

## OGF Monthly Newsletter

December 2006

*A look at Past Accomplishments and A Preview of Coming Attractions*

### In this Issue:

[Headline News – OGF HPC Profile Interoperability Demonstration](#)

[Take Our Member Survey](#)

[New Document Published](#)

[New Area Director Named](#)

[New Organizational Members](#)

[Upcoming Events](#)

[Individual Members - Activate your 2007 Membership](#)

### Headline News – OGF HPC Profile Interoperability Demonstration

**Q&A with Marty Humphrey, HPC Profile WG Co-Chair**

*Tell us a little about what the HPC Profile team just accomplished at SC06*

The HPC Profile WG is an effort in OGF to create the profile and protocol specifications needed to realize the vertical use case of batch job scheduling of scientific/technical applications. Our WG's approach includes leverage two existing OGF WGs – the OGSA Basic Execution Service (BES) WG and the Job Submission Description Language (JSDL) WG. Basically, our HPC Profile WG is contributing requirements and approaches to the BES WG and JSDL WG to ensure that their respective specifications could be used in our particular use-case of batch-job scheduling.

By this past September (OGF 18 in Washington, DC), our WG had gotten to the point that we felt pretty good about the state of the specifications (the BES spec, the JSDL spec, and our document describing how to basically combine the two specifically for batch-job scheduling), so we decided that it would be great to have a big “interoperability-fest” at SC2006 – essentially having people aim their HPC Profile-compliant clients at other projects' HPC Profile-compliant services. We thought that this would be great fun and very satisfying! Most importantly, we believed that you get to a point in writing specifications where you think they're correct, but it takes a number of different, independent implementations attempting to talk to each other to identify and resolve some of the really tricky issues. We believed that we were at that point, so we really needed to attempt this “interop-fest” to improve the specifications. SC 2006 was a great forcing function for this!

We ended up having 12 groups participating: Altair Engineering, Argonne National Lab (Globus Alliance), CROWN, EGEE, Fujitsu Labs of Europe, HP, Microsoft, Platform Computing, Tokyo Institute of Technology, UK eScience (OMII-UK), University of Virginia, and Genesis II (UVA). Most groups showed independent demos, some focusing on server-side functionality, and some focusing on client-side functionality. Rich Ciapala of Microsoft came up with a great demo, where he submitted a job to one of the other participant's servers. Once on that machine, the “job” was actually an HPC Profile-compliant client, which “forwarded” the job to another HPC Profile-compliant resource, which forwarded the job to another HPC Profile resource, for a total of maybe 6 hops. I thought this idea was outstanding – it showed how the HPC Profile could be used to facilitate “super scheduling”, where a resource might get a job submitted to it, but then the resource off-loads it to someone else, either because it is too busy or it might not have the requested application currently installed, or whatever. In this case, because all resources “spoke” the HPC Profile, the client could communicate with any of the back-end resources that the job

ultimately executed on. The really neat thing about this demo was that there was probably an equal mix of Linux-based systems and Microsoft Compute Cluster Server (CCS) systems. That's what it's all about – interoperability and support for heterogeneity!

I know that a number of the groups showed this demo from their booths, and I heard a lot of good comments about it. A number of people were surprised that we could do this, and saw how this capability could pay off for them in the future.

*That seems like a major milestone for OGF*

Oh, I certainly think so! We created the HPC Profile WG with some pretty tight deadlines – that is, many people in the group are from companies (as opposed to academia/labs), so they wanted to get this working and somewhat stable as fast as we could. These people have products to ship, and their customers want interoperability. So they were very driven. This interoperability demo at SC2006 really showed how people could come together in OGF, with tight deadlines, and produce a set of specifications and a relatively large number of interoperable, independent implementations. This really said something about our OGF community!

*What are the top 3 things the team learned?*

If I were forced to choose three from my perspective as HPC Profile WG co-chair, I think I would say:

[1] An “evolutionary approach” is really good – focus on existing tooling or tooling that's arriving-very-soon. The ability to have common protocols or interfaces for batch scheduling is needed NOW, and continually waiting or anticipating next year's tooling or protocols is not appropriate for this particular effort. Rather, create a design that works today and can be updated without great upheaval in the future. Anticipate next year in today's design, but don't overly rely upon it coming, because it might never come. By doing this, we were able to create as many interoperable implementations as we did in time for SC 2006!

[2] Don't make the effort too broad – I think we were successful because we scoped-down the problem to a much more manageable level, specifically the execution of scientific/technical apps on batch job schedulers. If this were broader, I believe we would still be attempting to create the SINGLE protocol that accommodated EVERYTHING, and that just hasn't worked out in the past. There are too many moving parts, and the “fringe” parts tend to make the “core” parts too difficult to implement.

[3] build broad community involvement from the very beginning – the HPC Profile WG has many people involved in it, and from the start a lot of people have contributed to its formation. Getting people involved and making them stakeholders with something to gain if the WG is a success (and something to lose if the WG fails) is crucial.

*How was the demo received by the HPC audience?*

Collectively we spoke to a large number of people at SC 2006. In talking afterwards with some of the participants in the demo, I think people who saw the demo generally fell into one of three categories:

[1] “Interesting! I could see how this would help!” – The capabilities that the WG collectively demo'd were very compelling to those people who have not seen something like it.

[2] "It's about time!" – Some people we spoke to thought that at least on the surface, the ability to have a single interface/protocol to the multiple batch job scheduling systems absolutely made sense and did not seem too difficult, so they said they were surprised that it took THIS long for the community to do this. I think we generally replied that it's only now that Web services protocols and tooling are converging such that this type of thing is possible. We can't really speak to the efforts of others in the past, but we CAN try to step up and make a difference today. That's what we're doing.

[3] "A good start – but I need more!" We explained that there's some misconception here in the HPC audience. The HPC Profile WG (in particular, the HPC Profile WG use-cases document) identifies a "basic profile" and extensions. The "HPC Basic Profile" is the core functionality that we expect will be implemented by ALL batch scheduling systems. The "extensions" by definition are common functionalities that we expect will be implemented by two or more (even perhaps a large majority), but not by everyone. A perfect example of an "Extension" is file staging, for example when you want to move a file to the batch scheduling system before execution, so that the executable can use the file as input. We believe that this is a really important requirement addressed by most batch scheduling systems. But is a single standardized way to do this for ALL batch scheduler systems needed? No, we don't believe so. In fact, the HPC Profile interoperability demo at SC2006 showed this! We were able to do a lot of interesting things without a standardized way to stage files in and out. Input/Output file staging was still required for most of the demos, but there wasn't a single way to do it (some used FTP, some used HTTP, some used their own particular way). So we don't think this should be in the HPC Basic Profile. It will absolutely be one of the first "extensions" – and services will have the ability to assert that they implement the "data staging extension" as well as the Basic Profile. So when people said that they "needed more", we absolutely agreed! But this demo was just the HPC Basic Profile, not any/all of the extensions as well (which are clearly part of the plan). When we explained this to people, I think generally they understand and respected/appreciated this approach.

*What are your plans to do as a result of the interop?*

Well, we've already started the teleconferences to address important issues that were raised during the implementation and interop-fest itself. These issues include: clarifying the semantics of certain operations, clarifying some XML schema to make it more amenable to certain tooling, and specifying security in the Base Profile. We're hoping to get the Base Profile into public comment either this month (which may be optimistic) or January. And we're also going to start focusing on one or more extensions. One of the first that we'll start with is the File Staging extension that I mentioned above. We will also work on a compliance suite for the HPC Basic Profile, which I think is really important – how do you know if a candidate service/client complies with the spec? This can be a tricky question. The compliance suite will help here and we're looking to leverage the excellent work of WS-I here. Overall, I think we made GREAT progress via the interop-fest at SC2006, and we'll be looking to continue our aggressive schedule in the beginning of 2007! And of course we're always looking for more people to get involved!

*Anything else you would like to add?*

I'd just like to personally thank a few people. It's really tough to single out anyone, because it was really a team effort -- I was continually impressed by the number of people on the teleconference calls and contributing to the email discussions. But certainly Chris Smith of Platform, Rich Ciapala of Microsoft, and Glenn Wasson of the University of Virginia deserve a special mention – their energy and technical skills really helped us identify and battle through some tough issues! It was really great to work closely with these guys. Andrew Grimshaw was important as well – he wasn't directly involved in the HPC Profile WG, but he was very interested in the success of BES, so he made sure that the BES WG was open and responsive to the comments coming from our HPC Profile WG. And certainly the WG as whole really appreciates the vision, technical expertise, and management of Marvin Theimer, who was with Microsoft until recently. It was Marvin who really

pushed this effort and made it a success. Marvin is no longer engaged in this effort because of his new role at Amazon, and I really believe that we have a great momentum to successfully complete this effort, but without a doubt we would not be where we are today without Marvin.

### Take Our Member Survey

All OGF members and active participants are urged to spend 5 minutes to take our on-line member survey. Results will be used to improve our member services and help guide our strategy. We will publish the results for viewing by the entire community in January. The survey is completely anonymous. <http://www.zoomerang.com/survey.zgi?p=WEB225WPEGUMR4>

### New Document Published

#### GFD.083 – Firewall Issues Overview

This document identifies typical firewall scenarios of today's grid environments. It structures the scenarios into use cases and classifies these cases into general communication concepts that can be used by grid application developers and management personnel as guidance. The classifications will be used to propose new or recommend existing academic and/or standards based solutions to the grid community. Congratulations to the Firewall Issues Research Group! <http://www.ogf.org/documents/GFD.83.pdf>

### New Area Director Named

Our Area Directors manage and facilitate groups, milestones and deliverables in defined areas of expertise such as data, architecture, management, etc.

#### Erwin Laure, Data

Erwin is Technical Director of EGEE, where he coordinates the technical work of the EU funded EGEE and EGEE-II projects. He is currently co-chair of GIN-WG. He has research interests in wide-area distributed computing and parallel computing, particularly in the data management area and production Grids. Congratulations Erwin! He can be reached at [Erwin.Laure@cern.ch](mailto:Erwin.Laure@cern.ch)

### New Organizational Members

Welcome to our newest organizational members **Sun Microsystems (Gold) and Availigent (Silver)**. **Sun** is changing the nature of computing with Sun Grid Compute Utility, enabling users to purchase compute power over the network when and where needed. Sun Grid provides optimal flexibility in usage and provides zero barriers to entry and exit. <http://www.sun.com/>.

**Availigent** is a leading Application Service Management provider for dynamic data centers. Its Duration software enables organizations to deliver optimal application service levels, maximize the utilization of commodity-based computing infrastructure and minimize the costs and complexities of application deployment and system administration. Availigent is headquartered in San Jose, CA. [www.availigent.com](http://www.availigent.com).

### Upcoming Events

#### DMTF Management Developers Conference December 4-7 in Santa Clara, CA

OGF is a sponsor of this event and will be represented by Tom Roney, Ellen Stokes and Fred Maciel who will lead a session on behalf of the OGSA Resource Management Design Team. Visit the conference website for more information <http://www.mandevcon.com/>.

#### OGF19 Chapel Hill, North Carolina January 29 - February 2, 2007

Program will include a full slate of chartered group sessions, 2 days of eScience workshops, Enterprise sessions, and more. Register by December 8 and save up to \$100

[http://www.ogf.org/OGF19/events\\_regstrn\\_ogf19.php](http://www.ogf.org/OGF19/events_regstrn_ogf19.php).

#### OGF20 Manchester, U.K. May 7-11, 2007

Save these dates! OGF20, co-located with EGEE's 2nd User Forum and hosted by UK e-Science and the University of Manchester, will be the premier grid technologies event of 2007.

### **Individual Members - Activate your 2007 Membership**

Please remember to activate your 2007 membership. The fee is \$195 and registered members will receive a \$100 discount to all 2007 OGF events. Individuals who are not employed by OGF organizational members may join OGF at <http://www.regonline.com/112287>. Individuals that registered under the 2006 program are requested to reregister for 2007 at a reduced 2007 annual fee of \$100. Your 2006 member number will be required to obtain this discount.

The success of OGF depends upon member participation. All of the significant events, activities and accomplishments of the forum are member driven. Please contact any OGF staff member if you want to get involved [http://www.ogf.org/ggf\\_contact.htm](http://www.ogf.org/ggf_contact.htm). We welcome your input!