

The D-Grid Integration Project: How to Build a Grid for Different Communities

V. Gülzow¹, W. Gürich², K.-P. Mickel³, U. Schwiegelshohn⁴, K. Ullmann⁵, and G. von Voigt⁶

¹ Deutsches Elektronen-Synchrotron DESY-IT, Notkestraße 85, 22607 Hamburg, Germany

² Zentralinstitut für Angewandte Mathematik, Forschungszentrum Jülich GmbH, 52425 Jülich, Germany

³ Institute for Scientific Computing, Forschungszentrum Karlsruhe GmbH, Hermann-von-Helmholtz-Platz 1, 76344 Eggenstein-Leopoldshafen, Germany

⁴ Robotics Research Institute, Universität Dortmund, Otto-Hahn-Straße 8, 44221 Dortmund, Germany

⁵ DFN-Verein, Stresemannstraße 78, 10963 Berlin, Germany

⁶ Regionales Rechenzentrum für Niedersachsen, Leibniz Universität Hannover, Schloßwender Straße 5, 30159 Hannover, Germany

email: volker.guelzow@desy.de, w.guerich@fz-juelich.de,
klaus-peter.mickel@iwr.fzk.de, uwe.schwiegelshohn@udo.edu,
ullmann@dfn.de, vonvoigt@rrzn.uni-hannover.de

Abstract

The Grid world is highly dynamic due to the appearance of new resources, easier access and new Grid applications. Moreover, an increasing number of scientific communities begins to realize that Grids will become an indispensable tool for modern science. However, the requirements of those communities often differ significantly from each other. Therefore, it is a particularly difficult challenge to establish a national Grid infrastructure that serves as many communities as possible. To address this challenge, the German D-Grid has developed a concept that is based on development of independent Community Grids that are connected with a single integration project. Further, the process will involve several steps to consider dependencies between the various projects and their tasks. Presently, we are still at the beginning of the realization of a German Grid. Therefore in this paper, we first briefly explain this concept in this paper. Then we focus on the integration project and describe its tasks in more detail. This includes some first results and experiences that were achieved within the first year of the project. Finally, some plans for future steps are presented.

1 The D-Grid Concept

In 2005, the German Federal Ministry of Education and Research started a series of projects to provide many application sciences and enterprises with

modern Grid technology. The first phase began with several application projects from different scientific areas where Grids are either already well accepted tools, like in high energy physics, or in which Grid use is ready to start due as the members of this area expect benefits from the Grid, like in earth sciences. In the ideal case, such an application project - also called Community Grid - comprises most of the major German research institutions belonging to this scientific area together with some computer scientists who have experiences in grid technology. The balance between application and computer science is the crucial point in the composition of a Community Grid. If there are too few application researchers then the computer scientists may develop a well structured Grid that does not consider the requirements of the actual users and will never be accepted by the community. On the other hand, application researchers tend to develop quick and dirty solutions that solve their problems in the short run but often create major obstacles to future development and extensions of the Grid. It is therefore the task of the computer scientists in such a Community Grid to develop well structured and modular Grid tools and services that are needed by this community but are not yet generally available, that is, the Grid projects are application driven. But although upgrades and extensions of the Grid is now possible there is still the potential danger of producing several isolated Grids which do not share tools and resources. Therefore, an additional integration project was established to provide a common foundation for the German Grid initiative, also called D-Grid. In this paper, we discuss the internal structure and the tasks of this project.

2 The D-Grid Integration Project

The integration project is supposed to work closely together with the computer scientists of the Community Grids. Structurally, it is divided into four sections:

1. Basic software components
2. Core D-Grid and studies
3. Network and security
4. Dissemination and sustainability

The basic software component section has the task to provide a common software base for the Community Grids. As the integration project and the Community Grids had the same starting date and some Community Grids had no prior exposition to Grid technology not all requirements of the Community Grids were already available at this starting date. Therefore, the integration project made the decision to focus on three commonly used middlewares and some other basic components to start with. Higher level Grid services were postponed to the second phase of the integration project, see Section 5, when the requirements of the Community Grids are clearly defined.

The selected middleware systems are the Globus Toolkit Version 4.x (GT 4) [1], gLite [2] and UNICORE [3]. GT 4 is the de facto world wide middleware standard and runs on several operating systems. It supports the Web

Service Resource Framework (WSRF) [4] and mainly provides basic services. gLite originates from the LCG production Grid [5] mainly used by the high energy physics community. It comprises several higher and lower level services but is currently restricted to a specific Linux operating system and an older version of the Globus Toolkit (GT 2.4). There are efforts under way to overcome those restrictions. UNICORE is a vertically integrated complete middleware solution reaching from a user portal to basic services. However, only a limited number of users employ it as a production system at the moment. As all those systems have some advantages and some problems, it is not likely that all Community Grids will use the same middleware. Therefore, the integration project considers all those described middlewares in the set of basic software components. Based on community demand it will be decided at the end of the first phase whether some of those middlewares are removed from the software base and/or some new (commercial) middlewares are added.

The basic software also includes GridSphere [6, 7], a user portal solution, the Grid Application Toolbox [7], to allow transparent selection of appropriate Grid tools, and the data management tools SRM/dCache [8] which is particularly used by the high energy physics community, OGSA/DAI [9, 10] as a means for inclusion of data bases into the Grid, the Storage Resource Broker [11], and datafinder [12]. Finally, a group particularly addresses tools for the management of virtual organizations.

The participants of the basic software section started to organize training sessions and tutorials for the members of the Community Grids in order to provide them with enough information to enable an educated choice of a suitable basic software for their Grids. Those events were so well received such that some of them were even repeated due to high demand. In addition, general support in employing those systems is also provided to the communities. Further, Community Grids can request special support for specific problems related to the basic software. For instance, the climate research community asked for Globus support for NQS II. However, those requests have been rare during the first year of D-Grid. Therefore, members of the integration project are now actively inquiring within the Community Grids about problems with the basic software using means like questionnaires and presentations at workshops of the Community Grids. In addition, the integration project provides documentation and installation packages in addition to those that are already generally available. This particularly includes the porting of the basic software to new platforms. Also small interoperability gaps are filled on request on the Community Grids, while developing new middleware components has been explicitly excluded from the first phase of the integration project in order to emphasize the service character of this project. Larger gaps requiring development efforts are addressed in separate new D-Grid projects, see Section 4. Finally, Community Grids are expected to eventually transfer some tools that have been developed within their projects to the integration project if those tools may be beneficial for other communities as well. For instance, the astrophysics community has provided tools for easier handling of virtual organizations. Those tools are the first example of such a

component transfer and will be included into the reference configuration of the new D-Grid infrastructure, see Section 3.

In addition to basic software components, the integration project also establishes the so called core D-Grid. This is one of the tasks of the second section whose participants are various institutions with their own infrastructure who have accepted to deploy the basic software components and additional tools from communities in order to test interoperability and to determine potential problems in production use. Compute centers of large federal research institutions like the research centers of Jülich and Karlsruhe belong to this section as well as large regional compute centers like those of Hannover and Berlin and smaller university compute centers. The composition of the partner structure assures that concerns and restrictions of different resource providers are considered.

This section also handles issues that are generally agreed to be of vital importance for Grids but for which no commonly accepted tools exist yet. Study projects are defined to address those issues which include monitoring as well as accounting and billing. Those studies are executed in close collaboration with interested communities. The monitoring, accounting and billing activities provide the components and concepts for visualizing the state and for accounting the usage of Grid resources. These functionalities are necessary to satisfy sustainability and long-term operation of the Grid infrastructure in Germany. Particularly, accounting and billing of the resource usage is a requirement of any economically orientated production Grid system. This also includes dynamic pricing that enables economic load balancing of the resources and thus improves the efficiency of the Grid. To this end actual prices for resource usage are calculated on the basis of the present load average and usage history that is provided by usage records. This allows resource providers to use different pricing and market models.

On the other hand, efficient monitoring systems play a substantial role within the setup and administration of a Grid environment, enabling the dynamic and continuous monitoring of the availability of system components and the status of submitted jobs. The presentation of the current state of resources is necessary to enable dynamic service discovery of Grid middlewares and to facilitate error tracing to reduce the degradation or the loss of services. Most users are strongly interested in the status of their submitted jobs to analyze the state of job processing, to trace errors and last but not least to satisfy their own curiosity. The presentation of the monitoring information by a customer service management (CSM) provides the basis for the application of accounting procedures.

Consolidating the information gathered from monitoring systems, an accounting system generates usage records based on single user accounts, which can be grouped together and thus also indicate the resource usage of virtual organizations. These usage records can further be utilized to analyse the acceptance of resources for future deployment, investments, and maintenance. The provision of various middlewares and their future interoperation requires a comprehensive concept harmonizing the heterogeneous accounting systems.

Billing of Grid services within D-Grid depends on the evaluation of usage data. This includes accounting and prognosticating of those data with respect to hardware-, software-, distribution- and transmission costs for the benefit of both Grid resource providers and consumers. Accounting modalities for the use of services and resources are also part of service level agreements (SLAs).

The third section deals with network and firewall issues as well as security aspects in D-Grid and is coordinated by the Deutsche Forschungsnetz Verein (DFN). The network infrastructure in D-Grid is based on DFN's network X-WiN. According to the demands of the communities extensions to the existing network infrastructure are planned and developed. Based on the demands of some communities several specific links have been set up for D-Grid sites. This includes, for instance, an international 10 Gbit/s connection via X-WiN and Geant from the Forschungszentrum Karlsruhe to CERN on request of the high energy physics community. In parallel, a sequence of interviews with all D-Grid communities is conducted to gather other specific requirements. Up to now all identified demands can be fulfilled by the available X-WiN VPN options.

Moreover within the integration project, studies are executed in the area of firewall technologies for high speed access and specialized high speed oriented transport protocols which may become very important for some communities in the future. Security services in Grid environments are largely based on PKI and X.509 certificates. The focus is on authentication services but also includes parts of user authorization. The authentication mechanisms are well established with DFN and Forschungszentrum Karlsruhe being accredited certification authorities at the EUGridPMA. Those institutions have already set up a great number of Grid registration authorities in organizations participating in D-Grid thus providing an easy process of issuing Grid certificates to D-Grid participants. Unfortunately, authorisation in Grids is not as straight forward to deal with. A detailed study has been conducted to determine and analyze use cases for authorization in the D-Grid. It shows that only inhomogeneous authorization mechanisms are currently available in the supported Grid middlewares. On the basis of community interviews, a concept for an overall authorisation structure in D-Grid is developed.

Tasks and workflows of Computer Emergency Response Teams (CERTs) are already well established in general networks but Grids have created new challenges. Therefore, the existing CERT structures are extended to cover special Grid demands. There the focus is on the four areas coordination and cooperation, incident prevention, detection, and reaction. Some Grid software has already been examined with respect to its vulnerability, and the demands of the D-Grid communities are considered in the development of Grid CERT services which includes the setting up of a hotline and the design of special workshops and tutorials.

Finally, it is the responsibility of the last section to assure dissemination and sustainability. This particularly includes information of all communities potentially interested in Grid use about the results of D-Grid and particularly those of the integration project. Further, models are developed to address eco-

conomic aspects of Grids and Grid use across national boundaries. The integration project will also present D-Grid in European discussions about the creation of a European Grid Infrastructure (EGI).

3 The D-Grid Infrastructure

In the end of 2006, the German Ministry of Education and Research provided additional funds for establishing a basic Grid infrastructure in Germany as previous funds were entirely dedicated to human resources to support configuration and development. Therefore, all needed infrastructure was provided by the participating institutions on their own. However, all institutions hosting components of the new infrastructure had to agree to certain conditions regarding system software and use of the resources:

1. Each compute cluster must support all three middlewares (GT 4, gLite, UNICORE) of the integration project without requiring adaptation to other operating systems.
2. Each storage resource must be accessible by one of the tools (SRB, SRM/dcache, SRB) of the integration project.
3. All resources are open to all members of D-Grid.

Although all participating institutions are free to satisfy those conditions using their own approaches, the integration project also provides a reference configuration which will be presented to the institutions in January 2007. The preparations for the definition of this reference installation have already shown the benefit of the structure of the integration project. The reference installation will demonstrate transparency of the three middlewares supported by D-Grid. It will be possible to submit three applications via three different middlewares and let them run together on a single cluster. This is the first step towards efficient resource sharing among different communities which contribute their own resources to the Grid.

4 New D-Grid Projects

In the first phase, the focus of D-Grid was on scientific communities on top of the integration project. In the second phase, those projects will be complemented by three new types of projects.

1. New communities with a focus on commercial applications. Those communities include industrial partners and are also based on the software foundation of the integration project.
2. New service structures that offer services which have been determined as useful in the first phase of D-Grid. Those services are also based on the integration project but not tightly coupled with it. Therefore, they can be established as an independent project.
3. Projects addressing gaps that have been identified in the software of the integration project during the first phase and that are too large to be

handled in this phase, see Section 2. Those gap projects are tightly coupled with the content of the integration project.

The integration project will provide an organisational structure for the gap projects although they are formally independent. The first project of this kind deals with the interoperability of Shibboleth [13] and the Virtual Organization Membership Service (VOMS) which is part of gLite. This project has been jointly identified by the integration project and the Community Grids and includes partners from several projects. Further, there will be a monitoring project to complement the monitoring study and to provide Grid-wide visibility of the new Grid infrastructure. Moreover, the core D-Grid will be extended by two projects particularly dealing with the inclusion of smaller compute centers. In addition, there is the plan to establish a project dealing with service level agreements (SLAs). This project must be based on the foundation provided by the integration project and at the same time support commercial and scientific applications.

5 Continuation of the Integration Project

The integration project will complete its first phase in 2007. It is planned to add a second phase with a potentially modified focus. This focus will mainly depend on the requests of the communities. Particularly, the new Community Grids with a commercial focus, see Section 4 were not considered in the original concept of the first phase. Moreover, the well established scientific Community Grids may determine that some services of the integration project are no longer required while some new services must be established. A candidate for such a new service may be a resource management and brokering service with the goal to generate an interoperable service including existing commercial and non-commercial tools as well as systems already developed or selected by the Community Grids. Such a project must consider generally available tools as well as developments within existing communities which have already addressed this subject with respect to their specific requirements. Accounting and billing may be another candidate that will take into account the studies already performed within the integration project.

Acknowledgements

The authors are grateful to German Federal Ministry of Education and Research which funds the integration projects. There are many participants who are contributing to the integration project. The authors particularly acknowledge that without continuous efforts of those project members, a success of the integration project is not possible.

References

1. I. Foster. Globus Toolkit Version 4: Software for Service-Oriented Systems. IFIP International Conference on Network and Parallel Computing, Springer-Verlag LNCS 3779, pp. 2-13, 2006.
2. R. Berlich, M. Kunze, and K. Schwarz. Grid Computing in Europe: From Research to Deployment. Proceedings of the 2005 Australasian workshop on Grid computing and e-research, Newcastle, New South Wales, Australia, Vol.44, pp. 21 - 27, 2005
3. A. Streit, D. Erwin, T. Lippert, D. Mallmann, R. Menday, M. Rambadt, M. Riedel, M. Romberg, B. Schuller, and P. Wieder. UNICORE - From Project Results to Production Grids. Grid Computing: New Frontiers of High Performance Computing, L. Grandinetti ed., Elsevier, pp. 357 - 376, 2005
4. Web Services Resource Framework (WSRF) v1.2 Specification, OASIS, March 2006, <http://www.oasis-open.org/committees/download.php/16654/wsrf-cs-01.zip>
5. J. Knobloch and L. Robertson. LHC Computing Grid. The LCG TDR Editorial Board, December 2006, http://lcg.web.cern.ch/LCG/tdr/LCG_TDR_v1.04.pdf
6. J. Novotny, M. Russell, and O. Wehrens. GridSphere: A Portal Framework for Building Collaborations. Concurrency & Computation-Practice & Experience. Vol.16(5), pp. 503-513, 2004
7. G. Allen, K. Davis, K.N. Dolkas, N.D. Doulamis, T. Goodale, T. Kielmann, A. Merzky, J. Nabrzyski, J. Pukacki, T. Radke, M. Russell, E. Seidel, J. Shalf, and I. Taylor. Enabling Applications on the Grid: A GridLab Overview. International Journal of High Performance Computing Applications, Vol.17(4), pp. 449-466, 2003
8. T. Perelmutov, D. Petravick, E. Corso, L. Magnoni, J. Gu, O. Barring, J.-P. Baud, F. Donno, M. Litmaath, S. De Witt, J. Jensen, M. Haddox-Schatz, B. Hess, A. Kowalski, and C. Watson. The Storage Resource Manager Interface Specification. December 2006, <http://sdm.lbl.gov/srm-wg/>
9. B. Plale, D. Gannon, J. Alameda, B. Wilhelmson, S. Hampton, A. Rossi, K. Droegemeier. Active Management of Scientific Data. IEEE Internet Computing Magazine, Vol.9, pp. 27-34, 2005
10. K. Karasavvas, M. Antonioletti, M. Atkinson, A. Hume, M. Jackson, A. Krause, and C. Palansuriya. Introduction to OGSA-DAI Services. Lecture Notes in Computer Science, Springer-Verlag LNCS 3458, pp. 1-12, 2005
11. C. Baru, R. Moore, A. Rajasekar, and M. Wan. The SDSC Storage Resource Broker. Proceedings of CASCON'98, Toronto, Canada, 1998
12. E. Hoffmann. Datafinder - Architektur eines Managementsystems für wissenschaftliche Daten. December 2006, <http://www.dlr.de/as/desktopdefault.aspx/tabid-379/1626>
13. R.L. Morgan, S. Cantor, S. Carmody, W. Hoehn, and K. Klingenstein. Federated Security: The Shibboleth Approach. Educause Quarterly, Vol.27(4), pp. 12-17, 2004