Adopting GLUE 2.0 for an Interoperable Grid Monitoring System

Timo Baur[1], Rebecca Breu[2], Tibor Kálmán[3], Tobias Lindinger[4], Anne Milbert[5], Gevorg Poghosyan[6], Mathilde Romberg[2]

[1] Leibniz Supercomputing Centre, Boltzmannstr. 1, D-85748 Garching, Germany
[2] Forschungszentrum Jülich GmbH, Jülich Supercomputing Centre, D-52425 Jülich, Germany
[3] Gesellschaft für wissenschaftliche Datenverarbeitung mbH, Am Fassberg 11, D-37077 Göttingen, Germany
[4] Ludwig-Maximilians-University Munich, Oettingenstr. 67, D-80538 Munich, Germany
[5] Max Planck Institute for Gravitational Physics, Am Mühlenberg 1, D-14476 Golm, Germany
[6] Forschungszentrum Karlsruhe GmbH, Steinbuch Centre for Computing, Hermann-von-Helmholtz-Platz 1, D-76344 Eggenstein-Leopoldshafen, Germany

Abstract

Interoperability is an important issue, especially in Grid infrastructures that use multiple, inhomogeneous Grid middlewares. Monitoring such heterogeneous scenarios requires interoperability of different information services. The approach taken bases on an integration proxy together with a standardized information model. This proxy functions as a bridge between the information models used in different Grid information services by using GLUE 2.0 as the homogeneous common model. This paper outlines the chosen architecture and details the experiences gained with the GLUE 2.0 schema.

1 Introduction

Many of today's large Grid infrastructures, such as DEISA, EGEE and TeraGrid, are based on a single middleware with a single monitoring component in concert with additional, middleware-independent components or frameworks. Other Grid infrastructures employ different middleware systems, each with proprietary monitoring components. The German Grid Initiative's infrastructure (D-Grid), for example, conveys such diversity by providing access to computing resources through multiple middlewares at the same time [1]. Resource monitoring in these environments tends to get inhomogeneous and patchy and thus interoperability becomes an important issue.

In this paper, we describe an interoperable Grid monitoring system that bases on the GLUE 2.0 information model [2]. The system uses translators for the conversion of different information models into the standardized schema and vice versa. As a byproduct a data repository with homogeneous monitoring information of a heterogeneous environment can easily be achieved.

The remainder of this paper is structured as follows: In Section 2 the interoperability problem of Grid infrastructures is discussed as well as related work in the field of interoperable monitoring. In Section 3 we introduce the scenario of a comprehensive and interoperable monitoring service for D-Grid. The chosen approach and the resulting architecture is described in Section 4, and the experiences made are given in Section 5, followed by a summary.

2 Motivation and Related Work

Grid computing defined as *coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations* [3] implies homogeneous access to a pool of inhomogeneous resources shared among different Virtual Organizations. Several middleware systems have been developed independently of each other to realize this vision. Large Grid infrastructures such as DEISA, EGEE and TeraGrid have been set up specifically enabling scientific communities to tackle large, complex problems. These infrastructures are usually based on one middleware but challenges exist which demand the use of resources governed by more than one infrastructure. This raises the question of interoperability of Grid infrastructures and middleware platforms respectively. Because of their different development history and usage models, these middleware systems are not interoperable per se. Interoperability has to be achieved through additional efforts, it has to cover all aspe cts of an infrastructure: job submission, data transfer, accounting, job status monitoring as well as monitoring of sites, resources, services and providers. The latter is subject to the German research project D-MON (Horizontal Integration of Resource- and Service Monitoring in D-Grid, [4]) whose results are presented in this paper.

Interoperable Grid monitoring has been a topic since many years [5], [6], and standardization efforts have been made by e.g. the Open Grid Forum with respect to monitoring architecture [7] and data schema [2]. Furthermore, middleware-independent monitoring frameworks have been developed for site and service monitoring aspects. Examples are Nagios [8] and INCA [9] which detect Grid infrastructure problems by executing periodic, automated testing of Grid software and services using middleware and infrastructure specific procedures. These tools allow Grid operation centers to get an overview of the status of sites and services in an infrastructure even if it is inhomogeneous. An interoperable Grid information service is still missing. Additionally, most monitoring systems still provide data from a providers view, unable to transparently provide the data in community-specific ways and related to specific virtual organizations (VOs).

3 A Grid Monitoring Scenario

The operation of Grid infrastructures needs Grid monitoring as a management functionality. For example, failure identification and elimination, accounting, performance analysis and scheduling rely on Grid-wide status information.

All components of a Grid monitoring system must be able to interoperate efficiently to achieve good overall performance as well as consistent and complete data. The monitoring of heterogeneous infrastructures requires data to be gathered from information services with proprietary interfaces and data models. In addition, homogeneous interpretation and presentation of the data is needed.

As an example for a Grid with heterogeneous middleware components the environment run by the German Grid initiative D-Grid is considered where many resources are available through multiple Grid middlewares [1]. Site monitoring, service monitoring, and job monitoring are done in several different, partly concurrent ways[10], and an easy access to combined data is not available. The broad spectrum of monitoring tools ranges from monitoring systems based on different middlewares to monitoring systems based on middleware-independent architectures. Beside the co-existing information systems for the different middlewares Globus, Glite and UNICORE (e.g. MDS[1], BDII[2], CIS[3]) there exist monitoring tools for special purposes [11], for monitoring levels [9], and for user-centric monitoring for special application environments, e.g community specific job monitoring [12]. This led to a very heterogeneous, non-interoperable set of monitoring tools within D-Grid.

To resolve this situation, Grid standards and best practices have been investigated but the question remains how interoperability could be realized for the monitoring of dynamic heterogeneous Grid infrastructures. Especially when adding a new monitoring system refactoring of all the Grid components has to be avoided.

4 The Integration Proxy Approach

The emerging GLUE 2.0 information model [2] is a suited candidate for a common monitoring data model of a non-intrusive, interoperable Grid monitoring system. At the time being it is proposed as as standard by the Open Grid Forum. The GLUE 2.0 schema describes the main characteristics of a Grid infrastructure (see Figure 1):

- VO modeling (UserDomains),
- mapping and access policies (MappingPolicy, AccessPolicy),
- allocation of resources and services to VOs (Endpoint),
- resource and service scenarios (ComputingService, ComputingManager), and
- resource providers (AdminDomains).

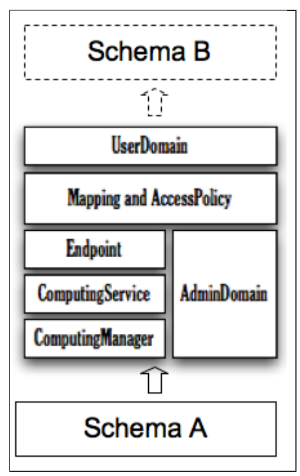

Fig. 1: Schema Mediation

[1]Meta Directory Service
[2]Berkeley Database Information Index
[3]Common Information Service

In addition to the common data schema the key to achieving interoperability is an integration proxy, which uses non-intrusive adapter modules to connect to the different Grid monitoring systems in place. The proxy integrates data from the different monitoring systems through a three step process called ETL (Extract – Transform – Load). It connects to the native interfaces, extracts data in XML format, transforms the data using XSLT[4] procedures, and uploads the transformed data into a repository. As repository an SQL database is used which implements the GLUE 2.0 schema.

Particularly, the GLUE 2.0 schema is used for the mediation of Grid resource and service monitoring data as illustrated in Figure 1: All data gathered from the connected monitoring services using schema A is transformed into GLUE 2.0. From that format, the data can be transformed into any other schema B. Thus, it is possible to cross-provide monitoring data, for example from MDS4 into the interoperable monitoring service and from there into CIS or vice versa.

A standardized provisioning of the collected monitoring data is achieved by an OGSA-DAI[5] interface on top of the database. The OGSA-DAI interface developed in D-MON provides access to the monitoring data for Grid services and programs. For end-users and Grid administrators a web interface is provided, which consists of portlets for the Gridsphere framework. The system also supports Virtual Organization-specific views on the data by using Grid-wide Policy Decision Points which provide a mapping of resources and services to virtual organizations [13].

Figure 2 shows the integrative architecture that realizes data integration as discussed in this section.

5 Experiences

The architecture as described in Section 4 has been implemented in the D-MON project. For the transformations D-MON focused on MDS4, BDII, and CIS as sources. MDS4 delivers its data in GLUE 1.1 and GLUE 1.3, BDII in LDAP Data Interchange Format (LDIF) using GLUE 1.2, and CIS in GLUE 2.0. For BDII one extra transformation from LDIF to XML is needed using Directory Service Markup Language. The data transformations from GLUE 1.1, GLUE 1.3 and LDIF show that not everything can be mapped. There are differences in semantics and thus the transformations may cause loss of information. As we are dealing with different releases of the GLUE schema we have been able to gather and transform the important data for the relevant values without a loss of accuracy. Examples for transformed data items are the measurements for ComputingResources, ComputingServices, and StorageResources and also site-related information, like site location, description and contact addresses.

The database schema of the repository is extended with respect to the GLUE 2.0 model by data provenance information. So every data item which is transformed is labeled with the source information service, the time and date of

[4]eXtensible Stylesheet Language Transformation

[5]Open Grid Services Architecture - Data Access and Integration

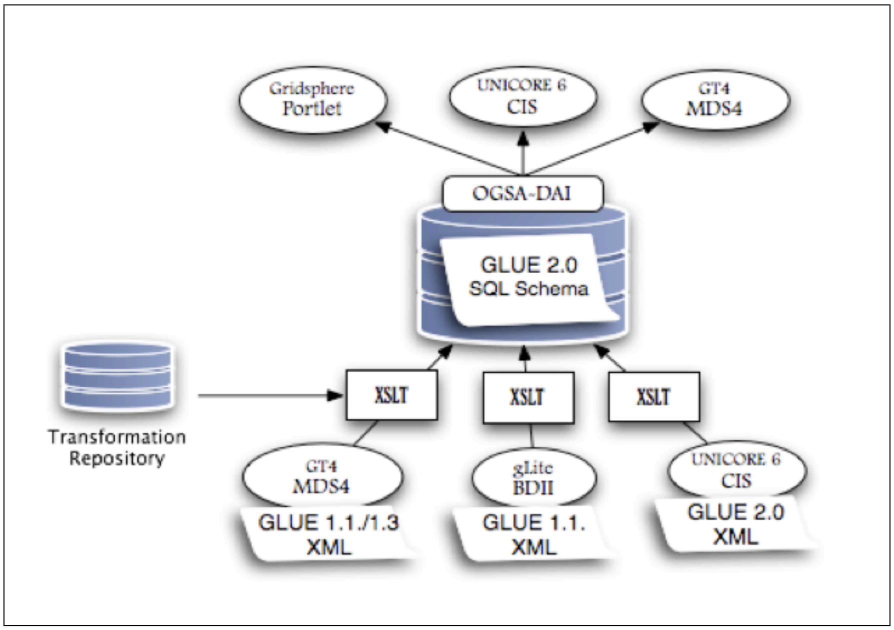

Fig. 2: Integrative Architecture (interoperability part)

retrieval, as well as the IP-address of the source. Figure 3 shows an example of database entries referring to ComputingServices, which have been gathered from the different information services CIS, MDS4, and BDII (given in the fourth column titled informationService).

The prototype system provides VO-specific views on the data exploiting the mappings of resources to VOs available from GRRS[6], a D-Grid-specific database for VO resource management. It acts as policy information base for the resource - VO mappings. This information allows to generate GLUE 2.0 AccessPolicies, which are then used to generate appropriate database views. As shown in Figure 4, the developed OGSA-DAI client takes a reference to a VO as an environment variable and prints out the related entries.

One of the results is that the integrated data discloses structural shortcomings of the monitored multi-Grid configurations. Often duplicate entries exist, which reference the same underlying resource, but not in the same way. For example, site identifiers are not uniquely defined throughout all middlewares. Nevertheless, unique Grid-wide identifiers are vital for statistics, efficient searches and scheduling. To warrant a consistency in labeling, multi-middleware and multi-infrastructure Grid projects need dedicated decision points which assign such names or numbers.

This is an organizational issue similar to issues which have already been

[6] Grid Resource Registration Service

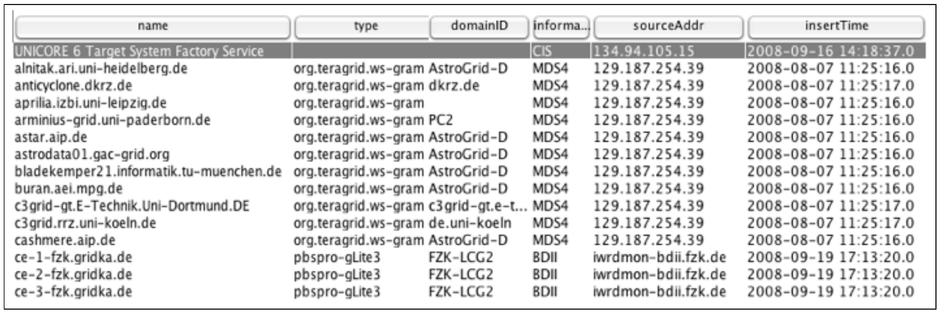

name	type	domainID	informa...	sourceAddr	insertTime
UNICORE 6 Target System Factory Service			CIS	134.94.105.15	2008-09-16 14:18:37.0
alnitak.ari.uni-heidelberg.de	org.teragrid.ws-gram	AstroGrid-D	MDS4	129.187.254.39	2008-08-07 11:25:16.0
anticyclone.dkrz.de	org.teragrid.ws-gram	dkrz.de	MDS4	129.187.254.39	2008-08-07 11:25:17.0
aprilia.izbi.uni-leipzig.de	org.teragrid.ws-gram		MDS4	129.187.254.39	2008-08-07 11:25:16.0
arminius-grid.uni-paderborn.de	org.teragrid.ws-gram	PC2	MDS4	129.187.254.39	2008-08-07 11:25:16.0
astar.aip.de	org.teragrid.ws-gram	AstroGrid-D	MDS4	129.187.254.39	2008-08-07 11:25:16.0
astrodata01.gac-grid.org	org.teragrid.ws-gram	AstroGrid-D	MDS4	129.187.254.39	2008-08-07 11:25:16.0
bladekemper21.informatik.tu-muenchen.de	org.teragrid.ws-gram	AstroGrid-D	MDS4	129.187.254.39	2008-08-07 11:25:16.0
buran.aei.mpg.de	org.teragrid.ws-gram	AstroGrid-D	MDS4	129.187.254.39	2008-08-07 11:25:16.0
c3grid-gt.E-Technik.Uni-Dortmund.DE	org.teragrid.ws-gram	c3grid-gt.e-t...	MDS4	129.187.254.39	2008-08-07 11:25:17.0
c3grid.rrz.uni-koeln.de	org.teragrid.ws-gram	de.uni-koeln	MDS4	129.187.254.39	2008-08-07 11:25:17.0
cashmere.aip.de	org.teragrid.ws-gram	AstroGrid-D	MDS4	129.187.254.39	2008-08-07 11:25:16.0
ce-1-fzk.gridka.de	pbspro-gLite3	FZK-LCG2	BDII	iwrdmon-bdii.fzk.de	2008-09-19 17:13:20.0
ce-2-fzk.gridka.de	pbspro-gLite3	FZK-LCG2	BDII	iwrdmon-bdii.fzk.de	2008-09-19 17:13:20.0
ce-3-fzk.gridka.de	pbspro-gLite3	FZK-LCG2	BDII	iwrdmon-bdii.fzk.de	2008-09-19 17:13:20.0

Fig. 3: Computing Services from Different Middlewares in the Integrated Database

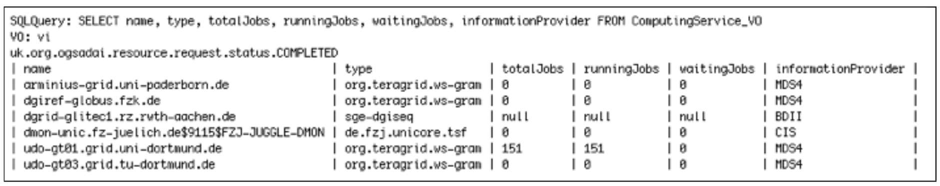

```
SQLQuery: SELECT name, type, totalJobs, runningJobs, waitingJobs, informationProvider FROM ComputingService_VO
VO: vi
uk.org.ogsadai.resource.request.status.COMPLETED
| name                                  | type                 | totalJobs | runningJobs | waitingJobs | informationProvider |
| arminius-grid.uni-paderborn.de        | org.teragrid.ws-gram | 0         | 0           | 0           | MDS4                |
| dgiref-globus.fzk.de                  | org.teragrid.ws-gram | 0         | 0           | 0           | MDS4                |
| dgrid-glitec1.rz.rwth-aachen.de       | sge-dgiseq           | null      | null        | null        | BDII                |
| dmon-unic.fz-juelich.de$9115$FZJ-JUGGLE-DMON | de.fzj.unicore.tsf | 0    | 0           | 0           | CIS                 |
| udo-gt01.grid.uni-dortmund.de         | org.teragrid.ws-gram | 151       | 151         | 0           | MDS4                |
| udo-gt03.grid.tu-dortmund.de          | org.teragrid.ws-gram | 0         | 0           | 0           | MDS4                |
```

Fig. 4: OGSA-DAI View on Resources and Services of a Specific VO

solved in 1984 for the Internet, when the Internet Assigned Numbers and Names Authority IANA was established. The IANA was responsible for the assignment of IP-addresses to autonomous systems (AS) and the registration of names in the Domain Name System. Common and large-scale Grid infrastructures also will require common naming standards or policies, assigned at clear policy decision points.

6 Conclusions

GLUE 2.0 is an adequate information model for resource and service monitoring in Grids. It is suitable for Grid scenarios: It allows the definition of VO's as UserDomains and the definition of AccessPolicies, which together eases the generation of views for different VOs. Integration proxies can be used to cache and exchange resource and service monitoring data in an interoperable way using GLUE 2.0 as a mediation schema. The data gathered from different monitoring systems shows, that there are differences in the labeling of sites and the addressing of services, which complicate a consolidation of the data for accounting and scheduling purposes. This is a managerial issue that has to be solved before taking full advantage of the developed interoperable Grid monitoring system.

Acknowledgements. This work has been partly funded by the D-Grid Initiative of the German Federal Ministry of Education and Research (FKZ01IG07010A).

References

1. M. Alef, T. Fieseler, S. Freitag, A. Garcia, C. Grimm, W. Gürich, H. Mehammed, L. Schley, O. Schneider, G.L. Volpato: *Integration of multiple middlewares on a single computing resource*; Future Generation Computer Systems, Available online 20 May 2008
2. S. Adreozzi (Ed.), S. Burke, F. Ehm, L. Field, G. Galang, B. Konya, M. Litmaath, P. Millar, J.P. Navarro: *GLUE 2.0 Specification V2.0 (draft 33)*; Open Grid Forum Specification, 2008
3. I. Foster, C. Kesselman, S. Tuecke: *The Anatomy of the Grid: Enabling Scalable Virtual Organizations*; International J. Supercomputer Applications, 15(3), 2001.
4. *D-MON*, `http://www.d-grid.de/index.php?id=401`, last accessed November 2008
5. S. Zanikolas, R. Sakellariou: *A taxonomy of grid monitoring systems*; Future Generation Computer Systems 21 (2005) 163188
6. A.P. Miller, G. Stewart, G.Cowan: *Monitoring with MonAMI: a case study.*; Journal of Physics: Conference Series 119 (2008) 062037
7. B. Tierney, R. Aydt, D. Gunter, W. Smith, M. Swany, V. Taylor, R. Wolski: *A Grid Monitoring Architecture*; Open Grid Forum Document GFD.7, 2002, http://www.ogf.org/documents/GFD.7.pdf
8. *Nagios*, `www.nagios.org`, last accessed November 2008
9. *INCA*, `http://inca.sdsc.edu/drupal`, last accessed November 2008
10. G. Poghosyan and M. Kunze *Monitoring for multi-middleware grid*; Parallel and Distributed Processing, 2008. IPDPS 2008. IEEE International Symposium on, 14-18 April, IEEE, Miami, Florida, USA.
11. M.L. Massie, B.N. Chun, D.E. Culler: *The Ganglia Distributed Monitoring System: Design, Implementation, and Experience*; Parallel Computing, Vol. 30, Issue 7, July 2004.
12. D. Lorenz, S. Borovac, P. Buchholz, H. Eichenhardt, T. Harenberg, P. Mattig, M. Mechtel, R. Müller-Pfefferkorn, R. Neumann, K. Reeves, Ch. Uebing, W. Walkowiak, Th. William, R. Wismüller: *Job monitoring and steering in D-Grid's High Energy Physics Community Grid*; Future Generation Computer Systems, Available online 28 May 2008.
13. T. Baur: *Functional Analysis and Architecture for Interoperable and DVO-specific Grid Monitoring Services.* In: Proceedings of the Fourth IEEE International Conference on eScience (eScience 2008). IEEE Computer Society Press (2008)