Grid Interoperability: The Interoperations Cookbook

Laurence Field
CERN

Laurence.Field@cern.ch

Markus Schulz
CERN

E-mail: Markus.Schulz@cern.ch

Abstract. Over recent years a number of grid projects have emerged which have built grid infrastructures that are now the computing backbones for various user communities. A significant number of these communities are limited to one grid infrastructure due to the different middleware and procedures used in each grid. Grid interoperation is trying to bridge these differences and enable virtual organizations to access resources independent of the grid project affiliation. This paper gives an overview of grid interoperation and describes the current methods used to bridge the differences between grids. Actual use cases encountered during the last three years are discussed and the most important interfaces required for interoperability are highlighted. A summary of the standardisation efforts in these areas is given and we argue for moving more aggressively towards standards.

1. Introduction

Interoperability is defined as "*The ability to exchange information and to use what has been exchanged*" [1]. Interoperation is defined as "*The use of interoperable systems*" [2]. It is important to be aware of the differences between these to concepts. Interoperability is only the first step towards interoperation.

To understand the problem of *grid interoperation*, we first need an explanation of a grid. In their influential paper,[3] Ian Foster et al. defined a grid as being "*coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations*". Today the word grid means different things to different people. Virtualized services, clusters, campus grids and data centres have all been given as examples of grids. As the fundamental problem addressed by each example is significantly different, confusion can arise when discussing grids as a general concept. One of the most fundamental aspects differentiating types of grids is whether they are intra-organisational or cross-organizational. Due to the different view points it is very important to define the context of the grid which is being discussed to avoid any confusion. Throughout this paper grids are discussed in the context of multi-institutional infrastructures for e-science.

Traditionally, users of resources and the resources themselves are located within the same organization. An organization represents an administrative domain where the organization has complete control over everything within its domain. Within this domain an organization will provide its users with access to its resources according to its own policies. Users from different organizations collaborate to achieve common goals and wish to increase efficiency by pooling the resources available to them, splitting tasks by specialty and sharing common frameworks.

In the original definition of a grid there were three fundamental entities; resources, institutions (organizations) and virtual organizations. The key concept introduced was the concept of the virtual organization. A virtual organization is a group of users from multiple institutions who collaborate to achieve a specific goal. Institutions support virtual organizations and hence allow users, who may belong to different organizations, to access the resources.

As each institution is autonomous, resource access many require the use of specific mechanisms. In order for members of a virtual organization to access resources located at an institute, grid middleware is used to provide an interface at the boundary of this administrative domain. Grid middleware follows the "hour glass" model.[4] At one end there is a diverse set of systems and at the other end there are many virtual organizations that have their own applications. The applications can gain access to the heterogeneous systems though a small set of well defined interfaces.

Over recent years a number of grid projects, many of which had a strong regional presence, have emerged to help coordinate institutions build *“multi-institutional infrastructures for E-science*”. Today, we face a situation where a number of infrastructures which used different middleware and procedures. The original solution proposed by Globus was the set of common interfaces provided by the Globus Toolkit[3]. However, many of the grid projects found that some of these interfaces did not meet the production requirements demanded by the virtual organizations and started developing alternative interfaces. As the infrastructures evolved independently from each other, different interfaces were developed. Although standard interfaces would have prevented this divergence, as highlighted by the experience with Globus, it is difficult to define a good standard without prior experience in the domain for which the standard addresses. Standardization can be a slow process and the standards require time to mature. The infrastructures were already being built to meet the timescales defined by the virtual organizations and could not wait for standards to appear. One of the advantages of this situation is that the experience with the difference approaches used is valuable input to the standardization process.

The computing interface is an example of a grid interface. The aim of this is to provide the virtual organization with a generic interface with which they can access the various batch systems deployed at the different institutions. Each grid infrastructure has defined their own computing interface and as a result there are now more computing interfaces than there are batch system implementations!

Members of VOs belong to the organizations which also provide resources. The organizations may participate in different grid infrastructures. Grid interoperation is trying to bridge these differences and enable virtual organizations to access resources independently of the grid infrastructure affiliation. Without grid interoperation the virtual organization would be limited to only one grid infrastructure. As different grids have their own middleware and polices, they can also be seen as an administrative domain. In a sense, the challenge of grid interoperation can be seen as a similar problem to that of users and institutions, but now with virtual organizations and grid infrastructures.

2. Interoperability

In order to overcome these differences, it is first necessary to understand each infrastructure. The fundamental aspect of the infrastructure is the grid middleware which provides the interfaces at the organizational boundary. An interoperability matrix (fig 1 below) shows the critical interfaces which are required for common tasks and highlights the different implementations used in each infrastructure. The matrix typically covers four main areas; security, information services, job management and data management. Once the differences have been understood, the process to overcome these differences is known as "*achieving interoperability*".

	ARC	**OSG**	**EGEE**
Job Submission	*GridFTP*	*GRAM*	*GRAM*
Service Discovery	LDAP	LDAP	LDAP
Schema	*ARC*	*GLUE*	*GLUE*
File Transfer	GridFTP	GridFTP	GridFTP
Storage Interface	SRM	SRM	SRM
Security	GSI	GSI	GSI

Figure 1: Interoperability Matrix

2.1. User Driven

The virtual organisations can themselves strive to achieve interoperability as shown in figure 2. They can access multiple grid infrastructures and either split the workload between the infrastructures or build into their frameworks the ability to work with each infrastructure. One of the problems with this approach is that is places significant effort on the virtual organisation. In addition, as each virtual organisation solves the problem, this results in a significant duplication of effort and loss of productivity. The effort required also increases with the number of grid infrastructures which the virtual organisation would like to use. In addition this results in a keyhole approach where the minimum common subset of functionality is used and handling failures can be problematic. This approach was used by the Atlas community to overcome the problem of interoperation with OSG, EGEE and Nordugrid.[5]

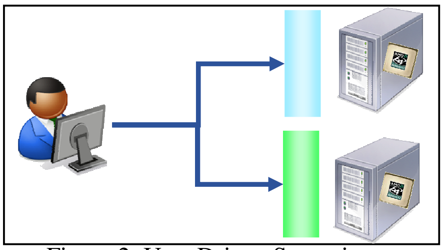

Figure 2: User Driven Scenario

2.2. Parallel Deployment

Institutions can achieve interoperability by deploying multiple interfaces as show in figure 3. The resource can be made available to multiple infrastructures by deploying the respective grid services that are required. This approach would enable seamless interoperation from the virtual organisations perspective; however, it is a significant overhead for the institute. The system administrator will need to become an expert in each grid service and each service requires resources that could have been used by the virtual organisation. The effort required also scales with the number of grid infrastructures that the institute site wishes to support and therefore this method is only recommended for large resource centres. Forschungszentrum Karlsruhe used this approach to overcome the problem of interoperation with EGEE, Nordugrid and D-Grid.[6]

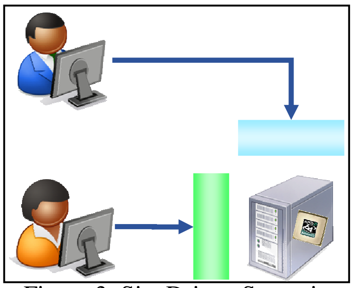

Figure 3: Site Driven Scenario

2.3. Gateways

A gateway, as shown in figure 4, is a bridge between grid infrastructures. It is a specific service which makes the grid infrastructure look like a single resource. This results in a keyhole approach where the minimum common subset of functionality is used and handling failures can be problematic. Gateways can also be a single point of failure and a scalability bottleneck, however, this approach is very useful as a proof of concept and to demonstrate the demand for achieving interoperability. This approach was used by Naregi in their interoperation activity with EGEE.[7]

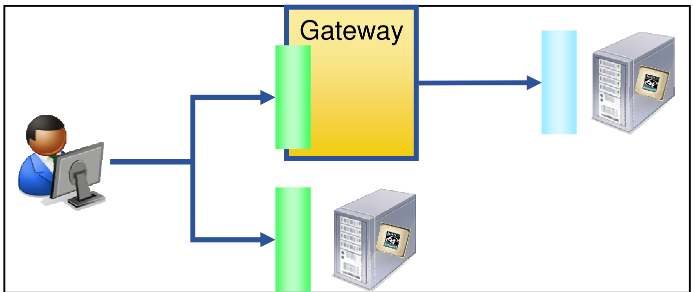

Figure 4: Gateway Approach

2.4. Adaptors and Translators

Adaptors, as shown in figure 5, allow two entities to be connected. Translators modify information so that it can be understood. Adapters and translators can be incorporated into the middleware so that it can work with both interfaces. This will require modifications to the grid middleware but it does mean that the existing interfaces can be used. Where and how the adapters and translators are used highlights the interfaces which need standardization. The ability to use multiple interfaces is a useful feature even when using standards to manage the evolution of the standard. A more detailed example of this approach is described in section 4.

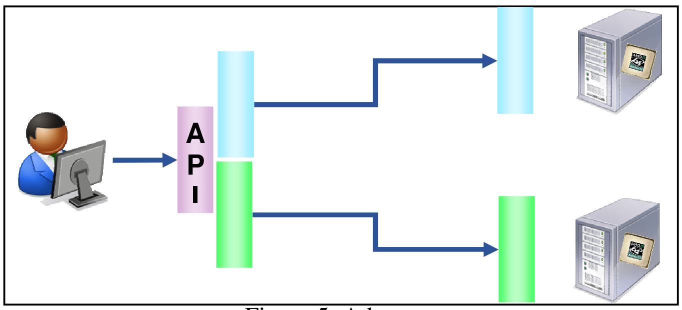

Figure 5: Adaptors

2.5. Common Interfaces

Common interfaces have already been proposed as the correct solution. However, with the absence of standards, which interface should be used? As a grid infrastructure has already heavily invested in one interface, moving to another interface may only be possible in the long term. Agreeing on which interface to use and the deployment of a production quality implementation across all infrastructures will take time. The following section explains how moving to common interfaces can help to achieve sustainable interoperation between grid infrastructures.

3. Interoperation

While interoperability can be achieved by the methods described or avoided by using the identical middleware, the problem of interoperation can not be avoided. There are many other issues which need to be addressed in order to achieve seamless interoperation. These can include, but are not limited to deployment, monitoring, support, troubleshooting, knowledge dissemination, accounting, policy etc. These issues can take significantly more time to address than achieving technical interoperability. Grid Interoperation is usually a bi-lateral activity between two grid infrastructures. Each bi-lateral activity brings us a step closer towards the overall goal of a uniform grid landscape. Recently, there has been an increasing emphasis on the need for interoperation and a number of bi-lateral interoperation activities have been initiated. The interoperation activity between EGEE[8] and OSG[9] was one of the first to achieve sustainable interoperation.

The first steps towards interoperability took place during November and December 2004. There was an initial meeting between the two projects where the problem was discussed. During this meeting an interoperability matrix was created showing the similarities and differences between the two middleware stacks. The information schema used in the information system was identified as the only major difference. The Glue[10] schema was originally created to facilitate interoperability between Grid3[11] and EDG[12], the predecessors of OSG and EGEE respectively. As such, OSG decided to move to the Glue schema on condition that a new version was defined which would then enable the OSG schema to be replaced. Over the course of the next few months a new revision of the Glue schema was defined.

In January 2005, a proof of concept was demonstrated. A test OSG site was deployed and the information schema changed to the Glue schema. This enabled a simple "Hello World" job to be submitted through the EGEE resource broker to an OSG CE and run successfully on an OSG WN. The EGEE clients were then installed using the standard mechanism for deploying application software and another job was submitted which carried out some basic data management operations. This proof of concept showed that submitting from EGEE to OSG was possible and the modifications that were required. The modifications necessary would need to be introduced into both the OSG and EGEE software releases. By summer 2005 the new Glue version was defined and all the necessary middleware changes had been incorporated in the respective releases.

Throughout August 2005 a number of steps were taken to include an OSG site into the EGEE operational framework. By November 2005 a number of sites had been included and the first user jobs from GEANT4[13] arrived on OSG sites. In March 2006 more discussion took place covering operations issues including the bootstrapping of the information system, the routing of trouble tickets, a joint operation virtual organization and joint operations meetings. By summer 2006 the CMS[14] physics collaboration were successfully using OSG sites via the resource broker. More discussions took place in summer 2007 on how interoperation could be sustained in the long term. One of the outcomes of the discussion was that a link is required in the software certification process to ensure interoperability is maintained.

4. Grid Interoperability Now

During an ad-hoc meeting at Super Computing 2005 in Seattle[15], representatives of more that eight different grid infrastructures decided to initiate the "Grid Interoperability Now" (GIN) Community Group[16] at OGF[17]. The aim of GIN is build upon the existing bi-lateral interoperation activities, share experiences and to work towards the common goal of grid interoperation. It was hoped that working solutions could be found on the short term and the experience gained will provide input to the standardization efforts. The GIN activity focuses on the four main areas of interoperability; security, information services, job management and data management. The first goal was to demonstrate various aspects of interoperability at Super Computing 2006. The information system area was very active and developed adaptors to query all the grid information systems from participating infrastructure. This information was translated into a common format and the result was inserted into a top-level information system aggregator. An interoperability matrix showed that the majority of infrastructures use the Glue schema and an LDAP based information aggregator. As such, it was decided to translate this information into the Glue schema format and insert it into a BDII[18]. This information was then used to plot the location of all the sites using Google Earth[19] and showed to which infrastructure each site belongs. This work highlighted the importance of a common schema and as a result the Glue schema is now an OGF working group.

5. Current Status

In order to minimize the interoperability problems and hence reduce the work required to interoperate, it is of critical importance that the most important interfaces are standardized. This section looks at the current status of standardization for the four main areas.

5.1. Security

The security model is the fundamental aspect of grid computing. Users belong to a virtual organization and do work on behalf of the virtual organization. A common security mechanism is required for all services in the grid infrastructure. The majority of grid infrastructures base their security model around X509 credentials. This is already an existing standard however, in order for this model to function, the root certificates of all the certificate authorities need to be managed and policies agreed. The work is coordinated by the IGTF[20] and has significantly reduced interoperability problems in this area. Further work is required on common methods for policy management with consideration for subgroup and roles within a virtual organization. Although this work has gone a long way to solving the policy problem, experience has shown that the current public private key approach can be challenging when it comes to performance.

5.2. Information System

For the information system it is important to separate the content, the interface and the topology. The schema defines the content, The Glue schema has helped to facilitate interoperation and is now an official OGF working group. Definition of version 2 is in progress and expected to be delivered soon. The GIN activity showed that LDAP is the dominant interface, 55% of grids and 95% of sites provide this interface. The other interfaces used are based on web services but these have shown problems with large query results. Although LDAP has been successful, the currently topology of existing information systems needs to be revised to address scalability limits.

5.3. Data Management

GridFTP[21] is supported in most grid infrastructures and has helped to reduce interoperability problems. The Storage Resource Manager[22] is a proposed interface to storage. During the development of the

SRM, there are been problems with different interpretations of the specification and incompatible implementations. Even though there has been a concerted effort to move from version 2.1 to 2.2, it has taken 18 months to get both the specification and implementation right.

5.4. Job Management

This is an area where a great deal of work on interoperability is neededFag as there are as many computing interfaces as batch systems. A number of efforts are underway in the OGF to address this area including JSDL[23] and OGSA-BES[24]. OGSA-BES version 1.0 is currently in draft and a number of prototypes already exist but are unproven in production. However, the current specification does not provide all the functionality required and a number of vendor specific extensions have been made which break interoperability.

6. Conclusion

As there are many different views on the definition of a grid, it is important to put "grids" into context before any discussion. The context of a grid is defined by the problem addressed. This paper has described grids in the context of *"Multi-institutional e-Science Infrastructures*".
Interoperability is "*The ability to exchange information and to use what has been exchanged*" and interoperation is "*The use of interoperable systems"*. As such grid interoperability is the ability of grid middleware to work together and grid interoperation is the ability of grid infrastructures to work together.

Grid interoperability is a second attempt at the original problem. The original solution was to provide common interfaces, most crucially at the organizational boundary. The solution is still to provide common interfaces however, as *different* common interfaces are in existence, the only real way forward is standardization. The most important part is to agree as the initial choice only defines the starting point and production feedback will ensure that the standard works. Interoperability can be overcome short term but only standards are sustainable in the long term. Infrastructures need to focus more on the standards and less on specific implementations.

Grid interoperability is an avoidable problem but grid interoperation is not. Once technical interoperability has been achieved, it is important to start looking at grid operations. Grid operations cover everything that is needed to operate a grid infrastructure. This includes deployment, monitoring, support, troubleshooting, knowledge dissemination, accounting, policy etc. The support teams within the different infrastructures may rely on different software tools but it is not necessary to harmonize these tools, however, it must be ensured that the tools will work with the other infrastructure. The procedures used on each grid infrastructure need to be analyzed to ensure that the necessary operations procedures can still be carried out with the additional institutions and virtual organizations. For example, ways to route trouble tickets between grid operations centers needs to be investigated.

The current grid paradigm is to provide a federated grid or "grid of grids" with the different grid federations working together to provide a seamless grid infrastructure. Even with technical interoperability assured, a truly federated grid brings a whole new set of operational challenges.

7. References

[1] *IEEE Standard Computer Dictionary: A Compilation of IEEE Standard Computer Glossaries*, 1990
[2] *ATIS Telecom Glossary*, 2000
[3] I. Foster et. al. The *Anatomy of the Grid: Enabling Scalable Virtual Organizations*, International J. Supercomputer Applications, 15(3), 2001.
[4] *Realizing the Information Future: The Internet and Beyond*, National Academy Press, 1994
[5] *Analysis of the ATLAS Rome Production Experience on the LHC Computing Grid*, First International Conference on e-Science and Grid Computing, 2005

[6] S. Gabriel, *GridKa Tier1 Site Management*, Proc of CHEP 2007
[7] *Job invocation interoperability between NAREGI Middleware Beta and gLite*, HPC Asia, 2007
[8] http://www.eu-egee.org/
[9] http://www.opensciencegrid.org/
[10] S. Andreozzi et. al. *Experiences with the GLUE information schema in the LCG/EGEE production Grid*, Proc of CHEP 2007
[11] http://www.ivdgl.org/grid3
[12] http://eu-datagrid.web.cern.ch/eu-datagrid/
[13] *Experiences in the Gridification of the Geant4 Toolkit in the WLCG/EGEE Environment,* Nuclear Science Symposium Conference Record, 2006
[14] *Experiment Dashboard: the monitoring system for the LHC experiments*, Proceedings of the 2007 workshop on Grid monitoring, 2007
[15] http://sc05.supercomputing.org/
[16] http://www.ogf.org/gf/group_info/view.php?group=gin-cg
[17] http://www.ogf.org
[18] L.Field, et. al. *Gird Deployment Experiences: The path to a production quality LDAP based information system* Proc of CHEP 2004
[19] http://earth.google.com/
[20] http://www.gridpma.org/
[21] *GridFTP Protocol Specification,* Global Grid Forum Recommendation, 2003
[22] *The Storage Resource Manager Interface Specification Version 2.2*, 2007
[23] *Job Submission Description Language version 1.0*, Global Grid Forum Recommendation, 2005
[24] *OGSA Basic Execution Service Version 1.0*, Open Grid Forum Recommendation, 2007