Editors:
Dr. David Snelling, Fujitsu
9 September 2006

# Technical Strategy for the Open Grid Forum 2006-2010

## Status of This Document

This document provides information to the Grid community about the overall technical strategy of the Open Grid Forum (OGF). It does not define any standards or technical recommendations.

This is an evolving document and as a result even formally published versions are subject to regular revision. The contents reflect the rough consensus view of the Technical Strategy Committee, but are intended to restrict or direct activities within the OGF. Rather this strategy document points in a direction that appears to the Technical Steering Committee (TSC) to be valuable and important at this time. It is being provided to gain input and feedback with respect to its structure and content. To facilitate this feedback, this document is available to all of the OGF membership.

## Copyright Notice



## Abstract

This document describes the overall OGF technical strategy spanning a three-year timeframe from 2006 to 2010. The technical strategy ultimately describes the output of the OGF standards working groups as well as the requirements (in the form of use cases) that serve as the inputs to standards working groups. The technical strategy is represented in the form of a roadmap of standards working group output over time with specific short-term milestones and target deliverables. Working group output and requirements inputs will be correlated.

## Contents

## 1. Introduction

OGF's mission (www.ogf.org) is to build an international community to accelerate Grid adoption by providing an open forum for grid innovation and developing open standards for Grid software interoperability; and, this mission is based on that belief that standards-based grid computing is critical to enabling business value and scientific discovery.

In the spirit of openness as well as enabling business value and scientific discovery, the TSC was established to ensure that there is an alignment between the OGF's technical strategy and the needs of the business and scientific community. Formulation of a technical strategy is the primary output of the TSC. The TSC meets on a regular basis and consists of members of the OGF community who represent the Grid community of users, architects, developers, vendors, etc.

### 1.1 Purpose of this Document

This document, the Technical Strategy Document (TSD), is intended to capture and communicate OGF's technical strategy.
"A strategy is a long term plan of action designed to achieve a particular goal, as differentiated from tactics or immediate actions with resources at hand" [1]. The OGF TSD is intended to:

1. Provide a concise view of the OGF technical direction and priorities.
2. Provide a mechanism to align key stakeholder requirements with OGF technical directions and priorities
3. Provide an indication of where more effort is needed, and what actions are needed to promote specific standards within the industry.

### 1.2 Document Structure

The structure of this document is as follows. In Section 1.3, we present a concise statement of the three-year goal of the Open Grid Forum. In Section □, we outline a number of high value use cases that serve to focus attention of which technologies are needed in the near term, Section 5, outlines the tactical priorities in the form of a roadmap for several identified specifications.

### 1.3 Background

The long-term vision of Grid can be summed up as follows: "Distributed computing across multiple administrative domains." The notion of *distributed computing* as used in this definition includes a wealth of highly complex technologies, some still in the focus or research. This definition complicates matters further by including operation across multiple domains of administrative control. The security, privacy, economic, and political aspects of Grids increase ~~by orders on magnitude~~ significantly with the introduction of multiple administrative domains.
However, there are many valid uses of the term "Grid" that do not fully encompass the above definition, therefore we believe that the term should be refined into three categories of Grid

1. Collaboration Grids: The above definition applies here. These Grids involve multiple organizations and individuals, security domains, protocols, discovery mechanisms, etc. This is the long-term vision of Grids.
2. Data Center Grids: These Grids are in most ways as complex technically as Collaboration Grids and involve the complete dynamic life cycle of service deployment, provision, management, and decommissioning as part of their normal operation.

However, the multiple domains are either absent or highly integrated, at least at a political level. These are the, frequently heterogeneous, production Grids of major data centers. One of the envisioned use cases for Data Center Grids within the Enterprise is Utility Computing, where there are many individual political/security domains and application types deployed on an infrastructure that is managed with grid protocols. In this case, there are multiple domains sitting on top of an integrated domain, which leads to a different sort of world entirely.

3. Cluster Grids: Aimed at high performance/throughput computing, these Grids are mostly workload scheduling environments. They tend to be statically, rather than dynamically, deployed. The services are either generic in nature, e.g. a job submission service, or provide the same service all the time. They do not typically support the whole service provisioning life cycle.

All three uses of Grid are valid, but the distinction is very important to avoid confusion in the industry. The Open Grid Forum embraces all three categories of Grid computing. In this document we place our use cases into one or more of these categories.

## 2. Goal of the Open Grid Forum: 2006 – 2010

Concisely put, the Goal of the Open Grid Forum for the 2006 to 2010 time frame is given below.

> **The Open Grid Forum should commit all its available resources to the goal that before this decade is out, commercial and academic organizations will build real operational grids using OGF defined components.**

No other single technical goal can more completely focus the activities of our newly united organization or more clearly define its success and no other goal will be more challenging or difficult to achieve. Furthermore, achieving this goal will require us to draw energy from all stakeholders within the organization.

One important aspect of ~~this~~ goal is that it is defined in terms of specific use case patterns and the specifications needed to enable these use case patterns. In some cases, the development of a particular specification may still be in a very early and immature state - more of a collection of community-initiated best known methods (BKMs). Thus, it is anticipated that each high-level use case pattern may contain both concrete specifications and some gaps that must be filled by BKMs until further technical and/or political maturity occurs. To build a strategy around the goal we must:

- Identify and focus on the main, common use-cases, patterns, and scenarios that commercial and academic grids require.
- Provide best practice and other documents and meetings that allow communities to evaluate and adopt Grids today and provide a pathway for the standards process.
- Identify and complete the core architectural standards required to build robust, commercially viable, grid solutions.
- Mobilize the whole of OGF, all the work, community and research groups to meet this challenge.

- Encourage software developers, in the open source community and the commercial software providers to adopt and implement these standards in products. They must do so early and often, as this is part of the standards process.

- Hold regular alignment summits where the OGF functions and key stakeholders review the TSC strategy and update based on lessons from OGF activities and broadly the complete Grid community. Here we would review lessons from the Grid interoperability work in OGF (eScience) , OGF best practice documents as well the Enterprise Voice of Community and other forums.

## 3. Technical Strategy Alignment Process

**Error! Reference source not found.** depicts, at a very high-level view, the OGF Technical Strategy Alignment process.

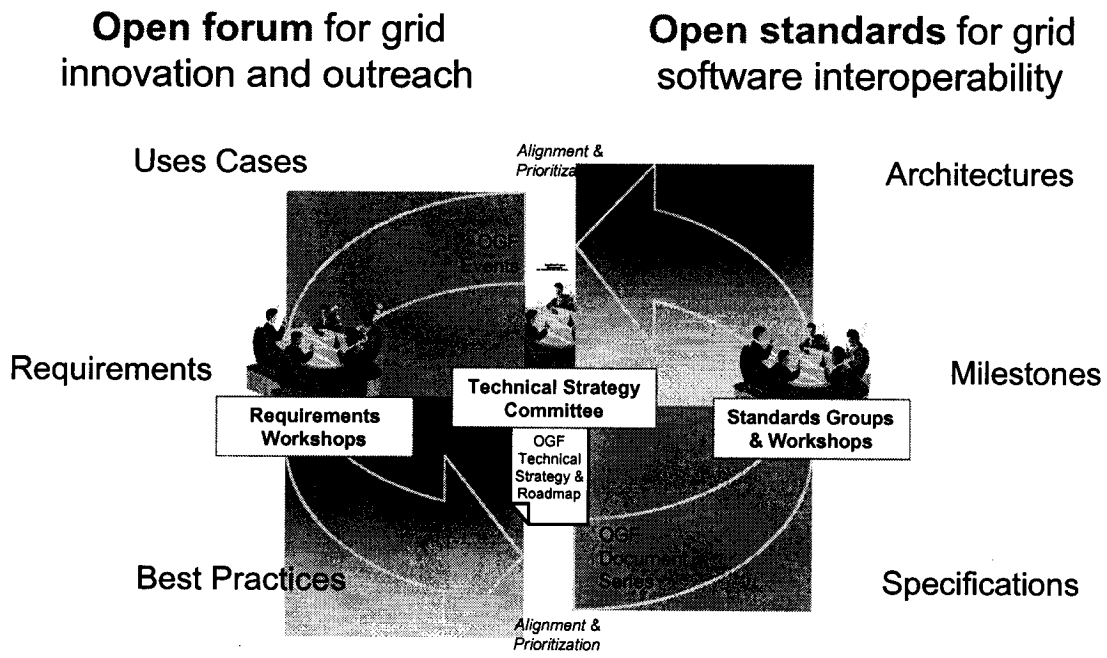


**Figure 1 : Technical Strategy Alignment Process**

The right half of the alignment process is concerned with the standards working groups and their production of specifications and reference architecture. The left half of the alignment process is concerned with the inputs to the standards working groups. This input represents requirements from the Grid community at-large and may be represented in the form of use cases or best practices. The requirements are gathered from various requirements gathering groups such as the Enterprise Grid Requirements Research Group (EGR-RG) or the Storage Networking Community Group (SN-CG), or the Telco Community Group. Each of these groups meets to capture requirements that are particular that group. Each group's requirements are then rolled-up into a merged and prioritized list which is then brought forward to the TSC.

The TSC represents the key part of the alignment process where requirements are matched-up against the current technical direction of the standards working groups, and a determination is made as to the degree of alignment/misalignment that exists between the requirements and the technical direction. The three parts described, the requirements gathering, the standards work and the technical alignment all operate simultaneously and in parallel.

When merged and prioritized requirements are brought forward to the TSC, assessment criteria will be applied to each requirement in order to determine the appropriate response to that requirement.

The assessment criteria for incoming requirements are as follows:

1) What is the degree of alignment with OGF objectives, goals and current technical direction?

a) Is the requirement just an extension to what OGF is already doing?

b) Is requirement a minor tweak to the current technical direction?

c) Is there a work group that already exists that would be a natural fit to handle the requirement?

d) What is the interest level in working on the requirement?

2) How important and beneficial is this requirement to the community?

a) How universal is the requirement to the community at-large

b) Is this requirement a top-line priority to a particular segment of the community?

3) Do the resources exist to address the requirement?

a) Do the resources exist to actually work-on the requirement?

b) Do the skills/knowledge/expertise exist to help address the requirement?

4) Does this requirement exclude other industries or vendors

5) What is the magnitude (resources and time) of the effort needed to address the requirement?

6) How complex or risky is the requirement?

7) What is the timeframe in which the requirement needs to be addressed in order or it to be useful?

Once the assessment criteria for the incoming requirements have been applied the next step is to determine the appropriate action/response to each requirement. The range of actions/responses is as follows:

1) Send the requirement to an existing working group for whom the requirement would be a natural fit.

2) Start a new standards working group to work-on the requirement.

3) File the requirement as pending due to current lack of interest or resources.

4) Ignore as out of scope for OGF.

Do you mean "if"?

5) Since requirement is already being addressed by an existing standards working group, make the connection between the source of the requirement and the working group. Make certain to include in the TSD. If specification is ready, send an open letter to the vendor/developer community to suggest implementation.

6) Form a new standards working group to create a specification for existing technology

7) Form a new Enterprise Group to develop a Best Practices Document that might offer an interim solutions (during standards development) or may turn-out to be a prescribed permanent solution

8) Develop a New OGF processes to handle addressing the requirement.

9) Refer to an existing technology specification from outside OGF.

10) Declare the requirement out of scope if it is not consistent with OGF mission or objectives and may not even be a "Grid" issue.

## 4. High Priority Use Cases

Based on the current state of play within the standards development activities of OGF, the input available from the former EGA, the results of a recent community survey, and discussions with key stakeholders within the community, the following use cases and scenarios have been identified as a priority to meet the stated goal in the section "Goal of the Open Grid Forum".

### 4.1 Grid APIs

Re to what?

One of the current critical barriers to widescale adoption of the Grid paradigm is the continued evolution of the underlying programming interfaces. Developers of both end user applications and middleware services need programming interfaces that provide stability across both different middleware technologies and changes in the underlying protocols through either different approaches or versions.

Having started from a diverse set of use cases collected from the 'grassroots' OGF community, the Simple APIs for Grid Application (SAGA) working group has evolved an interface specification that is agnostic to the underlying middleware. The API includes functional support for job submission and management, resource discovery, and data management, access and replication. This is on top of generic support for asynchronous notification, error reporting and security. As the semantics of the generic API stabilizes and moves forward to standarisation, work continues on generating language specific bindings and the solicitation of new use cases to drive a second round of API development.

Implementation and support of these APIs across different middleware technologies, such as Java RMI, raw TCP/IP, or various flavors of Web Services is now essential in order to build on top of the initial prototype implementation. The development of these APIs by middleware providers is seen as critical in the short term to ensure the uptake of Grid technology in the wide business and technical communities.

### 4.2 Job Submit

Does this mean GFD 22 DRAMA?

The simplest job submit use case is a high-throughput compute cluster that is managed by a batch job scheduler and that is used only from within an organization. Aspects to consider include user interface (semantics only, not GUI issues), the state model, and resource descriptions. With respect to the user interface, users expect to be able to submit jobs, query the

or BES

status of running jobs, cancel a job, and list job belonging to them on the the job service. The state model need to capture, at a minimum, the concepts of running and finished, as well as a state before execution commences (pending or queued).

Users expect to be able to discover something about a job service before they attempt to use it. However, given the complexity of the resource modeling domain, only a small set of standardized resources can be specified, such as number of CPUs/compute nodes needed, memory requirements, disk requirements, etc.

There are a number of common use cases that extend this very simple one that should be addressed in some way. In particular being able to describe the service's fault tolerance model, to handle extended functionality offered by specialist schedulers, to provide notification of job status to the user, and to advertise and request other aspects of quality of service.

I think the TSC might want to think beyond FTP

### 4.3 File Movement

The TSC has identified a need to define an interface that standardizes the process of invoking the movement of large amounts of data. No standardization activity addresses this issue so far. This use case covers the problems of discovering data transport protocols available at the data's source and destination location and agreeing on one of them, and the actual invocation of the agreed data movement, including direct data movement and 3rd party data movement. Executing a data movement includes the invocation of the transport protocol itself, and applying the previously agreed parameters where appropriate. While the data movement is executing, control and management operations on the data movement are necessary, such as "cancel", "suspend", and "resume". Progress information, including general status information, must be provided to interesting parties as well.

Stemming from this from this use case are requirements for moving bytes over an internet connection, e.g. FTP and GridFTP.

I am doubtful this will be useful any time soon.

### 4.4 Application Provisioning

Job submission, and indeed any sort of workload manager, implies the ability to discover, describe, provision and manage the lifetime of an appropriate application code onto an identified computing resource. In many instances, this can be done at a very high level, but some scenarios will require very specific descriptions at the application layer. This, in turn, may place requirements for a specific operating system and version, possibly implying a certain patch level and hardware requirements. EGA's Reference Model describes the overall flow of activity involved in provisioning a high-level component and decomposing the required work into accessible quanta: ACS and CDDLM are specific proposals/WGs that attack the problems of describing and managing the lifetime of specific applications.

I think a simple but practical based on a few simple cases - WSRF, RPM etc. are the best

### 4.5 Data Provisioning and Data Grids

Data intensive grids are of increasing importance and require components to handle files, different types of databases, caching, transport, metadata, federation leading to managed data, information and knowledge. One needs to address provisioning and management at both the data and storage levels.

To do data provisioning, the GME (Grid Management Entity from the EGA Reference Model) must become the intermediary between the compute, switching and storage infrastructures that make up the set of grid resources. The dynamic nature of grid-based applications requires provisioning on several levels at once to achieve what may look to the end user like an atomic operation.

I think file movement & "data provisioning" are almost the same - and that storage is something different

Data provisioning typically requires at least three steps: initial population, keeping the data in sync, and cleaning up the data when it is no longer needed. If the container must be populated with an initial data set from somewhere, additional work is required. There may be an opportunity to use cloning technology to greatly enhance the efficiency of the copy operation.

After the initial provisioning step, the data may needed to be frequently snapshotted and/or replicated for Disaster Recovery (DR) or other purposes. At the end of the job, results may need to be copied elsewhere to a location of the client's specification. In addition, all temporary copies of any data may need to be securely shredded when the user or application is done with the container and its offspring. Yet the user's desire is simple: a data container conforming to some service level that the system has previously advertised.

This, then, is the grid data provisioning problem in a nutshell. This high-level view decomposes rapidly into a number of other problems—each a significant subject in its own right—as one begins to drill down. In addition, at bottom one needs APIs that actually perform the provisioning and monitoring operations in order to build a GME that can offer the convenient and dynamic abstraction of a grid, which holds so much promise.

Once that decomposition has been done, we are in a position to examine whether suitable interfaces into the actual grid resources exist, and if so, where. This paper will not attempt this, except superficially. However, in one area (Asset Discovery), we will take a deeper dive as an example of our intended approach moving forward.

### 4.6 Grid Security

There are two grid security scenarios seen as critical in the short term, standardized secure connections and standardized attribute based authorization.

#### 4.6.1 Secure Channel Establishment

The growing number of Web services specifications makes it important to understand and define the interaction and use of these specifications to ensure interoperability. The WS-I Basic Security Profile 1.0 **[Error! Reference source not found.]** defines a collection of normative profiles that provide guidance on issues of interoperability for secure communication of basic Web services based on such specifications. A profile is needed highlighting the specific specifications, and parts thereof that provides secure (authenticated) connections between entities on the Grid.

#### 4.6.2 Attribute Based Authorization

Web and Grid Services define operations for their invocation and means for publishing their internal state. Security is expected to be managed using orthogonal mechanisms.

There are a number of authorization systems currently available for use on the Grid as well as in other areas of computing, such as Akenti [3], CAS [4], PERMIS [5], VOMS [6]. On the abstract level these types of authorization services have similar semantics - they are given a description of the initiator (which might include the initiator's privileges), a description of an action being requested (including its argument), details about the target resource to be accessed, and any contextual information such as time of day, and they provide an authorization decision whether the action should be processed or rejected.

I think we should stay out of mechanisms and focus on message content & terms and error codes

This use requires the definition of an authorization service that allows services to make queries and receive responses in regards to access control on grid services. For example authorization is needed for accessing functionality over an exposed Grid Service portType. A client sends a request for an authorization decision to the authorization service and in return receives an

authorization assertion or a decision. A client may be the service itself, an agent of the service, or an initiator or a proxy for an initiator who passes the assertion on to the service.

## 5. Tactical Priorities and Roadmap

Table 1: Simplified Specification Roadmap includes a list of required specifications needed to meet the use cases outlined in Section 3. The table is organized as follows:

- Use Case: The section number(s) of the use case(s) or scenario(s) for which this specification is required.

- Specification Name: The short name of the specification where possible. If no specification exists yet, this entry is left blank. Note that there may be several specifications addressing a given use case.

- Current Status: The current status of this specification on the following scale from Concept through Deployment. The levels on the scale are roughly sequential, but not all steps are always taken.

    - Concept: Concept exists and (proprietary) proof of concept implementations exist.
    - WrkGrp: Working Group formed to create the specification.
    - Draft: Draft specification exists.
    - Interop: Reference Implementations and Working Group lead interoperability tests.
    - Spec: Specification completed to OGF Proposed Recommendation.
    - Product: Available as a supported product, including Open Source based service contracts.
    - Deploy: Deployment observed in a production setting, commercial or technical.

- Milestone 1: The month, year and target status for the first milestone with respect to the specification. These milestones need not point to the next stage in the status list.

- Milestone 2: The month, year and target status for the second milestone with respect to the specification.

- Area: OGF technical area of responsibility for the specification.

**Table 1: Simplified Specification Roadmap**

| Use Case | Specification | Status | Milestone 1 | Milestone 2 | Area |
|---|---|---|---|---|---|
| Grid APIs | SAGA | WrkGrp | Dec 06: Spec | Dec 07: Product | Applications |
| | DRMAA | Interop | | | Applications |
| Job Submit | JSDL 1.0 | Spec 56 | Jun 07: Product | Dec 07: Deploy | Compute |
| | OGSA-BES | Draft | Oct 06: Spec | Dec 07: Product | Compute |
| | HPC Profile | WrkGrp | Dec 06: Spec | Dec 07: Product | Compute |
| File Movement | DMI | WrkGrp | Dec 06: Spec | Dec 07: Product | Data |
| | ByteIO | Interop | Oct 06: Spec | Dec 07: Product | Data |
| | GridFTP | Product | Dec 06: Deploy | | Data |
| Data Provisioning and Data Grids | WS-DAI | Spec ## | Aug 07: Product | Mar 08: Deploy | Data |
| | WS-DAIR | Spec ## | Aug 07: Product | Mar 08: Deploy | Data |
| | WS-DAIX | Spec ## | Aug 07: Product | Mar 08: Deploy | Data |
| Application Provisioning | CDDLM | Spec 69 | Aug 07: Product | Mar 08: Deploy | Management |
| | ACS | Spec ## | Aug 07: Product | Mar 08: Deploy | Management |
| Secure Channel Establishment | OGSA-SBP-Core | Spec ## | Aug 07: Product | Mar 08: Deploy | Architecture |
| | OGSA-SBP-SecChan | Spec ## | Aug 07: Product | Mar 08: Deploy | Architecture |
| Attribute Based Authorization | OGSA-AuthZ-SAML | Spec 66 | Aug 07: Product | Mar 08: Deploy | Security |

Table 1 is not a complete list of OGF activity nor is it a statement of the overall importance of these specifications with respect to the rest of the work in OGF. These specifications merely address the priority use cases set out in section 3; other work in OGF continues independently.

The contents and schedule represented in Table 1 will change over time, based in input from stakeholders as to perceived priorities, chairs in terms of available resources to meet milestones, and general input from the community.

## 6. Security Considerations

All OGF document must have this section. With respect to this document, it would be a serious omission if security specifications were not part of the OGF short term roadmap and an identified priority. Noting that this is the case meets the requirement that this document address security in a way consistent with the nature of the document.

## 7. Contributors

The Editors listed on the title page are those committed to taking permanent stewardship for this document – receiving communication in the future and otherwise being responsive to its content. Their contact information is provided below:

Dr. David Snelling
Fujitsu Laboratories of Europe
Hayes Park, UB4 8FE
David.Snelling@UK.Fujitsu.com
+44-208-606-4649

Other contributors include all members of the Technical Strategy Committee, namely Steven Newhouse, Andre Merzky, Geoffery Fox, Robert Fogel, Hanoch Eiron, Ian Foster, Chris Kantarjiev, Tom Maguire, and Joel Replogle. The TSC would also like to thank Mark Linesch and the chairs of a number of working groups who's documents provided much needed perspectives and content.

## 8. Intellectual Property Statement

The OGF takes no position regarding the validity or scope of any intellectual property or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; neither does it represent that it has made any effort to identify any such rights. Copies of claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the OGF Secretariat.

The OGF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights, which may cover technology that may be required to practice this recommendation. Please address the information to the OGF Executive Director.

## 9. Disclaimer

This document and the information contained herein is provided on an "As Is" basis and the OGF disclaims all warranties, express or implied, including but not limited to any warranty that the use of the information herein will not infringe any rights or any implied warranties of merchantability or fitness for a particular purpose.

## 10. Full Copyright Notice



## 11. References

1. Definition "Strategy" from Wikipedia, http://en.wikipedia.org/wiki/Strategy, August 2006.
2. A. Barbir, M. Gudgin, M. McIntosh, and K.S. Morrison (ed.): Basic Security Profile Version 1.0, Web Services Interoperability Organization, Working Group Draft, 29 August 2005. http://www.ws-i.org/Profiles/BasicSecurityProfile-1.0-2005-08-29.html.
3. Thompson, M., et al., "Certificate-based Access Control for Widely Distributed Resources," in Proc. 8th Usenix Security Symposium. 1999.
4. Pearlman, L., V. Welch, I. Foster, C. Kesselman, S. Tuecke, "A Community Authorization Service for Group Collaboration," Proceedings of the IEEE 3rd International Workshop on Policies for Distributed Systems and Networks, 2002.
5. Chadwick, D.W., O.Otenko, " The PERMIS X.509 Role Based Privilege Management Infrastructure", Proceedings of 7th ACM Symposium on Access Control Models and Technologies (SACMAT 2002).

6. "VOMS Architecture v1.1," http://grid-auth.infn.it/docs/VOMS-v1_1.pdf, February 2003.