

GridConnections

November 2007

News and Information for the Open Grid Forum Community

In this Issue:

[SC07 Preview – Interoperability Demonstrations](#)

[OGF21 Wrap Up and Survey Results](#)

[OGF and the Open Geospatial Consortium \(OGC\) Ink MoU](#)

[SAGA Tutorial at Grid2007](#)

[Documents Update – New Publications and Documents in Public Comment](#)

[Upcoming Events – OGF22, Globecom, and Management Developers Conference](#)

[Newsletter Contributors Needed](#)

[Membership Information – Budget Planning for 2008](#)

SC07 Preview

SC07 kicks off November 11 in Reno, Nevada. With HPC being the largest and most mature application for Grid technologies, the OGF community will have a very large presence at SC this year. First and foremost are the interop demonstrations planned many of which are described below:

High Performance Computing Profile (HPCP)

Demonstrations will take place in multiple booth locations

HPCP is a proposed standard for grid interoperability in HPC environments that references existing specifications including OGF Job Submission Description Language (JSDL), OGF OGSA Basic Execution Service (BES), and WS-I Basic Profile. HPCP allows application and middleware software providers to target multiple HPC systems via a single protocol leveraging standard Web services protocols and development environments. The demonstration shows interoperability between multiple 3rd party resource managers and Web services platforms using the HPC Profile. Tasks can be submitted from clients located in one booth to compute clusters located elsewhere via the HPC Basic Profile specification. The demonstration includes the submission of tasks to a resource manager, the retrieval of a task's execution status and the retrieval of information about an HPC system's resources. The demonstration shows how organizations can benefit by integrating commercial and open source products into existing and new HPC systems, interfacing with HPC systems at remote locations, and by leveraging generic HPC utilities.

Grid Interoperability Now (GIN)

Demonstrations will take place in multiple booth locations

GIN coordinates efforts among production Grids interested in interoperating in support of applications that require resources in multiple Grids. GIN plans and implements interoperation using solutions for which there are working implementations available leading to a more seamless usage of different Grid infrastructures by applications.

Grid Storage Management (GSM)

Lawrence Berkeley National Laboratory Booth #351

Large scale Grid computing requires dynamic storage allocation and management of large numbers of files. However, storage systems vary from single disk pool to complex mass storage systems. A standard middleware specification for dynamic storage reservation and management of the files in the reserved spaces has been developed over the last seven years - referred to as the Storage Resource Management (SRM) specification. This demo will show the interoperability of different SRM implementations around the world based on the specification. It will show the ability to put, get, and copy files between any of these storages systems using the uniform SRM interfaces. Many of these SRM-fronted systems are now used in large Grid projects, including the high energy physics Worldwide LHC Computing Grid (WLCG) Project, Open Science Grid (OSG), and the Earth System Grid (ESG) project. Contact Alex Sim (asim@lbl.gov) or Arie Shoshani (shoshani@lbl.gov) for information.

Grid Remote Procedure Call (GridRPC)

AIST Booth #765

GridRPC is a programming model based on a Remote Procedure Call (RPC) mechanism tailored for the Grid. The programming model provided by GridRPC is that of standard RPC plus asynchronous, coarse-grained parallel tasking. The GridRPC API specification clearly and unambiguously defines the syntax and semantics for GridRPC, thereby enabling a growing user base to take an advantage of multiple implementations. This demonstration shows how to develop and execute flexible, robust, and efficient Grid applications using Ninf-G, a reference implementation of the GridRPC API. The demonstration is targeted at application developers who are interested in developing large-scale sustainable Grid applications. Experimental results of large-scale molecule simulations on AIST-TeraGrid will also be presented. Contact Yoshio Tanaka (yoshio.tanaka@aist.go.jp) for information.

Grid Information Retrieval (GIR)

KISTI Booth #2933

Grid Information Retrieval is for distributed search in a grid environment. The main elements of GIR are query processors, indexing services, and collection managers. GIR is intended to provide access to relatively small collections which are not necessarily accessible to public Web crawlers. Because the GIR collections can be based on different information retrieval software, the search experience can be more highly customized for particular users, searches, or data types. Members of the OGF's Grid Information Retrieval Working Group (GIR-WG) will give a presentation and discussion of a new GIR prototype on XX at YY. Information retrieval will take place between several systems at KISTI (Korea) and the Arctic Region Supercomputing Center (Alaska). Contact Greg Newby (newby@arsc.edu) or Yangwoo Kim (ywkim@dongguk.edu) for information.

Resource Namespace Service (RNS)

Center for Computational Sciences, University of Tsukuba Booth #2651

RNS provides general hierarchical namespace that manages name-to-resource mappings. It is extensible and can be used by Grid file system, execution management, and various application domains. Contact Osamu Tatebe (tatebe@cs.tsukuba.ac.jp) for information.

Information Dissemination (INFOD) and Simple API for Grid Applications (SAGA) are also planning demonstrations, but have not yet submitted information.

Other SC07 Activities

OGF Sponsor XtreamOS invites all interested developers and users to the Grid Operating Systems Community [BoF](#) at SC07 on Tuesday, Nov 13, 2007 from 12:15 - 1:15PM

to discuss the current state as well as future directions for grid operating systems and middleware. Grid middleware is typically built on top of existing operating systems but little has been done to extend the underlying operating systems to enable and facilitate Grid computing. In June 2006 the EC-funded project XtreamOS started to develop a Linux-based grid operating system to directly support virtual organizations for next generation grids.

OGF21 Wrap Up and Survey Results

Over 250 attendees from around the globe gathered in Seattle, Washington for OGF21 which was held October 15 - 19, 2007. This event was mainly focused on the Grid standards creation, publication and adoption work of OGF's chartered groups. A significant trend clearly seen at OGF21 was the added emphasis on demonstrating interoperability and documenting use cases and reference implementations. For example, the Distributed Resource Management Application API (DRMAA) and Grid Remote Procedure Call (GridRPC) teams were recognized for their recent accomplishment of achieving 'Grid Recommendation' status for their respective standards which signifies that these standards are being adopted within the distributed computing community.

Another example of this trend was the large number of OGF interoperability demonstrations discussed at OGF21 which are being planned for the SC07 conference being held in Reno next month as discussed in detail above.

Yet another example was the popularity of the 15 software and solutions sessions which enabled open-source and commercial software providers to dialogue with potential new and existing users. The software providers, many of whom have implemented OGF specifications, also participated in a round-table discussion on critical requirements to the interoperation of production grids around the world, as identified by OGF's GIN Working Group. This special session entitled "Software Providers meet GIN and Standards," facilitated communication between these two groups and identified missing components in the proposed and planned standards. The software providers expressed a desire to make reference implementations of emerging standards available to their users and the GIN and Standards participants were interested in shaping their efforts based on users' feedback. One of the major issues of discussion was that of agreement on security mechanisms and all agreed that it was critical to develop and document reasonable security profiles to promote interoperability. Standards relating to data movement across multiple grid deployments were another key discussion point and the panel advocated looking toward the efforts in OGSA Data Movement Interface (DMI) working group for leadership in this space. Ravi Madduri, from the Globus Alliance and Argonne National Laboratory, commented that this was "one of the best sessions I have attended in OGF and I think these are the types of sessions that we need more of."

OGF21 also marked the beginning of Craig Lee's 3 year term as OGF president. Craig opened the meeting with a presentation that discussed, among other things, how OGF can nurture its current constituency, while expanding and evolving the organization. He outlined OGF's goal to build participation in the research/academic community by promoting peer-reviewed publications related to work enabled by OGF; by co-locating future OGF events with established research grid conferences; and by collaborating closely with other organizations such as the Open Geospatial Consortium which recently signed a Memorandum of Understanding with the OGF.

Attendees at OGF's second Web 2.0 and Grid workshop saw a wide range of applications, from "blogging the lab" to sensor networks, all building on Web 2.0 and Grid technologies. Organized by OGF's eScience function, the two day workshop featured project presentations, future visions and an expert panel, drawing expertise from Grid, Web and Social Science communities. Under the theme of Web 2.0 infrastructure and applications, Savas Parastatidis

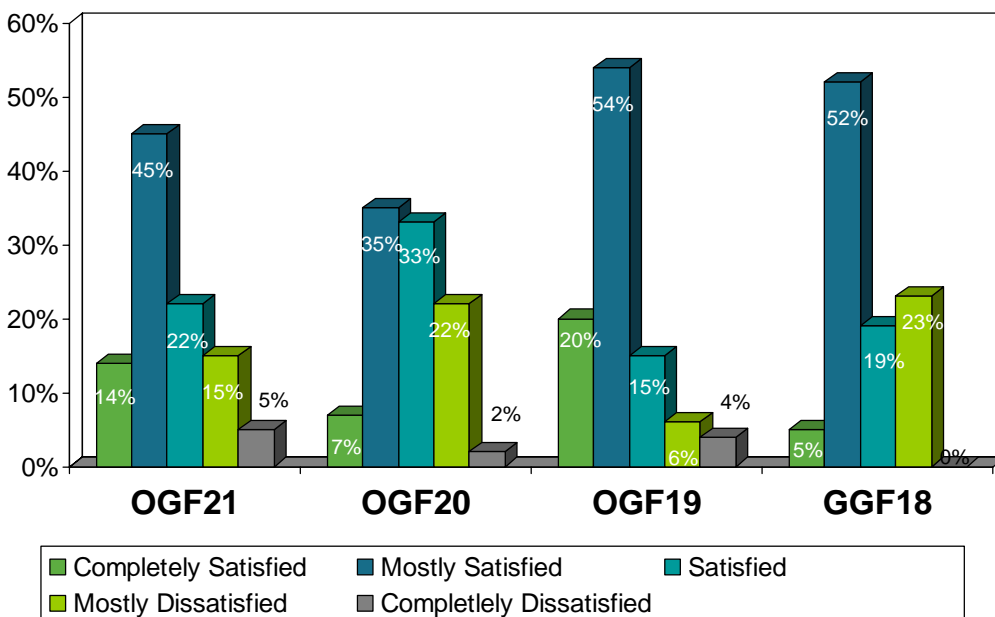
from Microsoft talked about "the cloud" and a possible paradigm shift away from research institutions maintaining their own infrastructure. A visionary talk by Geoffrey Fox from Indiana University (and OGF VP of eScience) described a future where Web 2.0 technologies make it easy to combine specialist software components to fully exploit the capabilities of multicore. Other talks included two social web sites for eScientists - myExperiment, which is being used by bioinformaticians and chemists, and ourSpaces for social scientists

OGF21 also focused on enterprise applications. The OGF21 enterprise program featured a special full day track on Grid in the IT data center. Sessions included 'Understanding Grid in the Distributed Computing Landscape', 'Grids and Virtualization', and "Data/Compute Affinity – Focus on Data Caching'. Presentations were made by representatives from Intel, Oracle, Boeing, Platform Computing, SAS, XenSource, Gemstone, and Gigaspaces. All of these sessions were well attended, but, as you might expect, the Grid and Virtualization session was by far the crowd favorite. A common theme from the session was that grid is a complementary technology to, and natural extension of, current virtualization solutions. In other words, the virtualization systems now being deployed are laying the groundwork for future extension into grid.

OGF21 was sponsored by Microsoft, Pacific Northwest Gigapop, HP, IBM, OMII Europe, Univa UD, Fraunhofer Grid Alliance, and GridToday.

OGF21 Event Survey Results

The results of the OGF21 event survey are in. A total of 64 attendees completed the survey. The following chart indicates the overall satisfaction that respondents had with the event and compares it to the last 3 previous events. The full set of OGF21 survey results can be viewed [here](#).



OGF and Open Geospatial Consortium (OGC) Ink MoU

OGF recently signed a memorandum of understanding with The Open Geospatial Consortium (OGC). OGC is the international standards organization for geospatial

(mapping) information and solutions that "geo-enable" the Web, wireless and location-based services, and mainstream IT. OGC has a suite of tools for managing and presenting geospatial data and wants to extend their tools with the capability for distributed resource management, i.e., grids. The initial goals of the collaboration include:

- 1.) Integrate OGC's Web Processing Service (WPS) with a range of "back-end" processing environments to enable large-scale processing. The WPS could also be used as a front-end to interface to multiple grid infrastructures, such as TeraGrid, NAREGI, EGEE, and the UK's National Grid Service. This would be an application driver for both grid and data interoperability issues.
- 2.) Integration of WPS with workflow management tools. OGC's SAGA draft standard is where multiple WPS calls could be managed.
- 3.) Integration of OGC Federated Catalogues/Data Repositories with grid data movement tools. OGC's GridFTP is one possibility that supports secure, third-party transfers that are useful when moving data from a repository to a remote service.

However, the real goal is not just to do science, but to greatly enhance things like operational hurricane forecasting, location-based services, and anything to do with putting data on a map. WPS is just a starting point for the collaboration. As the two organizations engage and build mutual understanding of technical requirements and approaches, many other things will be possible.

SAGA Tutorial at Grid2007

At the Grid2007 conference last month in Austin, TX, the Center for Computation and Technology (CCT) based group hosted a SAGA half day, hands-on tutorial helping interested people get up to speed with the C++ implementation. SAGA describes simple API's for commonly needed operations for file transfer and access, replica management, stream based interprocess communication, job submission and monitoring, and client side RPC calls. Several groups worldwide currently work on implementations for this specification, most notably the work on a Java based implementation done at Vrije Universiteit Amsterdam or the C++ reference implementation done at CCT in Baton Rouge, Louisiana, USA. Efforts have also begun to enable SAGA to interface with Platform LSF and Condor.

The tutorial was split into two parts. The first hour was dedicated to a high level overview about SAGA, the SAGA specification, and first SAGA related applications developed using the C++ implementation. The rest of the workshop was oriented to hands-on experience, where the participants were asked to complete an existing skeleton application with a couple of missing SAGA API calls. The idea was that all participants had to write about 20-30 lines of code, getting a really useful SAGA based Grid shell application in return. The tutorial, and a similar workshop held two month ago in Karlsruhe/Germany, attracted over 30 total participants primarily from industry, underlining the interest the community has in being able to use simple, standardized interfaces to the Grid.

Documents Update

Recently Published Documents

The following document was published in October. Congratulations to Ellen and the working group members involved in getting this important work accomplished!

DOCUMENT	TITLE	TYPE	AUTHORS	AREA
GFD.119	Execution Environment and Basic Execution Service Model in OGSA® Grids	Informational	E. Stokes	Architecture

Documents in Public Comment

Prior to formally publishing a document, OGF solicits “public comments” from the greater grid community, which is an important step in the OGF document process. The following documents are currently available for public comment. Please take a moment to provide your feedback.

[2nd International Workshop on Campus and Community Grids](#)

[Open Grid Services Architecture® Glossary of Terms version 1.6](#)

[Open Grid Forum Document Process and Requirements](#)

[The Storage Resource Manager Interface Specification Version 2.2](#)

Upcoming Events

OGF22 Cambridge, Massachusetts February 25-29, 2008

OGF launches its second full year of events with OGF22. We recently completed a very productive 2007 with accelerated document publication and new collaborations initiated. OGF22 will feature a series of Life Science/Pharma and Financial Services sessions to explain the opportunities and challenges facing these industries as they continue to leverage grid and related technologies. OGF is also planning a full day workshop exploring data management issues in both research and industry, with help from experts from the Storage Networking Industry Association (SNIA) and others. The Call for Participation for these and other topics will be launched soon.

- **OGF23 Barcelona, Spain June 2-6, 2008**
- **OGF24 Co-located with GridAsia08 Singapore September 15-19, 2008**

Other Events

- [Globecom 2007 Washington, D.C., November 26-30](#)

This is IEEE's 50th Anniversary of its Globecom Conference which spans voice, data, image, and multimedia communications technologies. OGF is organizing two panels for Tuesday Nov. 27: “Grids and Service-Oriented Networks” and “Grids, Networks and Virtualization”. This represents a great opportunity for OGF to expand contacts and engagement within the telecom community.

- [Management Developer's Conference December 3 - 6, 2007 Santa Clara, CA](#)

The Management Developers Conference is dedicated to Standards Based System and Network Management Technologies. OGF is once again an organization sponsor for this event and we will have a presentation on Tuesday, December 4 by Sergio Andreozzi on Implementing the OGF GLUE Information Model.

Newsletter Contributors Needed

The purpose of the OGF GridConnections newsletter is to inform and educate the greater grid community about our activities and accomplishments. If you have any news you would like to submit for the newsletter, please do not hesitate to do so. You, our members, drive all of the significant events, activities and accomplishments of our community and we would love to hear from you. Just send an email to the [GridConnections editor](#). We welcome your input!

Membership Information – Budget Planning For 2008

Many organizations do their budgeting for the upcoming year during this time. Please remember to allocate money for your 2008 OGF membership renewal. We have been making great progress in 2007 and we need your continued support to build on that momentum. Thanks!