

Open Science Grid

An Introduction

Overview

The goal of the Open Science Grid (OSG) is to transform processing and data intensive science through a cross-domain, self-managed, nationally distributed cyber-infrastructure that brings together community resources and enables effective computational resource sharing at the academic and research campuses. The OSG promotes science by enabling a framework of distributed computing and storage resources, a set of services and methods that enable better access to ever increasing computing resources for researchers and communities, and principles and software that enable distributed high through-put computing (DHTC) for users and communities at all scales. In addition, OSG represents and promotes these concepts and technology with other international partners on behalf of US science collaborations including the US-LHC experiments.

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; and, promote and support better utilization of local campus resources and improved access to the national cyber-infrastructure from those campuses.

To meet these goals, OSG provides: 1) a set of operational principles and methods that organize users and resources into effective and productive DHTC communities; 2) consulting and support for stakeholders including a set of common computing services; and, 3) integrated and tested software platforms which enable building of distributed computing environments at both the national and campus level.

Principles and Methods

Researchers at campuses and in collaborative communities who leverage shared computing resources are a key aspect of the Open Science Grid. This spans communities at all scales from a handful of researchers at an academic campus to the very large experiments.

An important concept in OSG is Virtual Organizations (VOs) which are a collection of researchers who join together to accomplish their goals; typically they share the same mission, but that is not a requirement. A VO joins OSG to share their resources, computing and storage, with the other VOs and to be able to access the resources provided by other communities in OSG as well as share data and resources with other peer and international computer grids. The resources owned by a VO are often geographically distributed; a set of co-located resources are referred to as a site and thus a VO may own a number of sites. There are two important sides of VOs: 1) the

user community within a VO 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.

At academic campuses and research institutions, OSG promotes and enables easier and increased access to computing by providing methods and technologies for resource sharing; this spans sharing within the campus, egress out to the national CI for access to additional computing, and ingress to shared campus resources from the national CI. Quite often, the computing clusters are dedicated to the mission of a particular research group at a campus and their users have exclusive access to these resources; this arrangement has many benefits that are related to control and ease of access to local data sets but the downside is the ability to access additional resources for peaks and uneven utilization of the local resources over time. OSG works with campuses to federate their resources into local DHTC frameworks and to enable sharing. Because of the intra-campus nature of this federation, certain authentication and security frameworks are often simplified. When users adapt their computation jobs to benefit from the local DHTC framework, it is a relatively small step from there to access the national CI leveraging OSG, XSEDE, and other partner grids.

A key principle in OSG is the autonomy of communities that bring their computing resources into the DHTC environment; this allows them to develop operational models that best meet the needs of their users. This autonomy applies both to how their users operate and how their computing resources are configured. Communities administer, manage, and support their own users and many communities provide common software infrastructure designed to meet the specific needs of their users. Providers of resources also have great autonomy in building and operating their sites; they use the OSG software stack to provide the “middleware layers” that make their sites ready for shared use and/or connection to the OSG. Sites set policies on how their resources will be used by their own users and other communities; the only requirement is that sites support at least one other community but the site controls the conditions under which their resources are shared. OSG does not restrict what hardware or operating system a community may supply or what software it may use to access OSG or provide resources on OSG; they are allowed to make such choices as long as they meet the minimum DHTC requirements. This autonomy allows communities and VOs to structure their computing resources to best fit their own specific needs and makes it more likely that a community will choose to join the national DHTC fabric because it doesn't have to compromise its own needs to do so.

Consulting, Support, and Common Services

The OSG provides a suite of services which broadly fall in three major categories: 1) education, training, and consulting in the best practices of DHTC; 2) technical support to users and communities in using and establishing distributed computing environments; and, 3) operation of common computing services which are operated by the OSG staff for the benefit of VOs and Campuses.

OSG offers education and hands-on training for students and researchers to learn about DHTC and provides consultation and support to assist communities, both campuses and research groups, in learning DHTC and implementing this framework in their own environment. We provide consulting services for the broader community that builds and operates their own distributed computing environments (e.g. LIGO) and support the specific DHTC environment within the OSG fabric. In addition, OSG User Support assists researchers new to OSG in adapting their applications to run in a distributed computing environment. Members of the OSG Campus Grids team reach out to campuses to communicate the benefits of federated resources and provide technical guidance and assistance in implementing campus grids that inter-operate effectively with the national CI.

To enable the work of the community, the OSG provides direct staff support and operates a set of common computing services for the benefit of communities that adopt the DHTC framework. These functions are available to all communities in OSG and provide a foundation for the specific environments built, operated, and supported by them; these include:

- Operational support including centralized ticket handling,
- 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,
- Site coordination and technical support for VOs and Campuses to assure effective utilization of grid connected resources,
- End-to-end support for simulation, production, and analysis to enable the science communities accomplish their goals.

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

Integrated Software Platforms

OSG provides software solutions that have been integrated and tested to enable communities to build campus grids or join the OSG with computing resources or as users. The primary goal of the OSG software effort is to build, integrate, test, distribute, and support a set of software solutions that meet the needs of user communities who wish to implement the various aspects of the DHTC vision; communities are free to select subsets of the offered software solutions that best enable their own plans for distributed computing. 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.

Software components (developed completely independently, many resulting from NSF and DOE supported programs) are evaluated, built, tested, configured, documented and packaged into a coherent suite of software solutions that researchers depend on. We undertake as little new software development as possible, to minimize ongoing local support, and focus on integrating existing tools from external software providers into a coherent whole. The goal is to provide and support a dependable, trusted, easy to deploy and low cost of ownership software infrastructure for DHTC that meets the diverse, evolving needs of the OSG stakeholders.

The key to making the OSG infrastructure work is a common package of software solutions provided and supported by OSG, sometimes 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 as 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 software. This OSG middleware allows the VOs and Campuses to build operational environments customized to their needs.

The OSG supports a varied set of Linux based operating systems and versions and provides tools that understand and adapt to the functionality that is available on each resource. These capabilities shield the users from much of diversity in the grid resources and allow the users to dispatch work to uniform interfaces that understand and manage the differences in these distributed resources. Through installation of the VDT and related packages, users and administrators operate in a well-defined environment and set of available services. Thus, scientists are able to send their jobs into a virtual computing cluster for execution; the DHTC fabric deals with the complexities of the distributed resources and manages the dispatch of jobs to remote sites, execution, and return of results to the user.

OSG Usage and Dimensions - June 2011

DHTC technology created and incorporated by the OSG and its contributing partners has now advanced to the point that scientific users are utilizing more simultaneous resources than ever before. Typical VOs now utilize between 15 and 20 resources with some routinely using as many as 40 – 45 simultaneous resources. The overall usage of OSG has increased by greater than 50% in the last year and continues to grow. Utilization by each stakeholder varies depending on its needs during any particular interval. Overall use of the facility for the 12 month period ending June 2011 was 424M hours compared to 270M hours for the previous 12 months period ending June 2010. During stable normal operations, OSG now provides over 1.4M CPU wall clock hours a day (~56,000 cpu days per day) with peaks occasionally exceeding 1.6M hours a day; approximately 300K – 400K opportunistic hours (~30%) are available on a daily basis for resource sharing.

While OSG does not own the computing, storage, or network resources used by the scientific community, these resources are contributed by the community, organized by the OSG facility, and governed by the OSG Consortium. Current OSG resources are summarized below:

Number of Grid interfaced processing resources on the production infrastructure	131
Number of Grid interfaced data storage resources on the production infrastructure	61
Number of Campus Infrastructures interfaced to the OSG	9 (GridUNESP, Clemson, FermiGrid, Purdue, Wisconsin, Buffalo, Nebraska, Oklahoma, SBGrid)
Number of National Grids interoperating with the OSG	3 (EGI, NGDF, XSEDE)
Number of processing resources on the integration testing infrastructure	28
Number of Grid interfaced data storage resources on the integration testing infrastructure	11
Number of Cores accessible to the OSG infrastructure	~70,000
Size of Disk storage accessible to the OSG infrastructure	~29 Petabytes
CPU Wall Clock usage of the OSG infrastructure	Average of 56,000 CPU days/day during May 2011