OpenGridForum
OPEN FORUM | OPEN STANDARDS

GridConnections

December 2007

News and Information for the Open Grid Forum Community

In this Issue:

SC07 Report - High Performance Computing Basic Profile

Steven Newhouse, Co-Chair, HPC Basic Profile Working Group

SuperComputing 2007 in Reno saw the second 'interoperability fest' involving the OGF's HPC Basic Profile (HPCBP) specification. Groups from academia, research and industry came together to demonstrate the interoperability of the core specification and to prototype extensions to the core specification used to support more advanced capabilities – such as file staging.

High Performance Computing Basic Profile (HPCBP) Background

The HPCBP specification has been developed by the OGF's HPC Profile Working Group over the last 18 months to define how to submit, monitor and manage jobs using standard mechanisms across different job schedulers or Grid middleware from different software providers. A significant milestone was passed in August 2007 when the HPC Profile Working Group's first specification – the HPC Basic Profile 1.0 – was published as a proposed recommendation. As part of the process of moving the HPCBP from a proposed to a full recommendation it is necessary to gather experiences of implementing the specification, and to demonstrate its interoperability between two or more independent implementations. The working group, comprising individuals from industry, academia and research, decided to use SuperComputing 2007 as the venue for an 'interoperability fest' of the published specifications. The HPCBP is focused on just managing job submission yet it provides sufficient core functionality to enable it use as a basis for integration into applications (both desktop and web based), its use as a basis for meta-scheduling and from within workflow engines. It leverages standards developed within the OGF (the Basic Execution Service and the Job Submission Description Language specification) and those from the broader web services community (WS-Security, WSDL and SOAP).

At SC06 last year over a dozen groups drawn from industry, academia and research demonstrated interoperability between their prototype implementations of this web service on their existing job submission infrastructures. This year, with the HPCBP specification now a proposed OGF standards recommendation, many of these groups returned to show how the longer-term integration of the HPCBP into their plans were proceeding. Several commercial organizations were using prototypes developed by their engineering teams as part of their plans for inclusion in their products. Teams from academia and research were showing how their implementations could be deployed and used with large scale Gird infrastructure deployments.

SuperComputing 2007 Demonstration Results

Participants in this year's activity included Altair Engineering, Inc, Microsoft, Platform Computing, the London e-Science Centre at Imperial College London on behalf of OMII-UK, the University of Virginia e-Science group, and representatives from the EGEE and NIC/FZJ groups within the OMII-Europe project.

All of the participants were demonstrating their implementations of the HPCBP, revised from last year's versions due to the final changes made in the specification during the public comment period, and verifying their continued interoperability with other participating implementations. This interoperability work had been greatly accelerated through the use of a web based compatibility tester developed by the University of Virginia e-Science group. The portal allows users to run a series of tests, derived from the HPCBP specification, to verify that their endpoint is compliant with the standard by seeing how it accepts and generates the XML messages, and responds to failures or incorrect messages.

From this ground work, the interoperability testing that took place in the run-up to SC07 and at the meeting was broadly very successful. Jobs were passing from clients to services located on different operating systems (e.g. different versions of Linux and Windows), using different web service stacks (including Windows Communication Foundation, gSOAP, Axis, XFire) and interfacing to different jobs scheduling interfaces (including PBS, LSF, Windows Compute Cluster Server v1 & v2, SGE, Torque, ARC, CREAM and Globus). Minor issues continue to be found relating to the interoperability of different web service hosting environments, (e.g. in how WS-Security is handled) rather than any fundamental problems with the HPCBP specification.

Commercial Adoption Plans

Several commercial organizations also announced their plans to provide implementations of the HPCBP compatible web services in their products.

Platform Computing have contributed their implementation to an open source project (BES++ hosted on SourceForge) that uses the gSOAP toolkit to submit jobs through the HPCBP into LSF. Work is ongoing at the University of Virginia to extend this software to submit jobs into other schedulers, such as PBS. This software will be integrated as part of Platform's product line in 2008.

Microsoft were demonstrating a prototype of an HPCBP compatible web service running on Windows HPC Server 2008 – the next version of Microsoft's HPC product due out in the second half of 2008. It uses the Windows Communication Foundation as its web service stack and it is currently planned for inclusion in the second beta of HPC Server 2008 due out in Spring 2008. Altair Engineering also had a prototype integration of an HPCBP compatible client in their PBS product that demonstrated the use of this specification as a meta-scheduler. Job submitted into PBS using the conventional command line utilities could be transferred to another HPCBP compliant resource. At SC07 jobs were being submitted into a PBS instance running on Linux on the Altair stand and being executed through the HPCBP web service on a Windows machine on the Microsoft stand.

Next Steps

After SC07 work is continuing within the working group to capture our interoperability experiences in order to support the migration of the HPCBP specification from a proposed to a full recommendation. These experiences are also being used to refine those implementations that will be emerging in 2008 as products. The practical experience gained at SC07 with the Data Staging and Activity Credential extensions will be used to develop a single extension focused exclusively on integrating File Staging operations that need different credentials with the current JSDL specification already used in the HPCBP. It is hoped that this specification will be submitted into the OGF process in early 2008.

SC07 Report - OGF Grid Interoperation Now Demonstrations

Q&A with Morris Riedel, GIN CG Secretary

What is GIN (and is there 'WHISKY' too)?

Many Grid projects have begun to offer production services to end-users during the past several years with an increasing number of application projects that require access to a wide variety of resources and services in multiple Grids. Therefore, the purpose of the Grid Interoperation Now (GIN) Community Group is to organize, manage and demonstrate a set of interoperation efforts among production Grid projects and e-Science infrastructures using computational or storage-related resources in more than one production Grid. The GIN group defines interoperation as what needs to be done to get production Grids to work together as a fast short-term achievement using as much existing technologies and workarounds as available today. Hence, this is not the perfect solution and different than interoperability that is defined as the native ability of Grids and Grid technologies to interact directly via common open standards in the near future. In other words, GIN demonstrates where standards are still missing in production Grids and provides lessons learned out of production Grids if standards have to be improved towards production usage. And sorry to say, but to the best of my knowledge there is no WHISKY group, but there was once a BOF for it.

Is GIN only focusing on job submission in cross-Grid scenarios?

In one word, No! Our interoperation efforts of production Grids typically involve a lot of different integrated components out of different technology areas like information, security or data. That's why GIN implements interoperation in specific areas. First, authorization and identity management (GIN-AUTH) deals with resource sharing among members of the GIN Virtual Organization (VO). Second, the data management and movement (GIN-DATA) area is working on the interoperation of different data management technologies currently in use of multiple e-Science infrastructures. These include the Storage Resource Broker (SRB), Storage Resource Managers (SRM) and GridFTP. Third, the job description and submission (GIN-JOBS) area focuses on job management across different Grid technologies and middlewares used in production. Another important area is the information services and schema (GIN-INFO) area, because the efforts conducted in this area basically provide the base for cross-Grid interoperation taking up-to-date information into account. These interoperations rely on information models such as Common Information Model (CIM) and Grid Laboratory Uniform Environment (GLUE) or information systems such as Berkeley Database Information Index (BDII) and Monitoring and Discovery Services (MDS). Finally, the operations experience of pilot test applications (GIN-OPS) for cross-Grid operations works on different applications that require resources from multiple Grid infrastructures.

Tell us about what the GIN team just accomplished at SC07

The idea of having interoperation demonstrations at the SC conferences is not something new and GIN presented a couple of interoperation demonstrations already at SC06 as one of its milestones. In a similar manner as last year, the GIN group again organized a set of demonstrations that have been announced through several channels. Finally, we ended up having about 9 different interoperation demonstrations that have been shown at booths at SC07 that either did daily demonstrations (e.g. Enabling Grids for e-Science, Hungarian Grid, NorduGrid Collaboration booths) or a dedicated Grid day (e.g. John von Neumann Institute for Computing booth) that include some of the demonstrations.

In more detail, members of the group presented 'Storage Resource Broker (SRB) and Storage Resource Manager (SRM) Island interoperability', including the transfer of files between an SRM and an SRB using OGF's GridFTP specification as the file transfer protocol. Another demonstration related to data was the 'WS-DAIR interface for the gLite AMGA Metadata Catalogue'. The gLite-AMGA metadata catalogue, developed within the EGEE project, provides

access to relational data on the Grid and is widely used in the scientific Grid community. The demonstration presented interoperation by implementation of a WS-DAIR compatible interface in AMGA. Another interesting demonstration was the ‘information system interoperation’. An attempt was made to translate information from all the existing production Grid infrastructures in order to populate a single information resource. The resource chosen was a Berkeley Database Information Index (BDII) and the common format used was the Glue Schema version 1.2. The result is a BDII that contains information from 9 production Grid infrastructures (EGEE, OSG, NDGF, Naregi, Teragrid, Pragma, DEISA, NGS, APAC). This information was used to show the location of the computing centres in Google Earth and the current Grid landscape for the production Grids.

There were also demonstrations that underlined how the standard landscape of OGF is evolving in production Grids. In the ‘UNICORE and gLite Interoperation’ demonstration, the interoperation of the gLite (European EGEE Grid) and UNICORE 5 (European DEISA Grid) production Grid middleware using proprietary protocols motivates the usage of standards from OGF in future. In this demonstration, users build a job on the gLite User Interface and submit it through the gLite Workload Manager to an interoperation Computing Element that forwards jobs to the UNICORE Grid for their execution on supercomputers. We see here a clear demand of a an open standard like the OGSA-Basic Execution Services (BES) that could be used in cross-Grid job submission scenarios if both middleware would use it in future. Therefore, the ‘VOMS-enabled and OGSA-BES-based job submits between UNICORE and gLite’ demonstration showed the evolution of the above mentioned demonstration for future usage in the production Grids using emerging open standards from OGF and OASIS. The UNICORE 6 BES and CREAM-BES of gLite are both OGF OGSA-BES-compliant services developed in OMII-Europe. However, lessons learned from GIN indicate that job submission standards alone are not enough for interoperation and security is also a major important point of it. Therefore, we used the new SAML-compliant Virtual Organization Membership Service (VOMS) of OMII-Europe as security setup. All in all, the demonstration showed one client that is able to submit jobs to both of these Grid middleware systems without changing the security setup or the job description language. This was not possible before and is a real success story coming out of OMII-Europe. Hence, there are components from OMII-Europe that provide technical interoperability between EGEE (through gLite) and DEISA (through UNICORE 6) in the near future.

End-users at SC2007 have seen the ‘GIN portals’ demonstration in action based on P-Grade. P-GRADE Portal (Parallel Grid Run-time and Application Development Environment) is a service rich graphical environment for the development, execution and monitoring of data-driven Grid applications. The P-GRADE Portal is used by users and application developers of several national (UK NGS, HunGrid, Turkish Grid, etc.) and international Grid based virtual organizations (EGEE, SEE-GRID, etc.). The demonstration included most important features and typical use cases of the environment as well as other outcomes of the developer alliance which is behind the P-GRADE Portal efforts.

Another demonstration showed the emerging interoperability between the accounting systems DGAS and SGAS achieved by the OMII-Europe project. DGAS is the accounting system of gLite used within EGEE and is currently augmented with an OGF OGSA - Resource Usage Service (RUS) compliant interface. SGAS is the accounting system of the Swedish Grid and Globus Toolkits (Tech Preview) and is also currently augmented with an emerging OGF OGSA-Resource Usage Service (RUS) interface. Here we demonstrated the exchange of accounting information by using an SGAS client that is used to extract usage records compliant with the OGF Usage Record Format (URF) standard from a DGAS server. Aligned with this demonstration we have shown an OGSA-RUS interface in UNICORE that uses the URF information for resource-level monitoring in an interoperable application called LLView that could leverage URFs via OGSA-RUS interfaces also from other systems such as DGAS or SGAS. We in GIN think that accounting and billing will get more important in the next years in

the production Grid infrastructures, especially in cross-Grid scenarios. That's why we have been interested in demonstrations in this rather new area of GIN.

An interesting scientific interoperation scenario was shown in the 'DEISA and Australian Grid Interoperation' demonstration. This demonstration showed the processing of a small workflow employing the NAMD molecular dynamics suite. Several DEISA sites participated and specific jobs have been formulated using OGF's JSDL specification. Within the workflow a number of compute intensive jobs have been submitted to DEISA sites with the help of the DESHL UNICORE command line interface that uses the OGF SAGA standard. At the same time a similar number of jobs have been submitted via GRAM to the AU Grid and executed primarily at Monash. Output from the Australian jobs will be moved to the DEISA GPFS filesystem via OGF's GridFTP and post-processed later. In another demonstration, a subset of the GridFTP2 protocol has been shown as implementation in dCache, Globus C, and Java client libraries. They are already used within production Grids and a demonstration was shown by the NDGF.

So, it seems that several major milestones for the group were accomplished at SC07

Yes, the interoperation demonstrations at the SC conferences are a highlight for the group itself showing short-term interoperation achievements or emerging standards that are mature enough to be considered for production usage within production Grids in the near-future (e.g. OGSA-BES, URF). However, the GIN community group always tries to transfer experience from production Grids into the standardization process in OGF and this is a good opportunity for OGF groups to gain lessons learned or even to collect specific standard requirements from real life interoperations in production Grids. For instance, at OGF 21 in Seattle, we organized a 'Software providers meet GIN and Standards' session that also brought together the people for discussions having different viewpoints. According to the feedback we got from participants, this session was very interesting and thus the interoperation demonstrations at SC07 are important milestones but not the only ones we have in mind.

What are the top 3 things the team learned from the SC07 demonstrations?

[1] That open standards have to evolve faster and must take into account the experience from real production Grids and also the feedback from developers that implement the standards in the software provider stacks that are finally used on these infrastructures. Hence, the implementation of the standard should not start after standardization but during standardization to prove that the standard can be actually used later within production Grids and their production Grid middleware.

[2] It makes no sense if numerous projects develop interoperability components via the integration of common open standards and they are not used in the production Grids afterwards. Hence, projects that adopt and implement the open standards from OGF in components of Grid middleware must try to get their implementations back into the major software provider stacks that are deployed on these large production Grid infrastructures.

[3] Focus on spin-off activities that deal with particular requirements. For instance, investigations during the SC06 and SC07 demonstrations showed that an achievement of interoperability in terms of information of production Grids requires an agreement on the information content and schema. As a result, the Glue Schema activity is now an OGF working group and is defining this common schema for Grid computing. There are also other specific requirements that need more focus or a spin-off activity like the worker node profile or dedicated security setups to be used with specific standards, for instance a common way of passing credentials through an OGSA-BES-based job submission for third party data staging.

What are your plans as a result of the interop?

There were many lessons learned and we will write them down as experience documents, which will also include lessons learned and approaches from the SC06 demonstrations. Also, at OGF22 we will host a SC07 debriefing session offering discussions with other OGF participants about how OGF standards can be used in production Grids and which standards production Grid software providers should integrate as soon as possible to get more sustainability in interoperation efforts of production Grids. Real interoperability scenarios requiring resources in multiple Grids are actually appearing more and more and thus the interoperability in production Grids must be more widespread than our several interoperation demonstrations.

Anything else you would like to add?

I would just like to personally thank all members of the OGF GIN group that have worked on the numerous demonstrators in the past several months. Also, thanks to all those national and international funding bodies that supported the numerous projects world-wide so that their members are able to participate within GIN. Finally, also thanks to OGF marketing in supporting us with flyers and announcements at SC07.

Now Open - OGF22 Registration/Call for Participation/Sponsorships

Registration

Registration is now open for OGF22 Cambridge, Massachusetts February 25-29, 2008. With the holidays quickly approaching, register now to save…advance registration will close January 11.

Call For Participation

The Call for Participation (CFP) for OGF22 is now also open. We are seeking content related to the following topics:

OGF Chartered Groups And BoFs

OGF Specification Adoption

Two weeks ago at SC07, OGF demonstrated interoperability of 12 OGF specifications and announced the implementation of the HPC Basic Profile in products being planned by Microsoft, Altair Engineering and Platform Computing. These significant accomplishments show how the work of OGF is being recognized as key to further adoption of grid technology. This multi-session track will build on the SC07 momentum by featuring contributions from developers and users showcasing their work implementing OGF specifications.

Pharma/Life Sciences

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 are being solicited.

Financial Services

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. We seek presentations and case studies that describe the emerging issues the financial services industry is now facing in Grid.

Other Community Content
We also welcome submissions related to HPC, Virtualization, Grid application and usage in eScience environments, Grids in the IT data center, Security, Networking, Visualization in a grid environment, and Grid application development

A successful event depends on great content which will come primarily from our community so please to make your submission today. Hurry, as the CFP will close December 21!

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.

Globecom07 Report
Craig Lee, OGF President

OGF hosted the “Grids and Service-Oriented Networks Panel” at Globecom 2007 on November 27. This panel explored where grids could support the emerging concept of service-oriented networks, i.e., networks that are used and managed as a set of services. The panelists were John Wittgreffe (ICT Research Chief, BT Group Chief Technology Office) Inder Monga (Director, CTO Labs, Nortel), and Gigi Karmous-Edwards (Principal Scientist, MCNC). A common theme from this session was that web/grid standards, Web 2.0, virtualization, and the convergence of digital media and appliances all point towards the notion of a "soft telco" -- that is to say, Service-Oriented Networks. BT's Web21C, for example, is a web services and software development kit with the goal of hiding complexity and making it easy for end-users to "mash-up" their systems. The real trick, though, is how to flexibly implement this in a wide-area, distributed environment. The contextualization of virtual machines may play a key role here in transparently managing complexity. There is clearly a need for a service-oriented network reference model wherein things like hierarchical SLAs can be composed to achieve desired end-user behaviors. Dynamic provisioning, advance co-reservations, and managing compute data affinity is also important. Follow-up work is being discussed for reference models and VM management.

Documents Update
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.
Defining the Grid: A Roadmap for OGSA® Standards v1.1
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**
- **OGF23 Barcelona, Spain June 2-6, 2008**
- **OGF24 Co-located with GridAsia08 Singapore September 15-19, 2008**

Other Events

- **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.

- **IGT Annual Event & Exhibition December 3, 2007 Hertzelia Art Center, Israel**

The IGT 2007 annual conference deals with the challenges of the Grid infrastructure virtualization in the next generation data centers.

- **International Winter School on Grid Computing 2008 Call for Participants**

The International Winter School on Grid Computing (IWSGC) '08 will examine the conceptual and practical underpinnings of today's grids. Experts will provide exciting practical exercises, discuss the challenges of building and sustaining e-Infrastructure, report its rapid influence on the way we research, design and make decisions and share their vision of the developments and challenges ahead. IWSGC will be delivered fully online: participants will be able to attend the School without traveling. Selection for the Winter School is competitive based on the information supplied on the application form and by an applicant's referee. Applications are now being accepted and must arrive by the 14th December 2007 closing date.

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 Organization Membership 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!