below:

- Storage Elements and Services across at least the US LHC Tier-1 and Tier-2 sites.

- An extension of VO management services to include role-based site authorization and access control.
- Access to computational and storage resources at the US LHC Tier-1 centers for OSG and LCG, with access to the resources governed by local site policies.
- Support for SAMGrid.
- Data management, including replica location services.
- An extension of the Grid3 monitoring, testing, and accounting services.
- Support for an OSG VO open to any OSG User. The development guidelines for the activities sponsored by the Consortium are to: Use and integrate existing Grid components and operational procedures; Minimize the impact of the OSG framework, except as needed to ensure consistent operation; Agree among all partners on the wrapping and/or re-engineering of components as needed.

ACKNOWLEDGEMENTS

I acknowledge the many contributions of the members of the Grid3 project, the participants in the OSG technical groups and activities, and the stakeholders and participants in the Open Science Grid Consortium.

REFERENCES

[1] Open Science Grid Consortium Participants
<http://opensciencegrid.org/consortium/partners.html>

- [2] <http://www.ivdgl.org/grid2003/>
- [3] <http://projects.fnal.gov/samgrid/>
- [4] <http://www.usatlas.bnl.gov/>
- [5] <http://www.uscms.org/>
- [6] <http://www.star.bnl.gov/>
- [7] <http://www-cdf.fnal.gov/>
- [8] <http://www-d0.fnal.gov/>
- [9] <http://www.slac.stanford.edu/BFROOT/>
- [10] <http://www.ligo.org/>
- [11] <http://www.sdss.org/>
- [12] <http://www.cs.wisc.edu/condor>
- [13] <http://www.globus.org>
- [14] <http://www.griphyn.org/vdt/>
- [15] <http://www.griphyn.org/>
- [16] <http://www.ivdgl.org/>
- [17] <http://www.ppdg.net/>
- [18] <http://acdc.ccr.buffalo.edu/>
- [19] <http://www.cs.wisc.edu/glow/>
- [20] <http://www.teragrid.org/>
- [21] <http://public.eu-egee.org/>
- [22] <http://lcg.web.cern.ch/LCG/>