OpenGridForum

OPEN FORUM | OPEN STANDARDS

GridConnections

January 2008

News and Information for the Open Grid Forum Community

In this Issue:

Message from Craig Lee, OGF President

As we enter the new year, I want to take a moment and thank you for your support, reflect on some of the highlights of the past year, and make you aware of some high-priority activities in progress.

Over the course of 2007, OGF continued to deliver many high-quality specifications that detailed standards-based approaches to essential grid interoperability issues. During the recent Supercomputing '07 conference, ten separate OGF specifications or specifications in progress were demonstrated in numerous show booths. These demonstrations showed how implementations of OGF specifications enabled interoperability across heterogeneous platforms - the ultimate goal of standards.

Another very significant milestone in OGF's history was announced and highlighted during the show. Microsoft, Altair, and Platform Computing announced that OGF's High Performance Computing Basic Profile (HPC-BP) would be scheduled into future releases of their software, enabling customers to enjoy the benefits of interoperability in a multi-vendor environment. This, plus the intent to use HPC-BP in numerous national and regional production grids serving eScience communities around the world, made the announcement a significant accomplishment for OGF.

Also this year, OGF hosted its best-attended event ever during OGF20 in Manchester, UK. This event, which was co-located with the Enabling Grids for E-SciencE (EGEE) User Forum, attracted over 900 attendees and resulted in grid users from research and commercial verticals productively interacting with those chartered to create best practices, standards, and solutions.

Meanwhile, data challenges in both commercial and research environments prompted OGF to increase formal engagement in two high-profile activities. First, OGF signed a MoU with the Open Geospatial Consortium (OGC) whose charter is to manage and distribute large quantities of geospatial data (i.e., anything that goes on a map). This initiative, along with a joint initiative with the Storage Networking Industry Association (SNIA) around developing standards and best practices in data movement and data management, has brought data to the forefront of our activities moving forward.

Building on these successes in standards adoption, events, and organizational engagement across the distributed computing landscape, OGF must also strive to explore promising research topics and emerging technologies that are important to the future of grids and our community. Clearly, data interoperability, virtualization, Web 2.0, muliticore processors, cloud computing, and service-oriented networks are all topics highly relevant to the evolution of the grid concept. There is still much work to be done to make grid capabilities easy to adopt and effective to use across the spectrum of application domains and requirements.

Finally, I cannot stress how much of our success is based on our member's financial and technical engagement. My sincere thanks to you all once again. People from all over the world depend on OGF to provide them an open forum for engagement and open standards for interoperability. As OGF continues to deliver these essential services, we hope to enjoy your continued participation. I am delighted to be at the helm of this great organization. Please do not hesitate to contact me with any questions or comments.

SC07 Report - OGF Grid Remote Procedure Call Demonstration

Yoshio Tanaka, Open Grid Forum GridRPC Working Group Secretary

Sustainable Multiscale Simulations using Grid Remote Procedure Call

In scientific Grid environments, you often have many supercomputers or clusters distributed geographically across the world processing problem sizes that are not constant and may vary during the execution. In such a scenario, how do you write application programs to effectively use the capacity of computing resources provided by thousands of processors for a long period of time like one month or even one year? One way is to ensure your programming model satisfies the following requirements:

- Flexibility in resource allotment: the application must be written to request additional computing resource on-the-fly according to the availability
- Fault tolerance and robustness: the application must be able to automatically recover from cluster node or interconnection failures
- Scalability in increasing the number of parallel activities: the application must be able to manage a large number of parallel activities effectively.

Message Passing Interface (MPI) is one well-known programming model for parallel applications, however simply extending MPI to Grid environments is not the answer. That's because a MPI program will terminate when one of the MPI processes fails and it is hard to allocate and execute many resources across clusters simultaneously using MPI.

GridRPC, a natural extension of remote procedure call (RPC) to the Grid environment, is the unique solution available that meets the critical requirements needed to write large scale applications for Grids. When viewed at a very high abstraction level, the programming model provided by GridRPC is that of standard RPC plus asynchronous, coarse-grained parallel tasking. At a more practical level, GridRPC provides a simple, yet powerful, client-server-based framework for programming on the Grid.

In September 2007, "A GridRPC Model and API for End-User Applications", which presents a model and API for GridRPC, was approved as a "Grid Recommendation" (or Standard) by the Open Grid Forum (OGF). "Grid Recommendation" status can only be achieved after a "proposed recommendation" has at least two documented interoperable implementations.

Last month at the SuperComputing 2007 show, the use of GRIDRPC for implementing large-scale, long-run, and adaptive Grid applications was demonstrated. In addition, simulation results of a relatively long H-diffusion path in an inhomogeneously nano-stressed gamma-alumina was presented. Utilizing 1129 processors distributed on Trans-Pacific Grid infrastructure consisting of supercomputers in TeraGrid and AIST, the simulation was executed for approximately 60 hours. Regardless of the system faults, the simulation was

continued by recovering from such faults and there was no manual restart of the simulation. The experimental results prove that GridRPC is a promising approach to implementing sustainable Grid applications.

The results of other large-scale empirical experiments have been shown at past Supercomputing conferences. Large-scale simulation of corrosion of Silicon under stress was presented at SC2004. Utilizing 1793 processors in AIST and Pittsburgh Supercomputer Center (PSC), the simulation ran for approximately 12 hours. In 2005, result of the larger-scale simulation of SIMOX (Separation by Implanted Oxygen) was presented. A joint research group by AIST, Nagoya Institute of Technology, Japan, and the University of Southern California (USC), achieved an automated execution of multi-scale simulation on a Grid consisting of 6 supercomputer centers in Japan (AIST, the University of Tokyo, and Tokyo Institute of Technology) and the US (PSC, National Center for Supercomputing Alliance, and USC), in which the number of processors changes dynamically on demand and resource are allocated and migrated dynamically according to both reservations and unexpected faults. The result of the 19 day experiment showed that GridRPC is able to implement a flexible (allow dynamic allocation and migration of resources), robust (detect faults and recover from such errors), and efficient (utilize thousands of processors) Grid applications.

The standardization of the GridRPC API will potentially make Grids easily accessible to thousands of users that are already familiar with using the RPC model for their scientific applications. In fact, this standard has already enabled a growing user base such as Ninf-G, developed by the National Institute of Advanced Industrial Science and Technology in Japan; GridSolve, developed by the University of Tennessee Knoxville; and DIET, developed by Institut National de Recherche en Informatique et Automatique in France to take advantage of multiple implementations.

The OGF GridRPC working group is now tackling the issues of service discovery, introspection, persistent data, and workflow management and their proper expression within GridRPC. For information, visit the OGF or Ninf websites.

SC07 Report - Internet2 and Collaborators Demonstrate OGF Standards

At SC07, Internet2, together with several partners and collaborators, demonstrated for the first time interoperability of its Dynamic Circuit (DC) Network with multiple regional and international networks as well as an equipment provider. The interoperability of the Inter-Domain Controller (IDC) protocol, being developed in a collaboration between GÉANT2, Internet2, and ESnet, is made possible through emerging network markup language (NML) standards being developed in the Open Grid Forum.

“For many years, Internet2 has been a contributor to the OGF standards process and we believe its plays a critical role in engaging and advancing the broad Internet development community,” said Eric Boyd, Internet2 deputy technology officer. “It’s gratifying to see the results of our community’s efforts instantiated in new standards-based software that will provide immediate benefits to our members and the global networking and grid community.”

The Internet2 demonstrations showcased interoperability with ESnet, GÉANT2 in Europe, NYSERNET in New York, the Great Plains Network (GPN), GRNET in Greece, HEAnet in Ireland, Merit Network, Northern Crossroads (NoX), a Nortel Network based in Ottawa, Canada, the PIONIER network in Poland, and the Phosphorus testbed at the University of Amsterdam via SURFnet’s NetherLight GLIF Open Lambda Exchange in Amsterdam.

In doing so, the SC07 demonstration marked a major first step in enabling the widespread adoption of dynamic circuit networking by showcasing how networks with different equipment, network technology, and allocation models can dynamically provision dedicated circuits across domains.

Dynamic circuit networks use web services to provide on-demand or scheduled dedicated point-to-point bandwidth in the spirit of grid computing to enable the most intensive applications for research and education to allocate the resources that they need when they need them. These include applications like massive terabyte-sized data transfers that will be critical for the high energy physics community when the Large Hadron Collider (LHC) comes online next year.

In addition to the NML standards, Internet2 together with several collaborators also demonstrated the perfSONAR framework which is built on standards developed under the direction of the OGF Network Monitoring Working Group (NMWG). The core of the perfSONAR framework comprises a set of open protocol standards for interoperability between measurement and monitoring systems. Overall, perfSONAR can be viewed as a set of open source web services that can be combined and extended to create a real-time performance-monitoring framework. As the protocols and models are open, deployments may choose between various implementations and may deploy as much or as little of the infrastructure as desired.

At SC07, perfSONAR was utilized to provide performance data on the Internet2 dynamic circuit network demonstrations. Because it is designed to understand the underlying topology of a network, including the complex hybrid IP/circuit network topologies that exist on the Internet2 Network, perfSONAR collected and reported performance data on the IP and dynamic circuit network links and provided visualization of the multiple gigabit data flows.

perfSONAR's ability to share network diagnostic information across administrative domains makes it an important tool for projects like the LHC, and other collaborations between multiple national and international organizations. perfSONAR simplifies the troubleshooting and evaluation of performance issues across networks—allowing network administrators to diagnose and verify network problems across complex multi-domain topologies in near real time.

Jeff Boote, senior network software engineer for the Internet2 Performance Architecture Team, noted "Developers of grid applications have long sought mechanisms to enable consistent and accurate feedback about network performance. perfSONAR allows the network to provide this critical information proactively which in turn allows developers and users to make intelligent resource decisions at application run time. By identifying performance issues on the spot, we believe perfSONAR can have a major impact on the future proliferation of grids and other high impact applications which depend upon network reliability and performance."

OGF22 Update – Advanced Registration Closing/Preliminary Program and More

OGF22, taking place February 25-29 in Cambridge, Massachusetts will feature a wide variety of new and exciting content. Please see the below important information and deadlines for this event:

Registration
Advanced registration will close January 11. Register now to save!

Preliminary Program
Keynote Sessions
Charlie Catlett will make the Tuesday (2/26) morning keynote address. Charlie founded the Global Grid Forum (predecessor to OGF) and is currently the CIO of Argonne National Laboratory, Director of Argonne's Computing and Information Systems Division, and a Senior Fellow at the Argonne/University of Chicago Computation Institute. As CIO of Argonne, Charlie will bring unique perspectives on operating a world-class data center while providing leading-edge technology solutions to computational scientists. Other keynote speakers are in being lined up. Watch for further announcements as their availability is confirmed.

OGF Chartered Groups And BoFs

Monday, February 25 through Friday, February 29

Activities are scheduled throughout the entire week where many of OGF's chartered groups will meet to make further progress on their work. In addition, the latest hot topic in distributed computing - "Cloud Systems" and how it may or may not impact Grids will be tackled in a special BoF session.

OGF Specification Adoption Track and Demonstration Session

Monday, February 25 and Tuesday, February 26

Many OGF working groups have recently demonstrated interoperability of their specifications. These significant accomplishments show how much of the work of OGF is shifting from developing specifications to seeing specifications being used and adopted into products. This multi-session track will feature contributions from developers and users showcasing their work implementing OGF specifications. In addition, live demonstrations will take place during the welcome reception on Monday evening.

Workshops

OGF22 will feature a number of interesting workshops that will address a variety of important topics in Grid:

- Enterprise Adoption Workshop – Tuesday, February 26
  This full day workshop - "Enabling the Next Generation IT Infrastructure" will focus on putting Grid in the context of the distributed computing infrastructure, including virtualization, service-oriented architecture, web services and next-generation networks. Emphasis will be on solving business problems and meeting business needs, with a focus on applications on-boarding and Managing Infrastructure. Presentations and case studies are expected from representatives from the 451 Group, VMware, Intel, SAS, Platform Computing, DataSynapse and others will show a variety of areas that can be addressed by Grid, in practical and achievable ways.
- Financial Services Workshop – Wednesday, February 27
  The financial services industry is one of the early adopters, and successful users, of Grid technology. As the technology matures, new (and old) issues emerge like expanding grids to support transactional workloads, data management and movement, and putting policies in place for charging and tracking usage. Presentations and case studies that describe the emerging issues the financial services industry is now facing in Grid will be featured in this half day workshop.
- Pharma/Life Sciences Workshop – Wednesday, February 27
  The pharmaceutical industry and life sciences in general have begun to depend on in silico experiments as critical elements in the path to discovery. This has enabled exploration that is only limited by their ability to access sufficient computing resources and to store and manage the resultant mountain of data. Presentations and case studies from those that have turned to Grid-based solutions to meet these needs will be featured in this half day workshop.
- Data Management Workshop – Thursday, February 28
  This full day workshop, co-hosted by OGF and the Storage Networking Industry Association (SNIA), will address data management challenges with a focus on file system and data movement. We will also explore potential synergies and common areas between OGF and SNIA.

Other Community Content

Monday, February 25 through Thursday, February 28

Sessions from the greater Grid community covering numerous topics and a variety of subjects will also be presented. These include workshops, tutorials, BoFs and presentations on marketing, education in emerging/developing countries, Grid Identity Providers, GridWay, and others topics.

Hotel Information
OGF has secured a block of sleeping rooms at the Hyatt Regency Cambridge (meeting location) at a special rate of $129 for single or double occupancy. Reservations must be made on or before February 8, 2008, to receive the discounted rate. Reservations can be made by calling the hotel directly (be sure to reference 'OGF22') or online

Student Scholarships
A limited number of student scholarships are available for qualifying students. Hosted by OGF and CoreGRID, the OGF22 student scholarship program provides a unique opportunity for students to participate in OGF program sessions and group sessions as well as network with distinguished community participants and leaders. Application details and eligibility requirements

Group Chair Scheduling
OGF Group Chair session scheduling will remain open until January 11. Session scheduling and document details

Sponsorship Opportunities
Sponsorship opportunities for OGF22 include Display Tables, Welcome Reception, Chair Appreciation Night, OGF Co-branded Promotional Item, Wireless, Lunches and Breaks, and Badge Lanyard. By sponsoring OGF22 your organization will gain greater visibility and recognition as a Grid thought leader in our global community. This will be our only OGF event in the USA for all of 2008, so please strongly consider sponsoring if you are interested in reaching this audience. The sponsorship prospectus provides additional information.

Documents Update

Recently Published Recommendations
The following documents were published in December. Congratulations to all the authors and working group members involved in getting this important work accomplished!

DOCUMENT	TITLE	TYPE	AUTHORS	AREA
GFD.122	Grid Network Services Use Cases from the e-Science Community	Informational	T. Ferrari	Infrastructure
GFD.121	OGSA® Data Architecture	Informational	D. Berry, A. Luniewski, M. Antonioletti	Data
GFD.120	Open Grid Services Architecture® Glossary of Terms Version 1.6	Informational	J. Treadwell	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.
OGSA® Resource Selection Services: Specification
Distributed Resource Management Application API 1.0 - IDL Specification

Upcoming Events

- **OGF22 Cambridge, Massachusetts February 25-29, 2008**
- **OGF23 Barcelona, Spain June 2-6, 2008**
- **OGF24 Co-located with GridAsia08 Singapore September 15-19, 2008**

Other Events

- **INGRID 2008 Lacco Ameno, Island of Ischia, Italy April 9-11, 2008**
  3rd International Workshop on Distributed Cooperative Laboratories: Instrumenting the Grid
  Call For Papers – deadline to make submissions is January 15th, 2008

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!

Please Renew your OGF Membership For 2008

Please remember to renew your 2008 OGF membership. We made great progress in 2007 and we need your continued support to build on that momentum. We also ask all organizations that are in the Grid community who are not currently members to join us in 2008. Memberships start for as little as $10,000 ($5,000 for eligible non-profits). You may contact Steve Crumb or visit the membership portion of our website for more information. Thanks!