Distributed Research Communities, GIS and the Open Science Grid

Ruth Pordes, Fermilab,

Open Science Grid Executive Director

January 2010

What is the Open Science Grid?

Starting from an initial grass-roots effort between NSF and DOE projects around 2001, a collaboration of domain and computer scientists, software developers, and IT came together to form the Open Science Grid Consortium (OSG) in 2005. The Consortium aims to support common end-to-end distributed computing solutions for a broad set of research communities. The OSG provides computing services, software and support for its members. Communities can build and operate their own autonomous Cyberinfrastructures (CIs), selecting which services and/or software they depend on from the OSG. The Consortium has also built a shared, nationally distributed infrastructure to which a community or computing facility can connect its resources. The federated model to which the OSG works allows for each CI to be independent or integrated as much as desired. OSG services and software are currently used by more than twelve research communities, more than seventy universities and laboratories, and more than six campus/regional infrastructures. While mainly in the US, there are several sites in Central and South America, Taiwan, Korea and, most recently, China.

The user communities benefit from the OSG software stack, operational services and support organizations that support high throughput computing jobs and data movement and storage. Typically the users' middleware and algorithmic applications are interfaced to the common client, management and monitoring tools. Each community can further decide if the resources they own are accessible through the OSG infrastructure. If they are made accessible, the owners maintain control of the management and policies for their use and sharing.

The OSG Consortium includes several of the large physics collaborations – including the STAR nuclear physics experiment, the LIGO Laser Interferometer Gravitational Wave Observatory, collaborations at the Large Hadron Collider at CERN, and the running Tevatron and Intensity Frontier high energy physics experiments – as well as the Cyberinfrastructure for Geographic Information group at the University of Illinois Urbana Champaign, Structural Biology Grid based at the Harvard Medical School, a Computational Biology group at the University of Connecticut, and an initial collaboration with the GlueX experiment at the Thomas Jefferson National Facility.

The OSG also supports its own community (Engage) that provides individual, small user groups and new communities support to adapt and run their applications, thus lowering the barrier to use of this typically unfamiliar computational environment - different from that of a local cluster. For example, since the resources accessible through the OSG are autonomous and fully distributed there is expected, guaranteed average, but non-deterministic availability of and throughput from individual resources. Also the resources are heterogeneous and dynamically gathered information is needed to determine which resources can best support any particular application. An average of 250,000 jobs a day are run through the OSG across more than 30 sites (with a peak of 350,000 jobs over the past few months). More than 25% of the throughput is "shared" - running on resources not owned by the user. While the LHC community's use is likely to ramp up over this coming year, our recent analysis shows there is sufficient headroom for the continued ramp up of other applications given that the policy of sharing is maintained.

The OSG Virtual Data Toolkit (VDT) provides a configurable middleware stack for OSG users, site and community administrators. OSG also supports the VDT for non-OSG use by the European Grids for EscienceE (EGEE), UK National Grid Service (NGS), Australia's APAC and other projects. The VDT provides a collection point for the integrated building and regression testing of the more than sixty software components across for the more than fifteen flavors of Linux in use. OSG provides packaging and distribution of collections of the toolkit as needed by the user communities. The OSG Integration Testbed provides access to ten distributed sites for system testing of the integrated software collections by the applications as well as internally by the OSG.

Domain and computational scientists work together to solve the end-to-end application needs based on advancing the principles, methodologies and frameworks of large-scale distributed computing. The teams apply the conceptual insights gained on the distributed infrastructures, to the benefit of all Consortium members. Over the past few years the OSG has made advances in: "resource overlay" job management technologies that allow transparent use of heterogeneous clusters as a uniform distributed facility; maintaining and evolving a production high-availability infrastructure; use of remotely accessible storage caches; dynamic sharing of the resources; and use and management of independent CIs in a federated model which allows for scaling of the whole system.

Cyberinfrastructure and Geospatial Information (CIGO) VO on OSG

The CIGI VO consists of about a dozen institutions with an international scope. Sites at NCSA and in China are accessible to the OSG and test jobs of GISolve have been run. The application goals are to characterize the interactions between human and natural environments, and understand geographic patterns and processes for the assessment of environmental and social impacts of global climate change. During 2010, CIGI will complete the interface to use OSG resources through the GISolve Toolkit and will increase the number of its' applications deployed.

Geographic Information Systems and Condor

OSG applications access the distributed resources through Condor client tools and many of the currently accessible resources run the Condor batch system. Several GIS applications use Condor distributed computing locally to manage workflows and/or process images from large datasets acquired by their sensors. Examples are: the Ordnance Survey Ireland; the interface of the Aerial Digital Sensor (ADS40) from Leica Geosystems GIS & Mapping (used by companies including the Hanjin Information Systems & Telecommunication CO., Ltd in South Korea); and the University of Newcastle in the UK for geodetic GPS processing research, running the analysis packages GIPSY, Bernese and GAMIT.

OSG and Its' Future

The OSG has had an effective partnership with the European Grid projects over the past ten years in support of many of its current customers. The European transition to a federated infrastructure – matching the architecture long promoted by OSG itself – will provide new opportunities for collaboration. The OSG increasingly collaborates on joint activities with TeraGrid. Several science gateway applications already support use of both OSG and TG and tests are ongoing for additional workflows of earthquake simulations from SCEC to use resources from both infrastructures.

The OSG will evolve to meet the needs for new capabilities and capacities of the current users - with more than a factor of two increase in throughput needed over the next couple of years. The OSG also continues to expand the range of applications supported as new customers approach us. Our multi-disciplinary teams continue to work on solving the research and operational challenges to improve the quality and usability of the technologies and methodologies. The OSG continues to work closely with computer science groups, software developers and computational scientists - providing requirements, testing, and feedback for software developed and used for distributed computing.

Such activities include: software engineered life-cycle processes and sustainability; improved methods and technologies for maintenance and protection of the production distributed systems in the face of service and software evolution, resource over-subscription and security incidents; provision and use of shared distributed storage and data management; end-to-end identity management integrating web-based collaborative tools with data access and job management; configuration management in a heterogeneous resource environment; extension of the job-overlay methods for more effective workflow management and throughput; increased support for small-parallelism applications on multi-core computers; preparing for advanced network technologies, such ESnet's 100 Gigabit testbed; support for commercial and scientific clouds; and provision of virtualization techniques for our high data volume and high processing throughput applications.