Interoperation Scenarios of Production e-Science Infrastructures

M. Riedel
Forschungszentrum Juelich
D-52425 Juelich, Germany
m.riedel@fz-juelich.de

E. Laure, L. Field, J. Casey, M. Litmaath, J.Ph.Baud
European Organization for Nuclear Research
CH-1211 Geneva, Switzerland

C. Catlett, D. Skow, JP Navarro
Argonne National Laboratory
Argonne, IL, USA

Th. Soddemann
Rechenzentrum Garching
D-85748 Garching, Germany

C. Zheng, P.M. Papadopoulos, M. Katz
San Diego Supercomputer Center
La Jolla, CA, USA

O. Smirnova, B. Konya
Lund University
SE-221 00 Lund, Sweden

P. Arzberger, F. Wuerthwein, A.S. Rana
University of California
San Diego, USA

V. Welch, T. Rimovsky
University of Illinois
Urbana, IK, USA

S. Newhouse
Microsoft
Seattle, USA

Y. Tanaka, Y. Tanimura, T. Ikegami
National Institute of Advanced Industrial Science and Technology
Japan

D. Abramson, C. Enticott
Monash University
Australia

R. Pordes, N. Sharma, S. Timm
Fermilab Computing Division
Batavia, IL, USA

G. Moont, M. Aggarwal, D. Colling, O. van der Aa
Imperial College London
London SW7 2BW, UK

A. Sim, V. Natarajan, A. Shoshani, J. Gu
Lawrence Berkeley National Laboratory
Berkeley, CA, USA

G. Galang
SAPAC, University of Adelaide
Australia

R. Zappi, L. Magnoni, V.Ciaschini
INFN - CNAF
Bologna, Italy

B. Cowles
Stanford Linear Accelerator Center
Menlo Park, CA 94025, USA

S. Wang
University of Iowa
Iowa City, USA

Y. Saeki
National Institute of Informatics
Japan

H. Sato, S. Matsuoka
Tokyo Institute of Technology
Tokyo, Japan

P. Uthayopas, S. Sriprayoonsakul
Kasetsart University
Bangkok, Thailand

O. Koeroo
NIKHEF
Amsterdam, Netherlands

M. Viljoen
National Grid Service
UK

L. Pearlman
USC Information Services Institute
Marina del Rey, CA, USA

S. Pickles
University of Manchester
Manchester, UK

G. Moloney
University of Sydney
Sydney, Australia

J. Lauret
Brookhaven National Laboratory
Upton, NY, USA

J. Marsteller
Pittsburgh Supercomputer Center
Pittsburgh, PA, USA

P. Sheldon, S. Pathak
Vanderbilt University
Nashville, TN, USA

S. De Witt, J. Mencak
Rutherford Appleton Laboratory
Oxon, UK

S. Phatanapherom, S. Sriprayoonakul
Grid Technology Research Center
Tsukaba, Japan

Abstract

Many production Grid and e-Science infrastructures have begun to offer services to end-users during the past several years with an increasing number of scientific applications that require access to a wide variety of resources and services in multiple Grids. Therefore, the Grid Interoperation Now (GIN) - Community Group (CG) of the Open Grid Forum (OGF) organizes and manages interoperation efforts among production Grid projects. The group demonstrated their efforts at the Supercomputing 2006 conference and since then keep continuing interoperation efforts to reach the goal of a world-wide Grid vision on a technical level in the near future and provide feedback to standardization. This paper describes several highlights and fundamental approaches taken within the GIN-CG to reach interoperation of production Grid and e-Science infrastructures in the areas of information services and modeling, authentication and authorization, job submission and management, data movement, as well as cross-Grid applications.

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 purpose of the *Grid Interoperation Now (GIN) Community Group (CG)* [7] 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 specifically 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) [20]. 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. These include the Storage Resource Broker (SRB) [23], Storage Resource Managers (SRM) [25] and GridFTP [15]. Third, the *job description and submission (GIN-JOBS)* area focuses on job management across different Grid technologies and middlewares used in production Grids today that have been augmented with OGSA-Basic Execution Service (OGSA-BES) [8] interfaces to conform to the High Performance Computing - Profile (HPC-P) [9]. 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. These interoperations rely on information models such as *Common Information Model (CIM)* [18] and *Grid Laboratory Uniform Environment (GLUE)* [5] or information systems such as *Berkeley Database Information Index (BDII)* [2] and *Monitoring and Discovery Services (MDS)* [24]. 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 multiple Grid infrastructures.

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 newly developed components and adapters that work today and include efforts within the most known production Grid and e-Science infrastructures. Hence, this paper basically describes various approaches that 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 enquiries.

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, e-Scientists 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 on cross-Grid scenarios and applications that need more than one Grid tomorrow.

This paper is structured as follows. After the motivation and introduction into the problem domain of interoperation and interoperability, Section 2 describes the fundamental process of providing 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 Modeling

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 [4] and OSG [11], NDGF [6] 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.

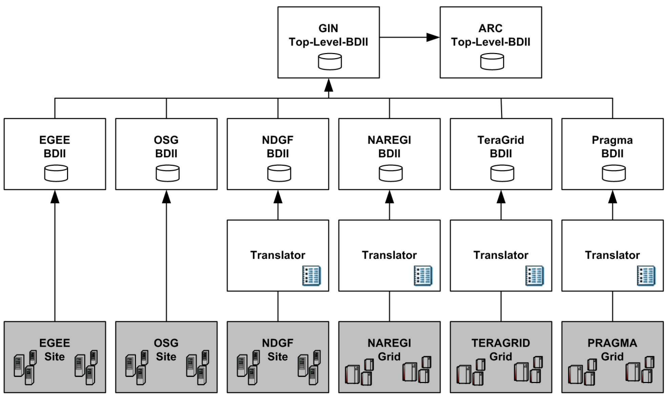

Figure 1. Information schema translators for production Grid interoperations.

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)*. Figure 1 illustrates the initial architecture of the BDII-based interoperation between different information services and schemas. BDII itself consists of two or more standard LDAP [10] databases that are populated by an update process. The update process obtains LDIF [13] from either doing an ldapsearch on LDAP URL lists or by running a local script that generates LDIF. Afterwards, the LDIF is then inserted into the LDAP database as an entry.

Within GIN-INFO, we initially developed numerous information translators for those Grids that not naturally support the BDII system. This raised a demand for a set of common minimal attributes, because the production Grids evolved over time and thus use a wide variety of technologies that deal with more or less the same pieces of information, but in different sets or named as different attributes using different information schemas. Therefore, when doing interoperation, the GIN-INFO group revealed that is is extremely important to agree on a *common minimal set of attributes*. Of course this common minimal set depends on the use case, for instance in the area of job execution the service endpoint attributes of a production Grid that offer job execution interfaces may be published by an information system. In addition, to the requirement of having a common minimal set of attributes for the content of information it is also necessary to map the different schemas to each other in order to put the contents into the right places within other schemas. Hence, translators have been developed that map the content of one attribute in schema A of the content of one attribute (or even a set of attributes) in schema B. They are available to be used within the Grid community.

As shown in Figure 1, we used a GIN-BDII as a top-level BDII with information (attributes) from all Grids in accordance to the GLUE schema in version 1.3 [5]. In particular, the GLUE schema has a *site entry* element that is able to provide such information needed during interoperation. We used mandatory attributes for essential interoperation scenarios such as a unique site identified, site location (longitude, latitude) or site name. In addition also several optional attributes such as site description, site location in a human readable format, administrator emails or site web page. Hence, all the mandatory attributes for a Grid site are provided by the GIN Top-level-BDII.

In this context, the GIN-INFO group faced another challenge related to a general *lack of information*. Not all information items or attributes can be found in all Grids that lead to missing data for one Grid site when using it in conjunction with other Grids that provide such information. Therefore, an cross-Grid application implementation may have errors that requires exactly that information. Of course these problems also arise if Grid sites not publish this information correctly (e.g. not schema compliant). In this context it seems reasonable to consider that this information could be information related to service discovery of job execution endpoints (e.g. OGSA-BES [8] interfaces). Thus, missing or incorrect data could lead to problems when using this data in real use cases such as data movement between sites or simple job execution. Hence, the lack of information disturbs the work of e-Scientists. In other words, the pieces of information required will be defined by the use cases and thus raise the demand to identify pieces of information required for cross-Grid use cases.

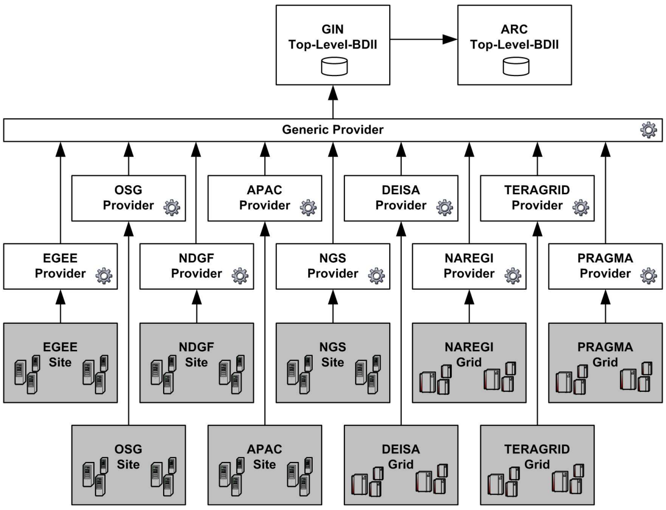

Figure 2. Information providers for different Grid systems.

The final approach uses the same set of attributes but is slightly different to the initial architecture due to administration overheads of providing a BDII for each Grid. Instead, Figure 2 illustrates that for each production Grid an information provider has been developed, re-using information from the translators. During production scenarios, the information provider queries the correspondent Grid site and outputs the information in LDIF in accordance to the GLUE schema. In turn these information providers are used to provide the GIN Top-Level-BDII with information from all participating infrastructures using a generic provider in between. In turn these pieces of information lay the foundation for the ARC-BDII that is also a Top-Level-BDII, but with information from all Grid infrastructures in accordance to the ARC schema. The ARC schema is used within the NDGF while its ARC middleware implements a scalable and dynamic distributed information system.

All in all, GIN-INFO provides components to fetch information out of nine production e-Science infrastructures that use different information services and schemas, namely APAC [1], DEISA [3], EGEE [4], NDGF, NGS, NAREGI, PRAGMA [12] and TeraGrid [14]. All use the GLUE schema, while NAREGI uses the CIM schema, NDGF relies on the ARC schema and NGS uses the MSD2.4 schema of Globus Toolkit 2. Finally, the described interoperations were demonstrated at the Supercomputing 2006 by using Google Earth showing information of all participating Grid sites. This basically include the common minimal set of attributes provided by various infrastructures and thus demonstrates that it is possible to interoperate to exchange pieces of information. More information can be found in the GIN-INFO area of GIN on GridForge [7].

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 unfortunately provide *no commonly accepted interoperable interface for job submission and management*. While the gLite [21] middleware of EGEE uses the proprietary *Job Description Language (JDL)*, the Globus Toolkit's GRAM [19] accepts job descriptions in a proprietary *Resource Specification Language (RSL)* format, and also UNICORE 5 [26] uses a proprietary job description named as *Abstract Job Objects (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 within e-Science infrastructures.

The *OGSA - Basic Execution Services (OGSA-BES)* [8] specification provides such an interface that accepts job descriptions in the standardized *Job Submission and Description Language (JSDL)* [16] format. Both specifications together with some HPC extensions to JSDL are named as the *High Performance Computing (HPC) - Profile* [9]. In the last months, many software providers have already started to support the HPC-Profile and thus many production Grids are already evaluating the implementation of this interface in the correspondent middlewares 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 usage within real application scenarios.

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. They basically agreed because the JSDL specification is already standardized while the HPC extensions and the OGSA-BES specifications are mature enough to become an OGF emerging 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 and news sections of online newspapers and reports. Many software providers and industrial vendors provided an implementation of the HPC-P with a dedicated version of OGSA-BES (version 26) and thus interoperation among a lot of technologies was possible for the first time in the important area of job submission and management. The GIN-JOBS group used *Transport Level Security (TLS)* in combination with the *WS-Security Username Token Profile* [22] as the security mechanism. Even if this kind of security can be significantly improved, the interoperation was focusing on the interface level of OGSA-BES and the HPC-P in order to be successful.

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 and provides stronger security mechanisms such as VOMS [17] during job submissions. 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 change its status from proposed standard recommendation to full standard recommendation within 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 [15] or data brokering technologies such as SRB [23] and SRM [25]. Therefore, the GIN-DATA area is working on interoperation of three different technologies. Firstly, the GridFTP technology that can be seen as lowest common denominator for data movement in Grids today. Secondly, SRM as an asynchronous standard interface to Grid storage systems to provide virtualization of storage resources. Finally, SRB as a shared collection management system of multiple (distributed) storage systems and their properties in order to provide virtualization of a whole shared data collection and thus data federation. All these three technologies are widely deployed in production Grid infrastructures today.

An GridFTP interoperation to ensure production level data transfers was 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, however basic interoperation was achieved.

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. By using this subset nine production Grid sites where able to interoperate. In more details, tools for validating and interoperation of SRMs for heterogenous storage systems have been developed within GIN. SRMs are based on a common interface specification, but nevertheless not all implementations provide all features of the SRM interfaces. The error response when asking for a non-implemented feature is in the most cases not clear enough to understand the real reason without contacting the administrator. Another challenge that the GIN-DATA group encountered was that the most SRM implementations within production Grids are tuned to use GridFTP as underlying data transfer protocol and thus run into problems when other protocols (e.g. simple HTTP) are being used in other production Grids.

Hence, SRMs can have different implementations and services for the underlying storage systems. The purpose of the newly developed SRM testing tool is to check compatibility and interoperability of SRMs according to the specification. Thus, the tool checks client-SRM adherence to the specification, as well as SRM-SRM remote copy capabilities. In more detail, the tests include the read access to a file in a remote Grid storage managed by SRM and also the write access to a registered user account in a remote Grid storage managed by SRM. Furthermore, a file replication for a registered user between two independent storage systems is tested. Finally, also space reservation and write access to the reserved space for a registered user in a remote Grid can be tested by the tool. To sum up, this tool ensures the possibility for production Grid interoperation with SRM technologies. It is being used by the SRM-collaboration to ensure the interoperation with SRM.

In addition to the efforts around SRM implementations, the access to whole data collections is one task of GIN-DATA by working on an SRB interoperation along with a trust establishment between different SRB-based data Grids. In more detail, the SRB interoperation tests initially focused on single file replication and subsequently on replication of data collections between multiple federated Grid infrastructures. This includes the federation of data Grids based on the SRB and replication of a collection into a federated data Grid. Another useful test is the use of remote procedures to extract provenance metadata and load the provenance metadata into a collection. Finally, tests have been done to test the use of extensible schema mechanisms to manage hierarchical metadata descriptions. To sum up, these tests were successfully run between 19 different Grid sites. The most identified problems and issues for interoperation during these tests were the establishment of trust between the Grids that basically requires manual intervention today, the interoperation between different versions of SRB, and the identification of what metadata will be shared is typically not straight forward.

Finally, the 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. The SRM and SRB demonstrations itself are not able to interoperate and therefore these demonstrations were named as *SRM islands* and *SRB islands*. More information can be found in the GIN-DATA area of GIN on GridForge [7].

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 possessions. 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 established a good baseline for VOs wishing 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. The AUP is publicly available at [7] for use in e-Science infrastructures that would like to engage in interoperation scenarios.

More challenges occur during the accreditation 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 lead 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 for 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

This section highlights several results of the interoperation efforts within the GIN-OPS area that focuses on the interoperation of real applications to help the different areas of GIN to verify their interoperation ability. One example of interoperation that was demonstrated at the Supercomputing 2006 conference was using the scientific program *Time-Dependent Density Functional Theory (TDDFT)* equation. TDDFT is a molecular simulation method in computational quantum chemistry, and it treats multiple electrons excitation directly, using techniques from quantum mechanics. The efforts include interoperation in different scenarios, for instance between a run of the TDDFT application across the PRAGMA Grid and TeraGrid. In particular this interoperation was achieved by running TDDFT on four heterogenous sites across both Grids, but also demonstrated that a level of interoperation is neither automatic nor unattainable.

In addition to TDDFT jobs, the GIN-OPS group used an data-intensive application named as the *Savannah fire simulation*. Typically, each 21 year simulation takes about 6-12 weeks to execute on a single processor. By using resources from PRAGMA and GIN, the execution time was significantly increased. These experiments also reveal a clear demand for job submission and management standard. Beside these issues the GIN-OPS group revealed challenges in software support environments. Hence, some Grids require site administrators to install a specific application (e.g. necessary libraries) and some Grids require applications that are running rather self-contained within a sandbox and thus e-

Scientists have to package all software that is needed. While the wider adoption of the sandbox method can be an option in interoperation, the GIN-OPS group worked on a *community software area (CSA)* where users can install and share software. However, this leads to some difficulties in management and performance. All ins all, these experiments provided valuable lessons for Grid infrastructure supporters and Grid application users.

Another area of cross-Grid applications are related to the *GIN resource testing portal* that provides access to the various GIN VO resources and monitors their availability. Hence, this in particular demonstrates that the GIN VO VOMS service can be used by applications that use portals as the base technology. To provide an example, the portal is used to create and execute *Chemistry at HARvard Mechanics (CHARMM)* and *MadCity traffic simulation workflows* to the GIN resources, while the portal can also be used to monitor its execution. In particular, these workflow jobs are utilizing GRAM of Globus Toolkit 4 through the GEMLCA architecture within TeraGrid, OSG and the UK's NGS as well as broker-based submissions to EGEE sites. The portal technology is based on PGRADE and GridSphere.

7. Related Work

There are many pair-wise efforts in the field of interoperation and interoperability within e-Science communities and Grid projects. These efforts typically include the interoperation between one specific version of a technology with another one.

One of the major activities in reaching the interoperation between two production Grids are conducted within the EGEE-II project. One particular activity within this project focuses on the interoperation between specific version of gLite and UNICORE. This means an interoperation using the non Web services-based gLite middleware of EGEE with the non Web services-based production UNICORE 5 of DEISA. The fundamental goal of this activity is to work on job submission mechanisms to enable interoperation between EGEE (HEP and other scientific communities) and DEISA (HPC community) on a technical level. A closer look reveals that this interoperation is not based on open standards and is fundamentally based on VOMS and a UNICORE condor bridge developed by the NAREGI project in the past. Within GIN we envisage the usage of the OSGA-BES interface and HPC-P in production middleware very soon and thus it will be used in production Grids in the near future.

In the context of GIN-INFO, there were also already bilateral activities between the production Grids EGEE and OSG since autumn 2005. The interoperability is achieved by the usage of LDAP-based information systems and the GLUE schema in both Grid infrastructures. However, there is a different boot strapping. While OSG site URLs are generated from an OSG GOC DB, the EGEE site URLs are generated from an EGEE GOC DB. In addition, the EGEE and NDGF production Grids are working on interoperability since summer 2005 as well. Both Grid infrastructures use an LDAP-based information system. But while EGEE uses the GLUE schema, the NDGF uses the ARC information system that is based on a different schema. Both activities were involved within GIN-INFO and thus the initial idea of having translators and an official mapping process of attributes comes out of these efforts, but have been extended to much more production e-Science infrastructures within GIN-INFO.

8. Conclusions

This paper describes the results of the GIN-CG until mid of 2007, mainly focusing on the highlights of interoperation demonstrations used on the Supercomputing 2006 conference. These initial demonstrations provided massive insights in the process of enabling interoperation among production Grid and e-Science infrastructures in all the different areas of GIN.

Within GIN-INFO, all used information systems are rather similar in their scope and partly also in their functionality. Therefore, the usage and development of information providers to populate systems was easier than expected and can be recommended for production Grids today. This was in particular the case, because the query mechanism to extract the data were often based on LDAP, basic WS calls or open standards such as WS-RF. To sum up, query an information system and populate another was rather straightforward. Much more challenges were revealed in the context of the output of queries, whereby mainly all pieces of information conform to a dedicated schema. In general, the Grid community can life with different information systems but not with different content schemas (e.g. GLUE or CIM). If there is a use case that needs to be done across all Grid infrastructures, then the information we need for these use cases must be present and in agreement. Also motivated by the efforts undertaken within GIN, the newly formed GLUE-WG of the OGF will work towards a general content schema acceptable by the broader Grid community.

To sum up, to demonstrate the feasibility of the interoperation of different information systems and schemas, the GIN-INFO area demonstrated the Grid site on a world map use case. Hence, the GIN-INFO area has successfully shown that Grids can be interoperable with respect to information sharing and exchange, because the GIN Top-Level BDII contains information about all participating production Grid infrastructures.

The GIN-JOBS area demonstrated that production Grids are able of use upcoming standardized interfaces for cross-

Grid job submission and management. This was particularly done by using the HPC-P and implied emerging standards OGSA-BES and JSDL. Hence, it makes no sense to work on adapters or hacks between the different job submission technologies when there is a reasonable technology soon standardized. At the time of writing, the HPC-P and OGSA-BES specifications as well as the HPC extensions of JSDL have been through the public comment process of OGF and will soon become an official OGF proposed recommendation. This will lead to even more adoption by middleware vendors in the future and thus a wider use of these interfaces within production e-Science infrastructures in the near future. In fact, the cross-Grid operations from GIN-OPS revealed that such a standardized interface is a necessary requirement for production e-Science infrastructures in the future, especially when dealing with grand challenges (e.g. protein folding) in interoperation scenarios.

The efforts of the GIN-DATA group have shown that interoperation between different Grid and e-Science infrastructures can be established via GridFTP, SRB and various SRM implementations. However, it is not a straightforward task and very error prone today. Apart from improvements on information on the service configurations (e.g. supported versions, location of storage areas, etc.), the main issues for productive interoperation are in the areas of network tuning (e.g. advanced reserved dedicated network connections) and firewall management.

The GIN-AUTH group has shown that the base functionality of an identity-based authentication and authorization system is in place. There is a working federation of the current major Grid credential generators (the IGTF) and production Grid infrastructures have established the necessary infrastructure for the distribution and tending of the necessary information updates. These processes still tend to be centralized and rote as the workload is tedious and non-trivial.

Technically, the described interoperation components can be used today by e-Scientists within these infrastructures even if there is a negotiation phase needed to agree on resource usage with the production e-Science infrastructures. To conclude, the GIN efforts and the Supercomputing demonstrations provided a good first start towards world-wide interoperable e-Science infrastructures and emphasizes that common open standards as defined within OGF, DMTF, IETF, W3C or OASIS are absolutely needed.

Acknowledgments

The work presented here is conducted by many talented computer scientists and e-Scientists that are funded by various public funding organizations that we thank and acknowledge, because this work would not have been possible without their continuous and sustainable support.

References

[1] Australian Partnership for Advanced Computing (APAC). http://www.apac.edu.au/.
[2] BDII. https://twiki.cern.ch/twiki/bin/view/EGEE/BDII.
[3] DEISA. http://www.deisa.org/.
[4] EGEE. http://www.eu-egee.org/.
[5] GLUE-WG. https://forge.gridforum.org/sf/projects/glue-wg.
[6] Nordic DataGrid Facility. http://www.ndgf.org/.
[7] OGF - Grid Interoperation Now (GIN-CG). http://forge.ogf.org/sf/projects/gin.
[8] OGSA - Basic Execution Services (OGSA-BES). http://forge.ogf.org/sf/projects/ogsa-bes-wg.
[9] OGSA - High Performance Computing Profile (OGSA-HPC-P). http://forge.ogf.org/sf/projects/ogsa-hpcp-wg.
[10] Open LDAP. http://www.openldap.org/.
[11] Open Science Grid - OSG. http://www.opensciencegrid.org/.
[12] Pacific Rim Applications and Grid Middleware Assembly - PRAGMA. http://www.pragma-grid.net/.
[13] RFC2849. http://www.ietf.org/rfc/rfc2849.txt.
[14] TeraGrid. http://www.teragrid.org/.
[15] W. Allcock. *GridFTP: Protocol Extensions to FTP for the Grid*. Open Grid Forum Draft - No. 20.
[16] A. Anjomshoaa et al. *Job Submission Description Language (JSDL) - Specification Version 1.0*. Open Grid Forum Proposed Recommendation, 2006.
[17] Y. Demchenko et al. VO-based Dynamic Security Associations in Collaborative Grid Environments. In *Proc. of the Int. Symp. on Collaborative Technologies and Systems*, 2006.
[18] DMTF. *Common Information Model (CIM) Infrastructure Specification*. DMTF DSP004, 2005. version 2.3 Final.
[19] I. Foster. Globus Toolkit version 4: Software for Service-Oriented Science. In *Proceedings of IFIP International Conference on Network and Parallel Computing, LNCS 3779*, pages 213–223. Springer-Verlag, 2005.
[20] I. Foster et al. The Anatomy of the Grid - Enable Scalable Virtual Organizations. In *Grid Computing - Making the Global Infrastructure a Reality*, pages 171–198, 2003.
[21] E. Laure et al. Programming The Grid with gLite. In *Computational Methods in Science and Technology*, pages 33–46. Scientific Publishers OWN, 2006.
[22] A. Nadalin et al. *OASIS Web Security Username Token Profile 1.1*. 2006. OASIS Standard Specification.
[23] A. Rajasekar, M. Wan, R. Moore, and W. Schroeder. A Prototype Rule-based Distributed Data Management System. In *Procedings of the HPDC Workshop on Next Generation Distributed Data Mangagement*, May, 2006.
[24] J. M. Schopf et al. Monitoring and Discovery in a Web Services Framework: Functionality and Performance of the Globus Toolkit's MDS4. In *Argonne National Laboratory Technical Report ANL/MCS-P1248-0405*, 2004.
[25] A. Shoshani, A. Sim, and J. Gu. Storage Resource Managers - Essential Components on the Grid. In *Grid Resource Management*. Kluwer Academic Publishers, 2003.
[26] A. Streit et al. UNICORE - From Project Results to Production Grids. In L. Grandinetti, editor, *Grid Computing: The New Frontiers of High Performance Processing, Advances in Parallel Computing 14*, pages 357–376. Elsevier.