

# The German Grid Initiative: A Uniform Accounting Service in Multiple Middleware Environments

Jan Wiebelitz<sup>1</sup>, Stefano Dal Pra<sup>2</sup>, Wolfgang Müller<sup>1</sup>,  
and Gabriele von Voigt<sup>1</sup>

<sup>1</sup>Regional Computing Center for Lower Saxony, Leibniz Universität Hannover,  
Schloßwender Strasse 5, 30159 Hannover, Germany  
{wiebelitz, wmueller, vonvoigt}@rrzn.uni-hannover.de

<sup>2</sup>Istituto Nazionale di Fisica Nucleare (INFN), Sez. di Pavola,  
Via Marzolo 8, 35131 Padova, Italy  
stefano.dalpra@pd.infn.it

**Abstract.** The development of business models is of great importance for the German Grid initiative D-Grid, which is trying to establish a sustainable and self-financing long-term operating e-science infrastructure all over Germany. To guarantee the continuity and future growth of this infrastructure it is necessary to become independent of limited public funding and achieve the required revenues by chargeable services, support and developments. The business model of the German Grid Support Organization comprises the development, provision and support of generic services for the usage by all Grid communities within D-Grid. The realization of these business models is based on the development of a comprehensive accounting service that enables a seamless recording of the usage of all resources. The special challenge lies in the heterogeneous middlewares, used within D-Grid. The paper describes the development of a highly integrative accounting system as a component of the service portfolio of the German Grid Support Organization.

**Keywords:** Accounting, D-Grid, DGAS, Grid, HLRmon.

## 1 Introduction

The German Grid initiative (D-Grid) was launched in 2004 and is funded by the German Ministry for Education and Research (BMBF). D-Grid currently comprises 27 projects, embraces over 100 partner organizations from science and industry, and is part of the high-tech strategy for Germany. In the first funding period (2004 - 2008), the D-Grid Integration Project (DGI) aimed to build a sustainable Grid infrastructure with a major focus on the support of the three widely used Grid middleware packages Globus Toolkit [6], LCG/gLite [2] and UNICORE [5]. At the end of the first phase, a core Grid infrastructure for the German scientific community was built that provides a portfolio of Grid services. The activities of the second funding period of the DGI aim to extend

and strengthen the production qualities of the infrastructure, and to develop and realize strategies to establish and maintain the D-Grid infrastructure and services even beyond the period of funding.

The focus on the sustainable - and therefore self-financing - operation of a Grid infrastructure requires technical and non-technical solutions. Non-technical issues to be solved are the creation of business models for the communities, and overall business strategies that contain market models for the provision of chargeable services. The technical issues address accounting, pricing and billing.

The D-Grid itself was transformed into a public company under German law, underwritten with the name D-Grid GmbH. After the funding period, it will gather its revenues by membership-fees from Grid resource providers (GRP), Grid service providers (GSP) and the community Grids (CG). Another source of income will be chargeable individual services, such as:

- Customization of generic software components of the Grid infrastructure to meet the individual needs of the GRP, GSP and CGs
- Support of software components of the Grid infrastructure, including training, consulting, helpdesk and second level support
- Providing of information on new developments, available services and conferences via the Internet
- Consulting services for GSP and GRP concerning the integration of resources, services and applications in the Grid infrastructure with special regards to the open standards
- Development of modifications of software components to guarantee their interoperability with the D-Grid infrastructure
- Consulting services for software engineers concerning the open standards of the OGF
- Providing a test bed for developers in order to evaluate software components, modifications and new applications
- Administration of central components of Grid infrastructures based on D-Grid SLAs
- Modification and adjustment of SLAs according to the requirements of the customers and their software components
- Consulting with providers concerning the setup of their Grid infrastructures with special regards to the compatibility with the SLAs of the D-Grid.

The German Grid Support Organization (DGSE) as a central part of the D-Grid Integration Project will provide these services. Besides these chargeable services provided by the DGSE users of the D-Grid resources have to pay for the resource usage. From this intention arises two requirements. The first requirement is the accounting independent from the used middleware. The second requirement is a comprehensive accounting of all used resources.

As the D-Grid project aims to support the job submission using the three mentioned middleware packages there are two possible solutions to account the resource usage. The first solution is to use multiple accounting systems and consolidate the accounting records to get merged accounting information for every user. The other solution is to apply a uniform accounting system to support the

accounting of the resource usage independent of the used middleware package. Because of a missing accounting service in the current deployed release of UNICORE it was decided to adapt one existing Grid accounting solution for the development of a uniform accounting service in D-Grid.

In the following section the decision for DGAS as the accounting system for D-Grid is reasoned. Furthermore the challenges for an uniform accounting system are summarized and the necessary modifications that enable DGAS to account all resource usage independent from the used middleware are described. The paper concludes with a description of further steps to a comprehensive accounting in D-Grid.

## 2 Uniform Accounting

One of the fundamental components of a sustainable Grid infrastructure operation is the precise accounting of the resource usage. In the context of the Grid infrastructure, resources can be compute or storage resources, as well as Grid services, applications or datasets. Currently available Grid accounting systems, such as SGAS [4], DGAS [11] or APEL [1] are designed to account the usage of compute resources for a specific middleware. SGAS developed in the Swe-Grid project is the appropriate accounting system for Globus Toolkit. DGAS and APEL are particularly developed for the LCG/gLite middleware. DGAS was developed originally in the DataGrid project. Further development is done within the EGEE project. DGAS is the accounting system of choice in the D-Grid project. Besides the use in D-Grid, the primary deployment of DGAS is in the INFN-Grid, the Italian partner in EGEE.

### 2.1 Challenges of Accounting in Heterogeneous Grid Infrastructures

When supporting the three middleware packages Globus Toolkit, LCG/gLite and UNICORE in the D-Grid project, the challenge arises to account resource usage independent from the used middleware. As most of the Grid accounting systems use the accounting information provided by the batchsystem, which enqueues the submitted jobs to the clusters of the worker nodes, the mapping of the user-account to a unique identifier of the user is the crucial point. In the Grid context the Distinguished Name (DN) of the users X.509 certificate is such a unique identifier. In Grids X.509 certificates are used to authenticate the user and, in conjunction with the belonging private key, to enable resources to act on behalf of the user [14].

The accounting systems DGAS and APEL interpret the gatekeeper logs of the LCG/gLite Computing Element to get information about the user-to-account-mapping as LCG/gLite provides a dynamic mapping to pool-accounts. Whereas SGAS requires the provision of accounting related information in the job description alternatively a mapping file could be used. In UNICORE 5 there is no accounting system implemented.

To comply with the demand of the BMBF of a minimal invasive adaption, DGAS was chosen as the accounting system in the D-Grid project because of

three benefits. The first benefit is the scalability reached by the hierarchical structure DGAS offers to build a distributed accounting system. Subordinated Tier 1 accounting servers at several sites are able to submit their accounting records to a superior, central accounting server. The deployment in the INFN-Grid for example comprises 43 sites in a two-layered hierarchical accounting infrastructure. The second benefit is the ability to include benchmarks<sup>1</sup> in the accounting records to gain comparable data of the resources usage, which is an important issue in a heterogeneous Grid infrastructure. For an interpretation of the accounting information it is necessary to get parameters which enable the comparability of different resources of the same type, an interpretation could be the charging of the users resource consumption. The third benefit is the adoption to the requirements in the D-Grid with only marginal changes to the original software.

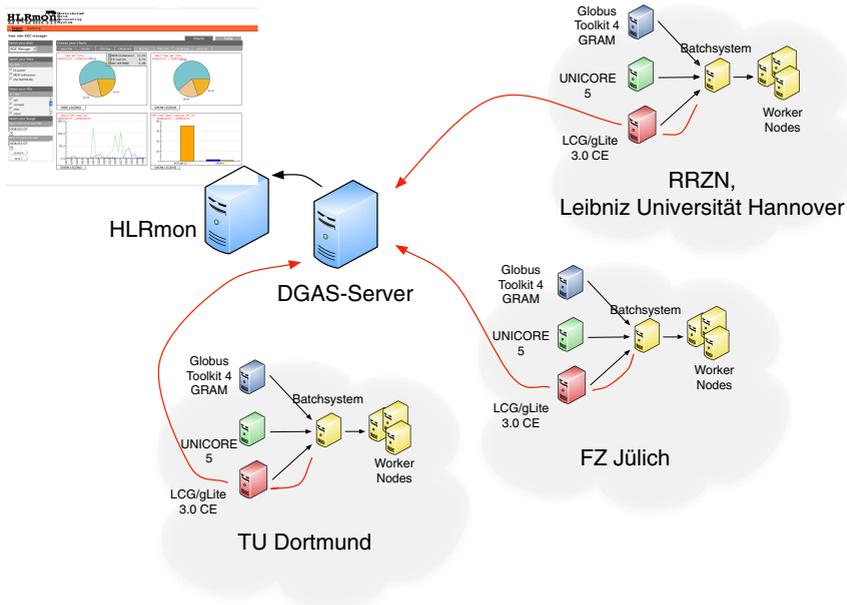
## 2.2 Deployment in the D-Grid

The DGAS accounting system is actually in the deployment phase and the sites TU Dortmund, FZ Jülich and RRZN from the Leibniz Universität Hannover are currently integrated in the uniform accounting. In the D-Grid infrastructure not all services are deployed to all sites. Generic services like monitoring, VO membership services and information services for example are located at selected sites. Every site in D-Grid has to provide access services for the compute- and storage resources. The access services for the compute resources are installed one dedicated *headnodes* for the three widely used middleware packages Globus Toolkit, LCG/gLite and UNICORE. These headnodes are connected to a central batchsystem server that enqueues the jobs to submit them to a cluster of worker nodes (see figure 1). The DGAS client is installed on the headnode for the LCG/gLite middleware the Computing Element (CE) having access to the accounting data of the batchsystem. The DGAS client merges the accounting information of successfully completed jobs with the users unique identifier, the DN and information concerning his VO-membership into usage records. These extended records are submitted to a central DGAS server and are stored in the *Home Location Register* (HLR), the central accounting database.

In D-Grid, users are mapped statically to accounts with a semi-fixed account name, which is constructed by a two-character prefix describing the organization (the flexible part), a two-character abbreviation of the virtual organization (VO) and a four-digit user-id. To account the resource usage independent of the used middleware, DGAS has to be modified. These modifications affect the mapping from the users account included in the accounting records to the users DN. This is realized by a static mapping of the local accounts to an extended version of the users DN instead of a dynamic mapping that results from the use of pool accounts. The use of this extended DN which includes the users VO membership is necessary because in D-Grid users have the possibility of a multi-VO-membership. The extension of the users DN has no effect to any other applications because of the constricted usage of these extended DN only for the accounting service.

---

<sup>1</sup> SPECint2000 and SPECfp2000.



**Fig. 1.** Current accounting infrastructure in the D-Grid project

### 2.3 Visualization of Accounting Information

HLRmon, the visual Web-based component of DGAS, enables access to accounting information for managers, administrators and users [3]. HLRmon gathers the accounting information via the DGAS command line tools to query the HLR accounting database and displays a variety of accounting information like Jobs per VO, Walltime per VO, CPUtime per day. Information is presented in a graphical or tabular form in order to satisfy different user needs.

Due to its ability to authenticate clients through certificate and access rights management, it can a-priori restrict the selectable items range offered to the Web user, so that sensitive information will only be provided to authorized people.

Access to the accounting information, provided by HLRmon, is granted to users whose identity is authenticated through a valid digital certificate, which must be released by a trusted Certification Authority and installed in the client's browser. HLRmon matches user's identity with his access role privileges and preferences. This allows one to directly offer activity reports for the proper subset of information related to role membership, and persistently keeps track of user's preferences (such as date interval or other selections) through subsequent visits. The following four different roles have been defined:

- Regional Operation Center (ROC) Manager  
Can access whole information about activities in every Site, VO and VO users (see figure 2).

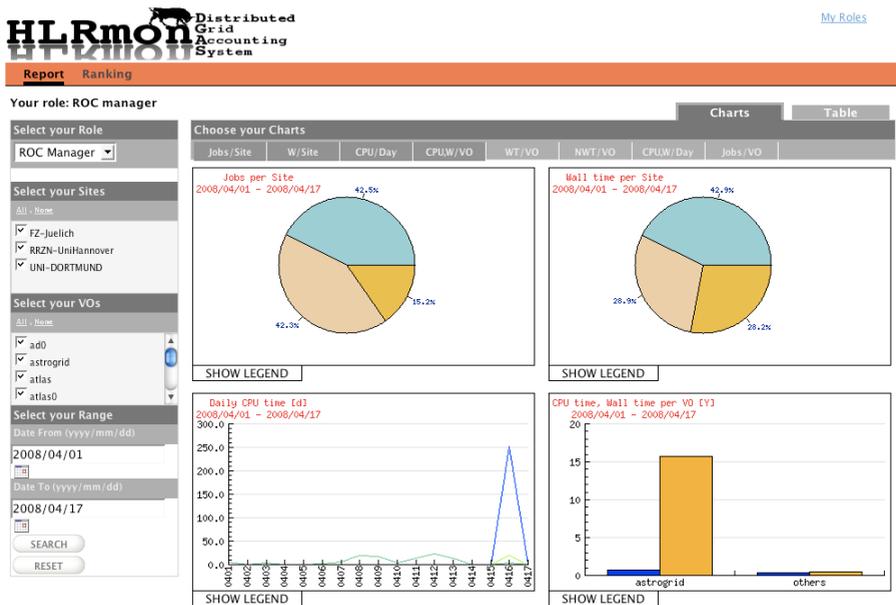


Fig. 2. HLRmon - Roc-Manager View

- Site Manager
  - Can access whole information about activities of every VO and VO users only on sites where the Site Manager role is granted.
- VO Manager
  - Can access information about activities of VOs and VOs users across all sites in the ROC, only on VOs where the VO Manager role is granted.
- VO user
  - Can access information about VO user activities only for VOs where the VO user role is granted. This is the only role that can only access information but its own activities.

Another important fact to consider is that the DGAS system is designed to robustly account against transient malfunctioning of Grid components. Assuming for example that a given Computing Element fails to deliver information to his HLR, this information will be retained at CE level and finally transmitted when things are back ok again.

Side effects regarding the visualization of accounting information could be caused due to the fact that these information were assigned to the date as the job execution starts and not when it is finished. To avoid these undesired side effects in HLRmon the nightly update routine is always executed to retrieve latest 60 days of activity in the past (this is of course a configurable parameter), and this makes sure enough that transient effects leave no persistent tracks in HLRmon once passed.

The HLRmon Web interface helps to reveal transient problems at site level by highlighting site names in red, from which no data are collected for six days. Moving the mouse arrow over that site-name, a tooltip would appear telling how many days have passed since the reported activity. A further feature implemented in newer releases of HLRmon permits site managers and VO managers to easily produce a set of monthly or quarterly reports, consisting of a set of twelve or four charts with the same type of chart, each relevant to a distinct month or quarter.

### 3 Future Work

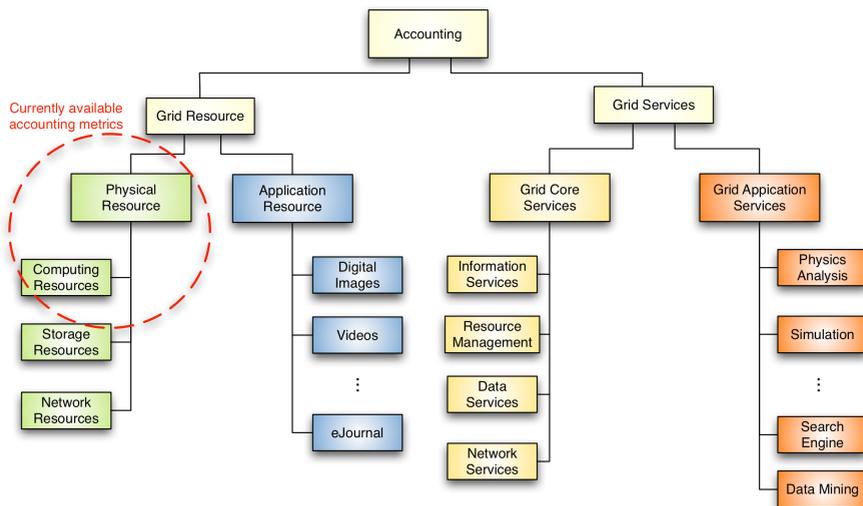
Today the accounting of compute resources usage is provided by every Grid accounting system. The accounting of the usage of storage resources is in progress [13] and has to be extended to support a broader range of storage systems. To provide comprehensive accounting information to enable a seamless charging of resource usage, further accounting metrics have to be identified, and sensors measuring usage have to be implemented. For the identification of these missing metrics the requirements of the scientific and industry Grid resource consumers and resource providers from both ranges have to be inquired.

As resources are not only compute and storage systems, but also for example applications and data, sensors have to be developed to measure the usage of these resources as well. Accounting metrics can be classified as follows [10] (see figure 3):

- Physical resources  
Resources used for compute, storage, network links, but also for example radio telescopes or industrial robots
- Application resources  
Data like images or video for example, which are processed with Grid applications
- Grid core services  
Services that enables Grid computing
- Grid application services  
Software and applications provided for the use by the Grid users

Physical resources and Grid core services are building the Grid infrastructure and have to be accounted by a general Grid accounting service. For a sustainable operation of a Grid infrastructure these metrics must be economically cleared. The application resources and the Grid application services are community specific and must be accounted and cleared within the VOs. To enable a comprehensive accounting by collecting usage information about the resource usage of all these classes, dedicated sensors must be developed that submit the data into a small set of database tables to reach a high standard of comparability. Development of standardized usage records, according to the Open Grid Forum standardized usage records (UR) [7] for compute usage is a necessary task for the near future.

Another focus of future activities is the integration of the resource usage service (RUS) interface [9], [12] to DGAS a deliverable of the OMII-Europe [8]



**Fig. 3.** Classification of Accounting Metrics (from [10])

project to ensure the interoperability with other accounting systems. The RUS, in conjunction with the UR, forms the basis of an accounting framework, which enables the integration of small Grids into a national Grid infrastructure.

## References

1. Byrom, R., Cordenonsi, R., Cornwall, L., Craig, M., Abdeslem, D., Ducan, A., Fisher, S., Gordon, J., Hicks, S., Kant, D., Leake, J., Middleton, R., Thorpe, M., Walk, J., Wilson, A.: APEL: An implementation of Grid accounting using R-GMA (September 2005), <http://www.gridpp.ac.uk/abstracts/allhands2005/apel.pdf>
2. Berlich, R., Kunze, M., Schwarz, K.: Grid computing in Europe: from research to deployment. In: Buyya, R., Coddington, P., Montague, P., Safavi-Naini, R., Sheppard, N., Wendelborn, A. (eds.) Proceedings of the 2005 Australasian Workshop on Grid Computing and E-Research - vol. 44, Newcastle, New South Wales, Australia. Conferences in Research and Practice in Information Technology Series, vol. 108, pp. 21–27. Australian Computer Society, Darlinghurst (2005)
3. Dal Pra, S., Fattibene, E., Gaido, L., Misurelli, G., Pescarmona, F.: HLRMON: A Role-based Grid Accounting Report Web Tool. In: Proceedings of Int, Conference in High-Energy and Nuclear Physik 2007, CHEP 2007 (to appear, 2007)
4. Elmroth, E., Gardfjll, P., Mulmo, O., Sandholm, T.: An OGSA-Based Bank Service for Grid Accounting Systems. In: Proceedings of the 2nd international Conference on Service Oriented Computing, New York, NY, USA (November 2004), <http://delivery.acm.org/10.1145/1040000/1035207/p279-sandholm.pdf>
5. Erwin, D.W., Snelling, D.F.: UNICORE: A Grid Computing Environment. In: Sakellariou, R., Keane, J.A., Gurd, J.R., Freeman, L. (eds.) Euro-Par 2001. LNCS, vol. 2150, pp. 825–834. Springer, Heidelberg (2001), <http://citeseer.ist.psu.edu/erwin01unicore.html>

6. Foster, I.: Globus Toolkit Version 4: Software for Service-Oriented Systems. In: Jin, H., Reed, D., Jiang, W. (eds.) NPC 2005. LNCS, vol. 3779, pp. 2–13. Springer, Heidelberg (2005),  
<http://www.globus.org/alliance/publications/papers/IFIP-2005.pdf>
7. Mach, R., Lepro-Metz, R., Jackson, S., McGinnis, L.: Usage Records - Format Recommendation, Draft Recommendation, Global Grid Forum (2003),  
<http://www.ogf.org/documents/GFD.98.pdf>
8. <http://omii-europe.org/OMII-Europe/>
9. Pace, M.: Preliminary Implementation of a RUS Interface for DGAS (July 2007), [http://www.to.infn.it/grid/accounting/techrep/DGAS-RUS-Prototype-1\\_0.pdf](http://www.to.infn.it/grid/accounting/techrep/DGAS-RUS-Prototype-1_0.pdf)
10. Pettipher, M.A., Khan, A., Robinson, T.W., Chan, X.: Review of Accounting and Usage Monitoring Final Report (July 2007), [http://www.jisc.ac.uk/media/documents/programmes/einfrastructure/jisc\\_aum\\_final\\_report\\_wth.pdf](http://www.jisc.ac.uk/media/documents/programmes/einfrastructure/jisc_aum_final_report_wth.pdf)
11. Piro, M., Guaresi, A., Werbrouck, A.: An Economy-based Accounting Infrastructure for the DataGrid. In: Proc. 4th Int. Workshop on Grid Computing (GRID2003), Phoenix, AZ (November 2003),  
<http://www.to.infn.it/grid/accounting>
12. Piro, R.M., Pace, M., Ghiselli, A., Guarise, A., Luppi, E., Patania, G., Tomassetti, L., Werbrouck, A.: Tracing Resource Usage over Heterogeneous Grid Platforms: A Prototype RUS Interface for DGAS. In: Proceedings of the Third IEEE International Conference on e-Science and Grid Computing, pp. 93–101 (December 2007),  
<http://ieeexplore.ieee.org/iel5/4426856/4426857/04426876.pdf>
13. Scibilia, F.: Accounting of Storage Resources in gLite Based Infrastructures. In: Proceedings of the 16th IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises, pp. 273–275 (June 2007),  
<http://www.ieeexplore.ieee.org/iel5/4407104/4407105/04407170.pdf>
14. Tuecke, S., Welch, V., Engert, D., Pearlman, L., Thompson, M.: Internet x.509 public key infrastructure(PKI)proxy certificate profile. RFC 3820 (2004),  
<http://www.ietf.org/rfc/rfc3820.txt>