Mem. S.A.It. Vol. 80, 497

Memorie della

Towards Seamless Grid Computing

The EGEE Experience on Interoperable Grid Infrastructures

L. Field and E. Laure

CERN CH-1203 Geneva, Switzerland
e-mail: erwin.laure@cern.ch
e-mail: laurence.field@cern.ch

Abstract. Over recent years a number of production grid infrastructure projects have emerged that are now the computing backbones for various user communities. These infrastructures comprise of application independent services, which enables the user communities to access the computing resources at various computing centres around the world. Many of these user communities are collaborating across the administrative domains which operate their primary infrastructure and as such it is important these infrastructures interoperate. Grid interoperation actives are trying to bridge these differences and enable the users communities to gain access resources independent of their grid project affiliation. This paper gives an overview of the different methods that can be used to bridge the differences between grid infrastructures and gives examples of where they have been successfully applied. Arguments for standardization are presented along with the identification the the most important areas for which this is required. A summary of the current standardization efforts in these areas is given and as well as a direction for future work.

Key words. Grids – Interoperability – eScience – eInfrastructures – EGEE

1. Introduction

Modern science is increasingly dependent on Information and Communication Technologies (ICT), analysing huge amounts of data (in the TeraByte and PetaByte range), running large scale simulations requiring thousands of CPUs, and sharing results between different research groups. This collaborative way of doing science, also referred to as *eScience* or *CyberScience* has led to the creation of Virtual Organizations (VOs) that combine researches and resources (instruments, computing, data) across traditional administrative and organizational domains **?**. Advances in networking and distributed computing techniques have enabled the establishment of such Virtual Organizations and this concept, which is also referred to as Grid Computing **???** is being adopted by an increasing number of scientific disciplines.

The experience gained of over the past few years has demonstrated the importance of basing Grid computing on a well managed infrastructure federating the network, storage, and compute resources across different institutions and making them available to different

Send offprint requests to: E. Laure

scientific communities via well defined protocols and interfaces exposed by a software layer (Grid middleware). This kind of federated infrastructure is referred to as *eInfrastructure* or *CyberInfrastructure* and many projects, both national and international, have created dependable eInfrastructures which support eScience. Notable efforts include the US based Open Science Grid (OSG) and TeraGrid projects, the Japanese NAREGI project, the European EGEE and DEISA projects, as well as (multi)national efforts like the UK National Science Grid (NGS), the Northern Data Grid Facility (NDGF) in Northern Europe, and the German D-Grid.

The vision of Grid computing which underlies these efforts implies a federation of resources across institutional, disciplinary and (often) national boundaries, this is in contrast to 'enterprise Grids' which often exist within individual companies.

Together, these infrastructure projects have world-wide coverage and utilize a diverse set of hardware resources. Particularly, infrastructures like EGEE and OSG primarily federate centres with clusters of commodity PCs while DEISA and TeraGrid federate supercomputing centres. However, the technologies deployed by the different projects vary widely exposing different interfaces and protocols and offer different service characteristics. This poses significant challenges to many applications that need to harness resources operated by different Grid infrastructures. To overcome these challenges interoperability efforts are pursued at multiple levels. The strategies used are outlined in Section 2 and Section 3 reports on the experiences gained during interoperation activities with EGEE. We close the paper with some summary remarks and an outlook to future work.

2. Approaches to Interoperability

Ideally, all Grid infrastructures would expose the same interfaces and protocols with common semantics and as a result the fact that resources are operated by different providers would be completely transparent to the end user. This is the case in today's networks and on the world-wide-web where different providers all offer the same standardized interfaces and protocols. Unfortunately, in the Grid domain standardization is not as advanced and it is unlikely that in the short term major breakthroughs will be achieved. Hence, pragmatic approaches to interoperability have to be applied and in this section we review the most common techniques.

The first step towards interoperable services is to understand the differences in each infrastructure. For the most typical usecases, four areas need to be covered: security to ensure proper authentication and authorization across infrastructures, information servcies to understand what services are offered by each infrastructure and in what state these are, job management to launch and control computational jobs and finally data management to store, retrieve, and transfer data. An interoperability matrix can be used to show the different interfaces and protocols employed by the different infrastructures. Once these differences are understood, there are several methods for achieving interoperability, each of which is as outlined below.

2.1. User Driven Interoperability

The users (or virtual organisations) can themselves strive to achieve interoperability. 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 (cf. Figure 1). One of the problems with this approach is that it places significant effort on the virtual organisations. 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. This results in a keyhole approach where the minimum common subset of functionality is used and handling failures can be problematic. To help applications with this approach several generic software tools have been developed to hide the complexity from the user. For instance, Condor-G **?** can submit jobs to a variety of different middleware solutions; portals, such as the P-Grade portal **?** also provide a common abstraction and can internally talk to multiple middleware interfaces. This approach was used for instance by the Atlas community to overcome the problem of interoperation with OSG, EGEE and NorduGrid **?**.

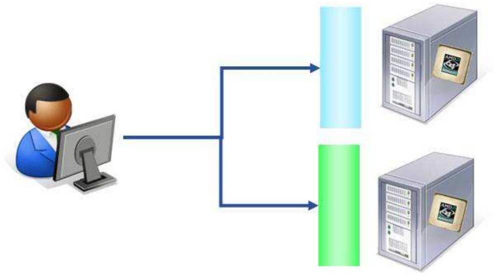

Fig. 1. User Driven Scenario

2.2. Interoperability through Parallel Deployment

Institutions can achieve interoperability by deploying multiple interfaces as show in Figure 2. 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 wishes to support and therefore this method is only affordable by large resource centres. For instance, the Forschungszentrum Karlsruhe used this approach to overcome the problem of interoperation with EGEE, Nordugrid and D-Grid **?**.

2.3. Interoperability through Gateways

A gateway, as shown in Figure 3, is a bridge between grid infrastructures. It is a specific service which makes the grid infrastructure look

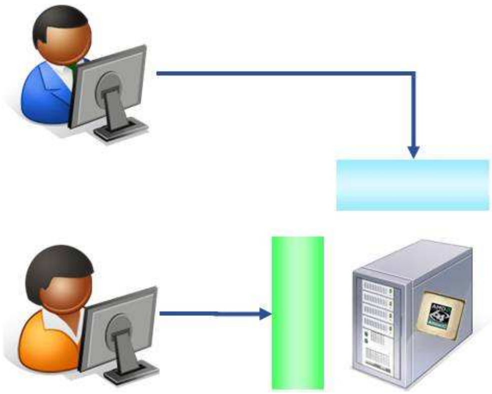

Fig. 2. Site Driven Scenario

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 **?**.

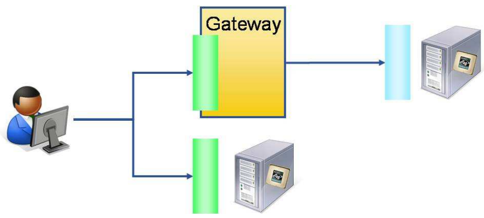

Fig. 3. Gateway Approach

2.4. Adaptors and Translators

Adaptors, as shown in Figure 4, 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.

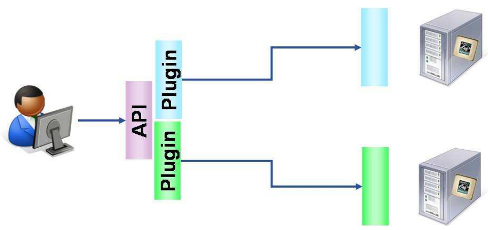

Fig. 4. Adaptor Approach

2.5. Common Interfaces

As already mentioned above, services that expose common interfaces, use the same protocols and offer similar semantics are the ideal way of achieving interoperability. The goal of the Open Grid Forum (OGF - http://www.ogf.org) is to harmonize theses different interfaces and protocols and develop standards that will allow interoperable, seamlessly accessible services. However, only a few established standards exist today and changing interfaces to services under production use is a difficult task. Hence this has to be seen as a more long term solution.

3. The EGEE Experiences

The EGEE (Enabling Grids for E-sciencE) project unites thematic, national and regional Grid initiatives in order to provide an e-Infrastructure available to all scientific research in Europe and beyond in support of the European Research Area **?**. The project is a multi-phase programme starting in 2004 and expected to end in 2010. EGEE has built a pan-European e-Infrastructure that is being increasingly used by a variety of scientific disciplines. EGEE has also expanded to the Americas and Asia Pacific working towards a world-wide e-Infrastructure. EGEE currently federates some 250 resource centres from 48 countries providing over 50.000 CPUs and several PetaBytes of storage. The infrastructure is routinely being used by over 5000 users forming some 200 Virtual Organizations and running over 140,000 jobs per day. EGEE users come from disciplines as divers as archaeology, astronomy, astrophysics, computational chemistry, earth science, finance, fusion, geophysics, high energy physics, life sciences, material sciences, and many more.

By working with these application domains as well as through close interactions with the peer infrastructures mentioned in Section 1 EGEE gained rich experiences in interoperability, in fact using all of the techniques mentioned above.

3.1. Case Study

As a case study we describe our experiences gained in working with OSG and NorduGrid to achieve interoperability between the infrastructures. As mentioned above the first step was to create an interoperability matrix as shown in Table 1

As can be seen all three infrastructures have many communalities, in particular in their security and data management tools. This facilitates interoperability, however, even though EGEE and OSG use the same interfaces and protocols for all aspects, differences in the exact version deployed can cause problems. For instance, OSG deploys both the pre-web-service and web-service version of GRAM while EGEE is only using the pre-web-service version; the information system schema used also differs.

The first step to achieve interoperability was the harmonization of the information systems. OSG and EGEE worked together on defining Glue schema version 1.2 and deployed the resulting implementation across their respective infrastructures. Nordugrid wished to keep the nordugrid schema as migrating their middleware and tools would require significant effort. As a result components were developed to translate information between the different information models. Both activities lead to the user having a full and consistent representation of the ser-

	NorduGrid	OSG	EGEE
Job Submission	GridFTP	GRAM	GRAM
Service Discovery	LDAP/GIIS	LDAP/GIIS	LDAP/GIIS
Schema	ARC	GLUEv1	GLUEv1.2
Data Transfer	GridFTP	GridFTP	GridFTP
Storage Management	SRM	SRM	SRM
Security	GSI/VOMS	GSI/VOMS	GSI/VOMS

Table 1. Interoperability Matrix

vices available on the combined infrastructure along with their meta-data and state.

The next problem to overcome were the differences in accessing compute resources. Currently, three different interfaces to compute resources are used and hence an adaptor approach has been applied. By using higher level submission systems that shield the differences in accessing comput resources seamless interoperability could be achieved. In particular, EGEE is using Condor **?** internally in its workload management system and thanks to Condor's ability to interface to all three systems jobs can be submitted to all resources available. Some applications also use Condor directly in their job submission tools or use other higher level services like GridWay **?** or portals that allow seamless job submission.

The final step was to ensure the clients needed by applications are available on the different infrastructures. This is either achieved by the infrastructure providing the different clients on their resources or by the users bundling the needed clients with their application code. Work towards defining the necessary subsets of clients required is currently ongoing.

Although achieving interoperability among these three infrastructures was relatively straight forward, it required significant and constant effort to maintain interoperability. Particularly, as the infrastructures evolve and services are upgraded or replaced interoperability might break. EGEE is testing interoperability as part of its standard certification process to ensure changes in the infrastructure will not interfere with the application workflows. This is clearly not a sustainable situation and hence common standards are needed to overcome these problems.

Apart from these interoperability activities, EGEE and OSG also engaged in interoperation, meaning efforts to harmonize the accounting, monitoring, and troubleshooting systems. A discussion of interoperation is out of the scope of this paper and the interested reader is referred to **?** for further details.

EGEE is also involved in interoperability activities with Unicore and NAREGI. Both have significantly different architectures with respect to information systems, job management and data management. Interoperability components have been created to overcome these differences using the interoperability methods, as outlined in section 2. Efforts are now underway to use the components to interoperation with the NAREGI infrastructure.

3.2. Standardization

As mentioned above, the usage of standardized interfaces and protocols is the best way of achieving interoperability. EGEE works closely with the Open Grid Forum (OGF—http://www.ogf.org) to bring its experiences into the standardization process and ensure upcoming standards are eventually usable. Note however that even if good standards are developed, a significant time is needed to get them adopted by major middleware providers and deployed across the infrastructure.

In addition to participation in various OGF working groups (see http://technical.eu-egee.org/standardisation.html), is very active and co-chairs the Grid Interoperation Now (GIN) community group. In this group major Grid infrastructures, such as EGEE, DEISA, TeraGrid, OSG, Naregi and many more, work to make their systems interoperate with one another through best practices and common rules. GIN does not intend to define standards but provide the groundwork for future standardization performed in other groups of OGF.

We give below an overview on the current status of the major GIN activities, more details can be found in **?**.

Security The majority of grid infrastructures base their security model on X509 credentials, however, some use proxy certificates while others use plain credentials which hampers interoperability. The root certificates of all the certificate authorities need to be managed and policies agreed. This work is coordinated by the IGTF 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.

Information System For the information system it is important to separate the content, the interface and the topology. The schema defines the content and the so-called *GLUE* schema has helped to facilitate interoperation and is now in the process of being published as an official OGF recommendation. While this standard is being more widely adopted, GIN produced a service discovery index of nine major Grid infrastructures, allowing users to discover the different services available from a single portal. This was achieved by defining a common set of necessary information and translators between the different schemata deployed.

LDAP is currently the dominant interface for managing the information, 55% of the infrastructures and 95% of their 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.

An information system topology defines how the information about resources is moved around the infrastructure and the aggregation points. By addressing the efficiency with respect to the movement of information around the infrastructures scalability limits overcome.

Data Management GridFTP **?** is supported in most grid infrastructures and has helped to reduce interoperability problems. The Storage Resource Manager (SRM) **?** is a proposed interface to storage and a recommendation of the OGF. While some applications prefer to use low level tools like GridFTP and SRM, others require more complete data management solutions such as the Storage Resource Broker (SRB) or iRODS that provide an integrated system with advanced functionality like rule based processing **?**.

Job Management This is an area where a great deal of work on interoperability is needed as there are as many Grid interfaces to compute resources as batch systems. A number of efforts are underway in the OGF to address this area including JSDL and OGSA-BES . 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.

4. Conclusions and Future Work

The concept of the Virtual Organization has enabled scientific collaborations to use resources from many administrative domains. Sustainable e-infrastructures have emerged to effectively manage the coordination of administrative domains on behalf of the VOs and ensure that service level agreements are met. As these infrastructure mature, VOs wish to use resources that may reside in different e-infrastructures. As a result of these infrastructure have evolved independently interoperability and interoperation issues arise. Interoperability can be overcome by using the techniques that have been outlined in the paper and a number of bilateral activities have demonstrated their effectiveness.

However, there are limitations to such methods and only standardization of the core interfaces can be sustainable on the long term. The key areas that need addressing are security, information systems, job management and data management. There is still a great deal of work that needs to be done in each area, each of which has had differing successes at standardization. Currently the most pressing need is to address the area of job management. Although a number of solutions have been presented, it still has to be seen whether they will meet the needs of the production infrastructures.

Acknowledgements. This work is co-funded by the European Commission through the EGEE-II project, contract number INFSO-RI-031688.

References

Allcock W. et al. GridFTP Protocol Specification Global Grid Forum Recommendation GFD.20, March 2003.

Campana S., et al. Analysis of the ATLAS Rome Production experience on the LHC Computing Grid In *e-Science and Grid Computing, 2005. First International Conference on e-Science and grid Computing*, Melbourne, Australia, December 2005.

Field L. and Schulz M.. Grid Interoperability: The Interoperability Cookbook, in *CHEP 2007*, Victoria, Canada, September 2007.

Floros V. et al. EGEE User Forum - Book of Abstract Technical Report EGEE-TR-2007-002, EGEE, 2007.

Foster I. and Kesselman C., editors. *The Grid: Blueprint for a New Computing Infrastructure* Morgan Kaufmann, 2003.

Frey J., T. Tannenbaum, M. Livny, I. Foster, and S. Tuecke. Condor-G: A Computation Management Agent for Multi-Institutional Grids *Cluster Computing*, 5(3):237–246, 2002.

Gabriel S., GridKa Tier1 Site Management. In *CHEP 2007*, Victoria, Canada, September 2007.

Huedo E., Montero R., and Llorente I., The GridWay framework for adaptive scheduling and execution on grids *Scalable Computing: Practice and Experience*, 6(3), 2005.

Foster I. and Kesselman C. and Tuecke S., The Anatomy of the Grid *The International Journal of High Performance Computing Applications*, 15(3):200–222, Fall 2001.

Kacsuk P. and Sipos G., Multi-Grid, Multi-User Workflows in the P-GRADE Grid Portal, *Journal of Grid Computing*, 3(3–4), September 2005.

Lamanna M.and Laure E., editors, *The EGEE User Forum - Experiences in using Grid infrastructures*, volume 6(1) of *Journal of Grid Computing*. Springer, March 2008.

Laure E. and Jones B., *Enabling Grids for e-Science: The EGEE Project*, CRC Press, to appear.

Nakada H., et al., Job invocation interoperability between NAREGI Middleware Beta and gLite, in *HPC Asia*, Seoul, Korea, September 2007.

Rajasekar A., et al., A Prototype Rule-based Distributed Data Management System, in *HPDC Workshop on Next Generation Distributed Data Management*, Paris, France, May 2006.

Riedel M., et al., Interoperation of World-Wide Production e-Science Infrastructures, *Concurrency and Computation: Practice and Experience*, to appear.

Sim A. and Shoshani A., The Storage Resource Manager Interface Specification Version 2.2, Open Grid Forum Recommendation GFD.129, April 2008.