# Interoperation Scenarios of Production e-Science Infrastructures

Morris Riedel
Central Institute of Applied Mathematics
Forschungszentrum Juelich
D-52425, Juelich, Germany
m.riedel@fz-juelich.de

Erwin Laure, Laurence Field, Gidon Moont, Mona Aggarwal, Olivier van der Aa, David Colling
Institution2
First line of institution2 address
Second line of institution2 address

Gerson Galang, Timur Perelmutov, Riccardo Zappi, Luca Magnoni, Balazs Konya, Aleksandr Konstantinov
Institution2
First line of institution2 address
Second line of institution2 address

Shaowen Wang, Liegh Grundhofer, Bob Cowles, Frank Wuerthwein, Abishek Singh Rana, Dane Skow
Institution2
First line of institution2 address
Second line of institution2 address

Yuji Saeki, Hitoshi Sato, Satoshi Matsuoka, Yoshio Tanaka, Yusuke Tanimura, Putchong Uthayopas
Institution2
First line of institution2 address
Second line of institution2 address

Matt Viljoen, Glenn Moloney, Jerome Lauret, Vijaya Natarajan, Junmin Gu, Oxana Smirnova
Institution2
First line of institution2 address
Second line of institution2 address

Sugree Phatanapherom, Somsak Sriprayoonakul, Cindy Zheng, Thomas Soddemann, Surya Pathak
Institution2
First line of institution2 address
Second line of institution2 address

Laura Perlman, Are Shoshani, Alex Sim, Somsak Sriprayoonsakul, Charlie Catlett, JP Navarro
Institution2
First line of institution2 address
Second line of institution2 address

Oscar Koeroo, Shaun De Witt, Jiri Mencak, Jean-Ph. Baud, Paul Sheldon
Institution2
First line of institution2 address
Second line of institution2 address

Olle Mulmo, James Casey, Maarten Litmaath, Vincenzo Ciaschini, Von Welch, Jim Marsteller
Institution2
First line of institution2 address
Second line of institution2 address

Stephen Pickles, Stephen Newhouse, Peter Arzberger, Phil Papadopoulos, Mason Katz, Tony Rimovsky
Institution2
First line of institution2 address
Second line of institution2 address

## Abstract

*TBD*
*TBD*
*TBD*
*TBD*
*TBD*
*TBD*
*TBD*
*TBD*
*TBD*
*TBD*
*TBD*


## 1. Introduction

Many Grid projects have begun to offer production services to end-users during the past several years with an increasing number of application projects that require access to a wide variety of resources and services in multiple Grids. Therefore, the prupose of the *Grid Interoperation Now (GIN) Community Group (CG)* of the *Open Grid Forum (OGF)* is to organize, manage and demonstrate a set of interoperation efforts among production Grid projects and e-Science infrastructures using computational or storage-related resources in multiple Grids.

Within this contribution, we define the difference between *interoperation* and *interoperability* as follows. Interoperation is specifially defined as what needs to be done to get production Grids to work together as a fast short-term achievement using as much existing technologies as available today. Hence, this is not the perfect solution and different than interoperability that is defined as the native ability of Grids and Grid technologies to interact directly via common open standards in the near future.

The GIN-CG group within OGF implements interoperation in five specific areas. First, *authorization and identity management (GIN-AUTH)* deals with resource sharing among members of the GIN Virtual Organization (VO) [**?**]. Second, the *data management and movement (GIN-DATA)* area is working on the interoperation of different data management technologies currently in use of multiple e-Science infrastructures. Third, the *job description and submission (GIN-JOBS)* area focuses on job management across different Grid technologies and middlewares used in production Grids today. One of the most important areas is the *information services and schema (GIN-INFO)* area, because the efforts conducted in this area basically provide the base for cross-Grid interoperation taking up-to-date information into account. Therefore this paper emphasizes on this important area, but also gives insights to the other areas. Finally, the *operations experience of pilot test applications (GIN-OPS)* for cross-Grid operations works on different applications that require resources from different Grids.

The contribution of this paper to the overall work in the field of interoperability and interoperation within the e-Science and Grid communities is that GIN provides components and adapters that work today and include the most known production Grid infrastructures. Hence, this paper describes work that has been made to basically enable e-Scientists to work in more than one production Grid tomorrow if they want to. This includes different working areas such as secure job submission, data transfers, or information enquieries.

Finally, it is important to mention that the GIN effort did not include any attempt to provide a common allocation or brokering of resources between production Grid projects. This is viewed as beyond the scope of the GIN efforts and resource allocation decisions are left to negotiations between e-Science projects and the individual e-Science and Grid infrastructures. Nevertheless, the work within GIN demonstrates that interoperation is feasible and technically possible today. Thus basically enabling e-Scientists to work with cross-Grid scenarios and applications that need more than one Grid tomorrow.

This paper is structured as follows. After the motiviation and introduction into the problem domain of interoperation and interoperability, Section 2 describes the fundamental process of prodiving a cross-Grid information system. Section 3 describes the particularly difficult process of achieving interoperation in the area of cross-Grid job submissions and potential solutions, while Section 4 summarizes the efforts that have been done in the context of data transfer and movement. The fundamental approach for VO-based identity management within GIN is given in Section 5, while Section 6 highlights some applications used during cross-Grid interoperations. Finally, this paper ends with a survey of related work and concluding remarks.

## 2. Information Services and Modelling

In order to identify appropriate resources for end-users within an e-Science infrastructure there must be some form of resource information conforming to schemas and access technologies using standard query mechanisms. The GIN-INFO area of GIN provides interoperation components (e.g. information schema translators, adapters, etc.) to deal with the usage of different information schemas within production e-Science infrastructures. The efforts in the GIN-INFO area are build upon previous bi-lateral successes such as interoperation efforts of EGEE and OSG [**?**], NDGF and EGEE [**?**] as well as other pair-wise interoperation efforts within production Grids. Hence, the major goal of GIN-INFO is to extend these pair-wise interoperations with a broader set of production Grids to identify a subset of information items that can be used as a common minimum set. This also motivates the development of translators for these information items used in different information schemas in production e-Science infrastructures today.

In more detail, the wide variety of Grid resources that are available in production Grids are described in a precise and systematic manner to be able to be discovered for subsequent management or use. Over the years several schemas evolved either in scientific-oriented production Grids or within business-oriented standard bodies such as the Distributed Management Task Force (DMTF). GIN-INFO currently provides components to fetch information out of nine production e-Science infrastructures that use different information services and schemas. These are APAC, DEISA, EGEE, OSG, PRAGMA and TeraGrid that all use the GLUE schema, NAREGI that is based on the CIM schema, NDGF that relies on the ARC schema and finally NGS that uses the MSD2.4 schema of Globus Toolkit 2.

GLUE etc.

TBD: PICTURE WITH INFORMATION PROVIDERS etc.

Finally, the interoperation of technologies described here where demonstrated at the Supercomputing 2006 by using Google Earth showing information of all participating Grid sites. More information can be found in the GIN-INFO area of GIN on GridForge [**?**].

## 3. Job Submission and Management

There are a lot of production e-Science infrastructures that all support wide varieties of Grid middleware platforms and technologies that typically provide no commonly accepted interoperable interface for job submission and management. While the gLite middleware of EGEE uses the proprietary JDL job description, the Globus Toolkit's GRAM accepts job descriptions in a proprietary RSL format, and also within UNICORE-based Grids the job description is typically based on AJOs, just to list some. Hence, there is currently no standardized job description format in use within production Grids and no well-defined interface for job submission and management broadly adopted.

The OGSA - Basic Execution Services (OGSA-BES) specification provides such an interface that accepts job descriptions in the standardized Job Submission and Description Language (JSDL) format. Both specifications together with some HPC extensions [] to JSDL are named as the High Performance Computing (HPC) - Profile [**?**]. In the last months, many software providers have already started to support the HPC-Profile and thus production Grids are already considering of evaluating this interface for production scenarios in the near future. Therefore, the GIN-JOBS area focused on this interface to provide a proof-of-concept interoperation demonstration before this interface comes into production.

It was commonly agreed within GIN that the use of HPC-P makes more sense than providing yet another interoperable short-lived adapters for Globus GRAM, UNICORE, or gLite environments. Also many other commercial vendors (e.g. Platform Computing, Microsoft, IBM, Fujitsu, etc.) agreed to provide such an implementation of this interface for their technologies. Finally, the JSDL specification is already standardized while the HPC extensions and the OGSA-BES specifications finally are through public comments and it can be expected that it becomes an OGF standard very soon.

Several interoperation efforts where demonstrated at the supercomputing 2006 and particularly these demonstrations lead to high visibility within the media in computer science. Many software providers and industrial vendors provided an implementation of the HPC-P with a dedicated version of OGSA-BES and thus interoperation among a lot of technologies was possible for the first time in the important area of job submission and management.

Finally, it is important that the OMII - Europe project augments currently gLite, UNICORE, and the Globus Toolkit with OGSA-BES interfaces to lay the foundation for its adoption by the middleware providers. Hence, this in turn lays the foundation for the usage of this interface and HPC-P profile within production e-Science infrastructures in the near future. This will lead to at least three independent implementations of OGSA-BES while there are a lot of other adopters and thus the OGSA-BES specification and HPC-P will become a full standard of OGF.

## 4. Data Management and Movement

In order to move and transfer data between production e-Science infrastructures they must be interoperable in terms of technologies that allow for data transfers with high performance such as GridFTP [**?**] or data brokering technologies such as SRB and SRM. Therefore, the GIN-DATA area is working on an GridFTP interoperation to ensure production level data transfers verified by test suites and an enumeration of clients. This is in particular a reasonable challenge, because the most production Grids use different versions of GridFTP. Beside these efforts for data movement, the efforts for data management are important as well and focus on the access to storage via standard interfaces such as SRM. Therefore the GIN-DATA area also achieved SRM interoperations using test suites and enumerations of different SRM versions finally leading to a SRM specification subset that was used for GIN interoperation scenarios. In addition, the access to whole data collections is provided by working on an SRB interoperation along with a trust establishment between different SRB-based data Grids.

More from GIN-DATA experience doc...

Finally, the highlights in interoperation scenarios described here and some other related efforts where demonstrated at the Supercomputing 2006 by using the GridFTP, SRM and SRB technologies from the different participating Grid projects. More information can be found in the GIN-DATA area of GIN on GridForge citeGINDATAWIKI.

## 5. Authorization and Identity Management

A functional authentication and authorization system is foundational to most of the other interoperation activities within e-Science infrastructures. The base use case is identity-based authorization at the service level with any further authorization restrictions enforces inside the service at the base internal system level. This includes scenarios by setting Grid permissions manually for members of specific groups or for end-users with a certain role possession.

These functionality is typically provided by a Virtual Organization Membership Service (VOMS) and thus GIN-AUTH provides a GIN VOMS service. VOMS is widely adopted in production Grids today and the two basic services that are provided by VOMS are the management of a list of members of a VO and the signing of attribute statements attesting to the subject's group/project membership or role possesions. However, initial tests with Grid interoperation leads to the demand of an Acceptable Use Policy (AUP) for VO membership.

In more detail, the creation of the GIN VO for testing purposes of a limited number of staff from the participating Grids introduced another point of confusion for end-users of the system. Frequently, it was mis-understood that membership in the GIN VO was the method by which one gained access to resources from GIN participating Grids to establish cross-Grid application interoperation. This was a persistent problem because part of the GIN baseline activity was a standard series of application tests to establish functional interoperation. This was also a problem, because the GIN VO had pre-negotiated access to all the participating Grids, a step which was viewed as a significant barrier to e-Science VOs wishing to get started with multi-Grid interoperations. Therefore, the GIN-AUTH group developed an clearly necessary AUP for the GIN VO which can serve as a model for other VOs wishing to establish serious multi-Grid interoperations. It was agreed among the participating Grids that this AUP met most of the requirements of them and establihshed a good baseline for VOs wishting to register with new Grids. There may be additional requirements for individual Grid or e-Science infrastructures, but those are typically few and deal with Grid-specific policies and procedures.

More challenges occur during the accrediation of CAs currently in use within e-Science infrastructures. Several of the Grids have internal processes for vetting CAs from whom they will accept credentials and there was no universal system for selecting or ranking a common set of CAs. Therefore, the GIN-AUTH team took the decision to concentrate on the International Grid Trust Federation (IGTF) set of regional Peer Management Authorities (PMAs) list of accredited CAs. Hence, these represent a common set of recognized identity providers. While this decision allowed us to clearly identify a common set of mutually acceptable identity sources and a process for accrediting new ones, there were a few residual problems which were uncovered.

Despite the agreement on credentials from IGTF sources being the commonly accepted set of credentials, end-users frequently made the presumption that because Grids X and Y are participating in GIN, that any credential which worked with Grid X would also work for communicating with Grid Y. Since there remain several Grids which recognize local CA's for internal purposes, that presumption is incorrect and leadt to much confusion and frustration in end-users getting started with establishing their interoperation between Grids. It is particularly difficult for end-users to recognize beforehand when dealing with service credentials issued by a local CA (non IGFT). In this sense, GIN-AUTH strongly recommend and encouraged e-Science infrastructure administrators that any service from multi-Grid interoperation use credentials issued by an IGFT accredited source to avoid such problems in the next years.

## 6. Cross-Grid Applications and Operations

GIN-OPS
Portal work etc.

## 7. Related Work

There are many pair-wise efforts in the field of interoperation and interoperability within e-Science communities and Grid projects.

EGEE and OSG, NorduGrid

DEISA and EGEE

## 8. Conclusions

TBD:summarized conclusions from the experience docs...

## References


[1] I. M. Author. Some related article I wrote. *Some Fine Journal*, 99(7):1–100, January 1999.

[2] A. N. Expert. *A Book He Wrote*. His Publisher, Erewhon, NC, 1999.