Report from GGF14: Finally, Open Grid Service Architecture (OGSA) is Real!

Alan J Weissberger
aweissberger@sbcglobal.net

1. OGSA WG Comes of Age

a] Executive Summary

Several very important developments at the GGF14 meeting in Chicago, June 27-30th have the potential to greatly accelerate grid middleware development. This could lead to interoperable distributed computing sooner then expected and on a much larger scale then has been seen to date.

First and foremost among the GGF14 happenings was the emergence of the **OGSA WG** as a viable standards setting body, rather then a "castles in the sky" committee developing abstract architectural concepts that could not be implemented. After years of discussion, the OGSA WG has now identified a core set of web services standards and described how they can be used together to form the basic building block for grids. This is specified in a new document titled, "**Basic Web Services Resource Framework (WSRF) profile."** Discussed below, this document can serve as the infrastructure for grid protocols to build upon. In addition to this profile, the OGSA WG submitted two other documents at GGF14 that will soon go out for public review:

-A **Profile Definition**, which is a guide to producing other OGSA profiles
-A **Standards Roadmap for OGSA**

The OGSA WG published a **v1.0 Architecture** document and a **Glossary** this past March. Those are both being revised with a target completion date of early 2006. Finally, a new OGSA **SCRUM WG** has been chartered to collaborate with other standards bodies and is described later in this section.

b] OGSA WSRF (Web Services Resource Framework) Basic Profile and Roadmap for OGSA

For several years, GGF and OGSA WG have stated that a grid service is a web service, with overlaid grid resources and the notion of state (and state changes) associated with those resources. Yet there was little or no progress on defining a minimal set of web services standards or specifications that would comprise this web services ecosystem/ infrastructure for grids to build on top of. This will now change with the introduction of the **first web services profile for grids**.

The **OGSA WSRF (Web Services Resource Framework) Basic Profile 1.0** (AKA "Basic WSRF Profile") describes uses of widely accepted specifications that have been found to enable interoperability. The specifications considered in this profile are, in general, those associated with the addressing, modeling, and management of state. The "Basic WSRF Profile" combines the WS-I Basic Profile 1.1, WS-I Basic Security Profile 1.0 with WS-Addressing (from W3C) and several standards from the OASIS WS Resource Framework and WS Notification TCs. In particular, WS Resource Lifetime, Resource Properties, and Base Faults are referenced from OASIS WS-RF TC, while WS Base Notification (scheduled to go to OASIS vote in early July) is referenced from the OASIS WS-Notification TC. HTTP over Transport Layer Security (TLS), and the TLS Protocol (better known as

SSL) from the IETF are also referenced. The "Basic WSRF Profile" then explicitly specifies requirements for exposing state, lifetime management, notifications and security. It specifies that if you are to implement a particular function, you shall do it in a very specific way (this is the essence of a profile, in this authors opinion).

OGSA had previously published a high level architecture document, a glossary and use cases. But the whole notion of OGSA as the "implementable reference architecture" for grids was very abstract, fuzzy, and vague. The **Roadmap for OGSA** clarifies the meaning and role of OGSA in grids and provides a framework for managing future GGF standards work. It identifies documents, schedules, dependencies on other documents, and other relevant standards organizations that might collaborate with OGSA WG. OGSA is described as both an architectural process and a set of normative (required) documents/standards, which will enable deployment of interoperable grids.

The OGSA Roadmap document was just submitted to the GGF Editor on June 24th. The plan is to publish v1.0 of this Roadmap after a 30-day public comment period, which has just commenced.

Editors Opinion: While the three new OGSA WG documents submitted are a great start, there are several work areas that OGSA urgently needs to initiate. These include: an activity to define conformance test suites for the Basic WSRM profile and other OGSA approved standards/ specs, plans for interop testing of OGSA compliant implementations, composability aspects of OGSA standards/ specs. Implementers need to know how the numerous GGF WG/RG specs work together, if at all? For example, how does Basic Execution Services work with Resource Management Services or with Job Submission Description Language? Which OGSA grid specs/standards will use the Basic WSRF Profile and how do they interface to it?

Note that there are many GGF WG and Research Groups (RGs). Please refer to http://www.ggf.org/ggf_areasgrps_overview.htm

c] What is a SCRUM? Let's Hope it's not a SCAM!

On Wednesday afternoon, June 29th, the first meeting of the just formed OGSA **SCRUM WG** was held. Officially titled, the **Standards development organizations Collaboration on networked Resources Management** (how do you get SCRUM acronym from this?), it is a "cross institutional" working group within GGF. The first deliverable will be a "**standards landscape document**" for distributed computing. That document is intended to provide information regarding the definitions, taxonomy and interplay of the various specifications of each respective organization. In addition to GGF leaders, several chairs from the IETF, DMTF (2), OASIS (2) participated in this inaugural session. The goal is increased efficiencies in the generation of standards, which will hopefully lead to more effective results that are delivered faster.

2. WSDM "Deep Dive" and Interoperability Demo's

On Wednesday morning June 29th, there was an excellent GGF session on the technical details of the **OASIS WSDM (Web Services Distributed** Management) specs- Management Using Web Services (**MUWS**) and Management of Web Services (**MOWS**). This was followed by an afternoon

session which provided an overview of managing grid resources using WSDM and two very impressive interoperability demos (HP and IBM) which illustrated management of a grid service as a web service. In one demo, an unauthorized program was detected by a Manager and prevented from being downloaded by a Blackberry. In the other demo, a simulated space –data collection station was taken off line for recalibration, with clients re-routed to another station using End Point References (WS-Addressing). Both sessions described how WS-RF and WS-Notification standards are used by WSDM to manage resources and send resource property change notification messages.

It's likely that the two WSDM standards –MUWS and MOWS- will be used together with WSRF standards, WS Base Notification and WS Topics to manage grid resources. This fulfills the mission of the former GGF Common Management Model (**CMM**) WG. In fact, the **Resource Management in OGSA** document (GFD 1.045) from the now inactive CMM WG references WSDM's MUWS and the DMTF Common Information Model (**CIM**). No further profiling work is required to use WSDM to manage grid resources. It was noted that WSDM makes use of both WSRF and WS Notification standards. The GGF CDDLM WG references WSDM for configuration and deployment of grid services. So we have several levels of dependencies:

CDDLM→WSDM→WSRF and WS Notification.

Separately, the June 2nd publication by IBM, HP, and Computer Associates of their jointly authored **Management Roadmap** was reviewed. This roadmap characterizes the challenges facing the IT community in light of diversity of systems, managers, firewalls, and router issues. It proposes the use of Web Services management (WSDM) as the foundation of grid management. WSDM is seen as the mechanism for managing grid resources and expressing management metrics. Using WSDM should simplify the integration of grid management solutions for existing systems and applications.

Several "infrastructure web services specifications," needed at run time, are identified in this Roadmap. These include: WS-Meta Data Exchange, WS-Coordination, WS-Atomic Transactions, WS- Resource Properties+, WS-Resource Lifetime+, WS-Service Groups+, and WS-Base Notification (from the OASIS WS-N TC).

+ The OASIS WS-RF TC developed these standards.

3. Globus Toolkit 4.0 (GT4) Being Effectively Used in Real Grid Applications

Ian Foster of Argonne National Laboratory & University of Chicago presented an in depth tutorial on the components of GT4 and improvements over GT3. This was followed by several testimonials from companies that were using GT4 in real grid applications. The GGF14 session abstract can be downloaded from:

http://www.globus.org/alliance/events/ggf14/

Background: On April 30th, the Globus Consortium released the Globus Toolkit Version 4.0 (GT4) as a stable, enterprise ready set of services and software libraries incorporating the latest web services standards, with new security and authorization features. The Globus Toolkit is "an open source software toolkit used for building enterprise-level Grid systems and applications."

Ian described GT4 as a toolkit to realize "service oriented applications." This involves wrapping the applications into services and composing applications into workflows. GT4 was described as a service enabling toolkit to build applications that exploit the grid infrastructure. The core web services standards supported in GT4 provide the infrastructure for building new grid services, including those for security, execution management, data management and monitoring. This "2nd Generation" web services implementation has been optimized for performance, flexibility, stability and scalability, according to Foster.

The main GT4 enhancement from GT version 3.0 is **combining grid protocols with emerging Web services standards**. In particular, GT4 includes support for: Web Services Interoperability Organization (WS-I) Basic Profile, OASIS Web Services Resource Framework (WS-RF) and Web Services Notification (WS-N) standards; Security Markup Language (SAML); Extensible Access Control Markup Language (XACML).

Foster stated that Grid FTP services in GT4 use a **Grid Resource Allocation Management Service** (**GRAM**), which is a type of execution management. GRAM provides an interface for process execution management in the form of a common web services interface to job schedulers.

This editor believes that GT4 will be a key component for accelerating the adoption of enterprise-class grids. This opinion was corroborated by several presentations from early adopters of GT4, who shared their positive experiences with the session attendees. Four speakers offered testimonials for GT4:

-Terry Harmer of the Belfast eScience Center talked about GT4 use in Gridcast- a commercial application for broadcast media- that is being delivered to the BBC. He stated that their grid focus was as an integration fabric which tied distributed sites together. GT4 demonstrated significant improvements in performance over the previously used GT3. It also was more reliable.

- Ravi Subramanian of Intel described how GT4 is used as a building block software base for Intel's enterprise grid (for integrated circuit design), "deployed across all machines." Ravi said that GT4 offered a quick migration from GT3, was easier to work with and provided performance improvements. On the negative side, he commented that GT4's migration away from User Hosting Environment (UHE) was not consistent with Intel's grid model of "per user container environment." [GT4 moved away from UHE because it was found to limit scalability and slowed perfomance.]

-Jaime Frey of the University of Wisconsin described experiences with GT4 in the Condor-G project – a high throughput computing system operating throughout the University of Wisconsin Madison campus. Condor-G used GT4, with Java client bindings, for job management and scheduling in grids. While still in an ongoing testing and debugging phase, Jaime liked GT4's scalability (one process for all jobs), WS Resource Lifetime features (job killed automatically when lifetime expires), and automated server recovery (no client action needed).

- Hidemoto Nakada of AIST talked about Ninf-G4- a Grid RPC (Remote Procedure Call) system based on GT4. With GT4, RPCs are realized on the grid through standardized interfaces.