

SuGI - Portal and Training Systems for Grid Middlewares

Joachim Götze, Bernd Reuther, Paul Müller
Dept. of Computer Science, ICSY Lab
University of Kaiserslautern
Kaiserslautern, Germany
{j_goetze, reuther, pmueller}@cs.uni-kl.de

Marc Seifert, Sebastian Breuers, Viktor Achter, Ulrich Lang
Regional Computing Center of Cologne
University of Cologne
Cologne, Germany
{marc.seifert, s.breuers, vachter, lang}@uni-koeln.de

Abstract

Grid computing is a research and development topic, currently getting much attention. As a result, new ideas as well as research and project results are being produced at a very high rate. Therefore, it is difficult for resource providers and researchers in this area to keep track of the development. Resource providers and researchers, new to the topic of Grid computing, have to overcome a huge inhibition threshold before work in this area can be started. The SuGI project is aiming to support providers and researchers with two approaches. First, the SuGI portal presents videos, documents, and links regarding many topics in the area of Grid computing for newcomers as well as for experts. Second, the training systems for Grid middlewares created by the SuGI project can help beginners to do their first steps without a complex installation process, while experts can easily try out different configurations without interfering with their productive systems.

1. Motivation

SuGI (Sustainable Grid Infrastructures) is a project of the German Grid Initiative (D-Grid)¹, which is a joint initiative of research and industry partners started in September 2005. Its goal is the development of a distributed, integrated resource platform for high-performance computing and related services to enable the processing of large amounts of scientific data and information. The major task of the SuGI project is to disseminate the knowledge of Grid technology and to enhance its use. Therefore, SuGI addresses all academic computing centers as well as enterprises, which still have not adopted Grid technology. They will be supported in providing Grid resources and services. During the project, research and technology experiences gained in the D-Grid projects will be made available to the intended audience ranging from experts of Grid computing over organizers of Grid workshops to users of Grid technology.

SuGI uses two approaches to introduce new users to Grid computing: a portal that gives theoretical background and

training systems that provide practice with real Grid middlewares. Thus, SuGI offers training courses, makes video and audio recordings of external courses available, provides training systems to allow administrators and users to try out basic configuration steps within a Grid middleware. All this is made available to the Grid community via a scaling training infrastructure.

In co-operation with several partners, e.g. DESY (Deutscher Elektronen-Synchrotron), e-learning and information components for Grid computing and Grid middlewares have been developed. The training systems contain basic configurations of the Grid middlewares Globus Toolkit and UNICORE, so that users can learn either how a basic installation of such a Grid middleware can be done or experience and test an already functioning middleware, which even can be modified to provide a full-featured Grid user interface. Information and training materials are provided online via the SuGI portal².

The focus is set on well known approaches of e-learning in order to provide high quality educational material for training and further education of staff members and users at computing centers. Heterogeneous communities, as well as the sustainability of processes and the created products are playing a major role in this environment. This includes aspects like the different previous knowledge of users and Grid experts, the different goals of computing centers of industry and research centers as well as the resource conserving provisioning and archiving of content and research experiences from the Grid environment.

This paper is organized as follows. In section 2, the functionality and the development process of the SuGI portal is described, while section 3 discusses the technology used and the possibility provided by the Grid training systems. The paper concludes with a summary and future work that is to be done in the SuGI project.

2. Portal

Grid computing is a topic of computer science currently receiving a lot of attention by the research community.

1. The German Grid Initiative (D-Grid) is funded by the German Federal Ministry of Education and Research (BMBF)

2. The SuGI portal can be found at <http://sugi.d-grid.de>

Hence, the knowledge about Grid computing and the amount of available tools and implementations in this area are growing rapidly during the last years. Neither scientists nor resource providers are able to stay up to date on all relevant topics. The SuGI portal aims at the provisioning of a single point of information on all relevant topics of Grid computing in order to enable resource providers and Grid users to either introduce themselves with the current state of the development or stay informed on topics of Grid computing important for their current research.

2.1. From an e-learning platform to an enhanced multimedia platform

The development model described by Gaiser & Werner [2] is the basis for developing, implementing, and evaluating the SuGI portal. This model focuses on a process consisting of a development of a strategy, and formulation of a catalog of specifications, including the identification of target groups which may be placed out of or across institutional borders. Further steps are a technical and a visual design of the portal, implementation and different steps of quality management as well as an introduction and evolution. All steps are iteratively repeated for each generation of the portal. That leads to a development process, created as a generation model where the results of the evaluation will pour in the development of any new generation which leads to a higher impact of evaluation results and enables a well-founded monitoring in the process of building up the portal

Target groups for the SuGI portal are experts of the D-Grid communities, organizers of training events, multipliers in computing centers (commercial and non-commercial) as well as users. All members of these highly heterogeneous target groups need to be networked with each other or with the respective communities.

These special requirements driven by the D-Grid as well as the process of development and implementation of the e-learning portal has led to a number of innovations. Thus, the SuGI portal differs from conventional e-learning solutions in some criteria.

Most of the usual e-learning solutions like ILIAS (e-learning portal of the University of Cologne), Prodo (University of Applied Science Cologne), E-Campus (JL University Gießen), etc., reflect the learning context of educational systems or universities. They support the lecture in a framework of an educational system established over decades which is characterized by well-defined structures like e.g. teaching classes with corresponding readings, detailed tutorials, etc. Support of such e-learning activities often is guaranteed by permanently funded members of staff. Directed to this situation, these e-learning solutions provide personalized member areas enabling learners to find those pieces of information, to manage their time schedule, and to check their learning success to only mention a few important

functions. That leads to an effectively supported form of learning which is, however, related to high mentoring and support effort.

Within the D-Grid, however, members are only funded for a restricted time. During the course of the project, a large amount of information and knowledge will be generated, processed, and published. This information needs to be handed down to succeeding scientists. Thus, the SuGI portal rather forms a mixture of conventional e-learning solutions and free video sharing platforms like YouTube. Such systems are driven by a community in terms of content. Any user is able to provide content to others and the contents are likewise monitored to a large extent by the community.

Access to the information provided is explicitly directed to a large amount of anonymous users. Thus, the SuGI portal maps the requirements of the D-Grid to a large extent. While video sharing platforms like YouTube are restricted to videos which need to match a certain format, such functionality is not adequate to match the requirements of an e-learning portal. Therefore, the SuGI project is pursuing an appropriated synergy of both described solutions, i.e. e-learning and video sharing platforms. Important features are – with respect to recent learning theories as well as to our target groups – the networking of contents or part of contents within the SuGI portal and with further e-learning material or information providing web sites like the e-learning portal of the SUN Microsystems Academic Initiative (SAI), the GridKa School Karlsruhe, or the International Summer School on Grid Computing (ISSCG). Contents are linked for example by authors, keywords, Grid projects, related in-class events, etc. Thus, concrete learning situations like talks at in-class events will be made accessible for revision and reevaluating by recording and providing them online. The knowledge gained in these situations is contextualized with contents of adjacent as well as overlapping content areas and media types and may thus be transferred associatively into new learning situations.

A result of those efforts is a portal solution enabling users to administer contents of manifold formats like e.g. texts, hyperlinks, videos in different file formats, learning modules with various levels of interactivity, created by different tools and editors, training systems, etc. At the same time, all provided information is easily accessible. As from generation 2 (the current generation is 1.5), supervision, support, and mentoring will be fully community-driven. Thus, invested work coming along with the arrangement of workshops, conferences, and summer schools, etc., is integrated directly in e-learning contents, that will be provided sustainable to an extended community during a longer time span. There is a significantly better protection of investments.

Employees of computing centers who are unable to visit lots of in-class lectures and conferences will be able to consume contributions they are interested in, suited to their working time schedule and to their personal work pace.

Participants of in-class events may re-experience the contributions on request at their personal computer. Thus, the SuGI portal forms an innovative, sustainable and scaling e-learning solution.

2.2. Ways to access the portal

One of the essential tasks of an e-learning portal is an easy and scalable access to the provided material. In the SuGI portal, this is achieved on the one hand by a target group based and/or category based access and on the other hand by the search functionality. The menu items of the portal represent categories of content, e.g. the item *events* provides an overview of all recorded events. This way of access is suitable for users already having a detailed idea of the information they are searching for. An image map on the title page supplying entrance points focusing on certain target groups and their expected desire for information. This way of access is suitable for portal users looking around or prospecting for a first overview on a certain topic. The search functionality is focusing on users looking for suggestions on a certain topic. Therefore, the search results are represented in a way that the content can be acquired quickly, so that users gain the greatest possible benefit.

Typically, search engines provide a list of references of web sites or documents allowing the user only a minimal assessment of the referenced content. The contents of the SuGI portal are augmented by metadata used to provide a strong relation between various aspects of the content, e.g. a certain document is enhanced with information about authors, keywords, abstract, and required experience level, but allows also a relation with the corresponding event the document has been presented at. These metadata is used to enable a better assessment of the search results.

Figure ?? shows how a single content element is presented in the SuGI portal in a compact way. This form of presentation is not only used for search results, but for every kind of listed information. The reference *details* enables the user to retrieve an enhanced overview for the selected element. A functionality that allows to show relevant sections of multimedia content is planned for the future. Hits referring to the same content element are identified and summarized. Search results are rated by the source of discovery, e.g. keyword, title, description, and ordered correspondingly. Furthermore, several types of metadata, e.g. authors, keywords, events, are linked, so that it is very easy to identify additional content of the selected criteria.

Additionally, a filter option has been included, enabling the user has the ability to further refine the currently viewed information. The filter can restrict a list view, e.g. by language, document type, and the level of difficulty. Furthermore, the list view can be sorted by the publication date of the entries.

When the user has found a suitable content element, the concept of document libraries can help to find further information regarding the selected topic. Document libraries are libraries of videos, documents and links concerning a special topic. A library is enhanced with a short introduction to the topic and information about its purpose, e.g. the dCache document library contains basic information about dCache [3] and a training system for first steps in dCache, but also additional detailed information about more sophisticated topics regarding dCache.

The development of the SuGI portal has been partitioned into three steps, which are based on an evolutionary concept enhancing the portal content and quality with every step. These steps are called generations:

- Generation 0 has been created to provide a foundation for the SuGI portal allowing the fast publication of content and material created so far. Furthermore, the early provisioning ensured the identification and solution of integration problems. The so gained knowledge is integrated in the following generations in order to ensure a steady improvement.
- Generation 1 included the expected functionality to a nearly complete extend and introduced multiple planed measures to enhance the user satisfaction, e.g. an extended full text search function. Additionally, templates have been created for the easy integration and representation of records and modules.
- Generation 2 will be created implying the experience of user feedback from previous generations and containing refined structures and functionality. Furthermore, sustainability is playing an important part, so that functionality will be introduced allowing the provisioning of new content by specially selected user groups.

3. Grid Training Systems

Grid middlewares provide a wide range of interacting services and therefore the installation and testing of the software is complex and error prone, especially for administrators new to the topic of Grid computing. In contrast to the larger computing centers, where Grid technology is getting more and more a common good, smaller computing centers are frequently lacking the men power to acquire a thorough understanding of new approaches towards computing. In order to enable administrators to introduce themselves with Grid technology without the need to overcome the complex problems of creating a testing installation, the Grid training systems of the SuGI project provide an easy to use access to common Grid middlewares:

- The practical and every day usage of a complex software system provides certain insights that cannot be compensated by theoretical descriptions in documents and presentations. The practical use of Grid software

includes an exceptionally huge effort especially for inexperienced Grid user. One of the problems to run a Grid middleware is the execution of configuration steps needing specific expertise. The possibility to test a Grid middleware with only little effort allows for example the evaluation of Grid middleware.

- Administrators and users can use the training systems based on the sandbox principle and try out various individual operations without any risk to the current productive system in operation. Administrators can therefore test configurations and software tools, while users can use the services of the training systems and need not run their tests against real Grid resources, e.g. to verify complex job descriptions without spending expensive compute time.
- Based on these training systems, administrators can easily provide tools for users to access Grid resources (Grid UI). Therefore, a training system only has to be extended by the user's certificate and allows users the access to complex client applications without the need for local installations or without the need to provide a central access to the applications.

The Grid training systems created by the SuGI project are using a hardware virtualization like VMware or Virtualbox. An advantage of such a solution based on virtual machines is the quick restoration of the initial Grid middleware configuration, because either a snapshot of a working configuration can be restored or the whole virtual machine can be erased and the user returns to the originally provided virtual machine. Such a virtual machine emulates a single processor system with 1 GB RAM, 8 GB harddisk and has network access via network address translation (NAT). The operating system is a standard OpenSuSE 11.1 installation, without many large end user applications. Based on this configuration, multiple trainings systems are provided containing various pre-configured Grid applications. Currently the Grid middlewares Globus Toolkit and UNICORE are supported. These Grid training systems contain the typical services used with the corresponding middleware (see 3.1 and 3.2). Both supported middlewares are currently undergoing major changes, i.e. UNICORE has been updated to a new major release (version 6) including multiple complex changes within the middleware architecture and Globus Toolkit has been updated to version 4.2 resulting in binary incompatibilities between services compiled with the new version and services compiled with the previous versions. As a result, the SuGI project aims to create training systems for all four middleware versions, so that administrators and users have the option to choose the training system most suited to their desires.

The training systems use special self-created certificates, which are only valid within these systems, and therefore allow the user to introduce themselves with aspects of

certificate management and of configuration issues regarding certificates. Using these training systems as a Grid User Interface (Grid UI) requires the substitution of the self-created certificates with valid user certificates signed by the corresponding Grid Certificate Authority (CA). Therefore, the training systems supports the import of user certificates from an external remote machine, e.g. the local host of the virtual machine.

Another aspect of the configuration of a Grid middleware is the integration of compute cluster into the Grid environment. In order to allow the administrator to learn and test how to configure the bridge between a Grid middleware and a compute cluster, the SuGI project has created a virtual cluster. The virtual cluster contains several virtual machines providing the functionality either of a cluster frontend (or head node) or of a cluster compute node (or worker node). The cluster is based on Rocks, a Linux based operating system, which is a rather common approach used by resource providers. The cluster is managed by the Torque resource manager, a PBS-based application, and the Maui scheduler. These virtual machines combined with the training systems for Grid middlewares enable administrators to try out the integration of a compute cluster into a Grid frontend. Additionally, administrators can experiment with different cluster configurations, e.g. regarding to queue management, without interfering with the currently used productive systems.

3.1. Globus Toolkit 4 Training System

The training systems provide a realistic view on a basic Globus Toolkit 4 (GT4) installation (see figure 1). Beside installing and configuring the Grid middleware itself, an administration must integrate the middleware in a specific environment. The later includes access to Grid resources, e.g. compute nodes, access to storage, participation in a monitoring infrastructure and mapping of user certificates to unix accounts using a Grid map file.

The GT4 training system is currently available in three variants for GT4.0.8³:

- The first variant aims to *train the configuration* of GT4. It consists of all necessary software to install a typical GT4 middleware node. Administrators can use this system to learn how to configure GT4 elements or test specific configurations before using them in a productive environment.
- A virtual machine with a *running middleware* consists of several configured and running GT4 services. Administrators can use this system to learn how to use and test the Globus middleware. It is possible to launch compute jobs locally by using the fork job manager instead of a batch system.

3. Similar systems will be available for GT 4.2 soon.

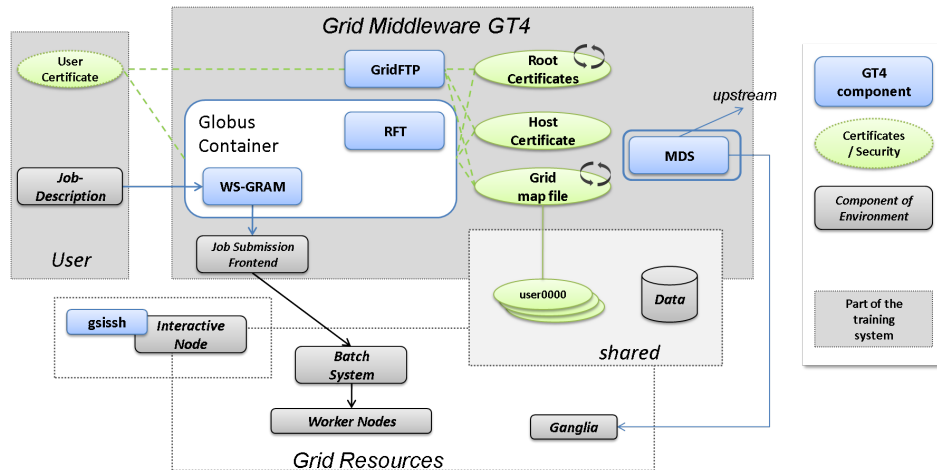


Figure 1. Administration view on GT4

- The *Grid UI* variant provides the functionality of a Grid user interface, including tools to submit jobs and transfer data. The Grid UI virtual machine can be adapted easily to be used as a frontend for a productive GT4 system.

GT4 is not designed to be used directly by an end user, but a toolkit for Grid application developers. GT4 offers several services and tools, where each application may choose to use different subsets of services and tools. Nonetheless some basic services and tools will be used by most GT4 applications, which are supported by our training systems:

- The Globus container, providing a runtime environment for several Globus services. The container also implements the Globus security mechanisms. The container activates about 50 services by default, the most important are: WS-GRAM, RFT and MDS.
- WS-GRAM is the Globus Resource Allocation Management with a Web Service interface. The WS-GRAM offers functionality for job submission and control, but is not a job scheduler.
- GridFTP offers high-performance file transfers. It is possible but not necessary to run a GridFTP service on the same host than the Globus middleware.
- The Relibale File Transfer (RFT) service controls and monitors file transfers using GridFTP servers, providing reliability and resume functionality for file transfers.
- The Monitoring and Discovery System (MDS) offers information about the availability and status of Grid resources.

Within the training system administrators can use and (re-)configure these services, e.g. to test their integration into a specific environment. Figure 1 illustrates several components that must be taken into account by a Globus administrator. The GT4 training system provide the user and Grid middleware components.

3.2. UNICORE 5 Training System

The goal of the UNICORE 5 training systems⁴ (see figure 2) are similar to the GT4 systems. The user interface of UNICORE is designed to be used directly by end-users. Thus the user interface of UNICORE is much simpler to use and to install. The focus of the UNICORE training system is to support administrators. There are two variants of the UNICORE 5 training system:

- A training systems with a *running middleware* is provided to test configurations and test the integration of UNICORE 5 in a specific environment.
- A virtual machine containing a UNICORE 5 *user interface* is provided for completeness only.

With regard to job submission the basic functionality of UNICORE 5 is similar to GT4, but concepts and services provided are different:

- A Gateway controls the access to a UNICORE site and handles the authentication of users.
- Within a UNICORE site several Grid resources are represented as a virtual site (Vsite). A Network Job Supervisor (NJS) controls and monitors jobs per Vsite.
- The Incarnation Database (IDB) adapts a job description to a specific Grid resource, while the Target System Interface (TSI) is an adapter for accessing the resource.

To adapt a UNICORE 5 system to a specific environment, administrators at least must define and configure a Vsite, must enable data transfer to and from the grid resource (this is supported by UNICORE), and must define a mapping of user certificates to unix accounts by the UNICORE User Database (UADB).

4. Unicoire 6 will be supported soon.

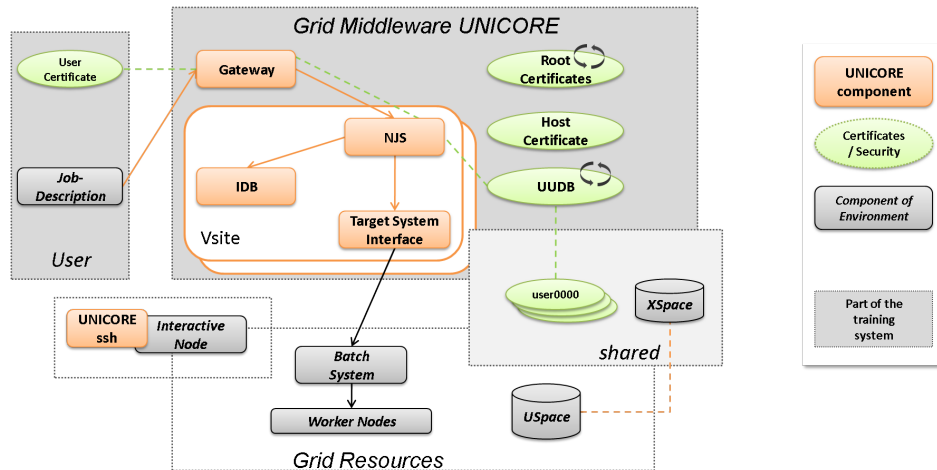


Figure 2. Administration view on UNICORE 5

4. Conclusion and future work

The main goal of the SuGI project is the support of sustainable Grid-Infrastructure in Germany. This should be achieved by reducing the obstacles of implementing and using Grid infrastructures. Thus the method of SuGI is the dissemination of knowledge about Grid technology. Therefore SuGI provides theoretical background via an information portal and training system to easily gain practical experience. The portals offers users several alternatives to find appropriate information, via a structured menu in combination with filters, a sophisticated search function, and linked keywords. Because of the user-friendliness of the portal, new users of Grid computing can easily find and retrieve information regarding their fields of interest. Content providers can describe their content in details and may structure the content, e.g. using containers to group related document. The training systems enables administrators and end users to use and test real grid middleware without the overhead of installing and configuring complex software.

In addition the SuGI Portal itself must be sustainable, this is achieved by enabling portals users to provide new content on their own. Therefore the portal is designed for:

- easy integration of new contributions, so that up-to-date content is ensured,
- flexible content handling, so that many different types of files and media are supported, ranging from single PDF files to structured video presentations
- a well-structured user interface, allowing new users find their ways quickly,
- using open source software and standard tools, ensuring stability and simplifying maintenance.

Currently the Portal contains 190 contributions, most are video presentations, many consist of more than one document, and more are in preparation. The generation 2

of the portal will be released soon. Then it will be possible for users to provide content on their own. Also the final evaluation of the Portal will be started with the release of the second generation. First feedback of grid experts has shown that the long term availability of their presentations and documents are highly appreciated and has resulted in cooperation leading to further contributions to the portal.

References

- [1] D. Clark, et al: New Arch: Future Generation Internet Architecture. Final Technical Report <http://www.isi.edu/newarch/>
- [2] Gaiser & Werner, 2007
- [3] <http://www.dcache.org/>