Functional Analysis and Architecture for Interoperable and DVO-specific Grid Monitoring Services

Timo Baur
Munich Network Management Team
Leibniz Supercomputing Centre, Garching/Germany
timo.baur@mnm-team.org

Abstract

Grid operations rely on monitoring services fed by a multitude of sources. These services provide relevant information about the status of the resources and the quality of the services in the infrastructure. In a Grid, they must disseminate this information across multiple administrative domains and heterogeneous technical platforms to users organized in different virtual organizations. In order to achieve this, data must be translated and homogenized dynamically. It must furthermore be tailored and processed in a controllable way, based on its allocation to virtual organizations. To resolve these requirements for the application of Grid resource monitoring, we propose functional modules, which conduct automated data migration and classification based on dynamic virtualization policies and define the composition of monitoring data related to virtual organizations. The resulting system enables an interoperable provisioning of specific monitoring data about the status of distributed Grid resources and services to dynamically changing virtual organizations.

Keywords: policy based Grid monitoring, virtual organizations, Grid interoperability, information sharing

1 Introduction

Grid operations rely on monitoring services fed by a multitude of sources of data and events (e.g. databases, logfiles, sensors, measurement points), which provide relevant information about the status of the resources and the quality of the services in the Grid's infrastructure. Such information services form an integrated data delivery infrastructure which spans multiple (ideally autonomous) resource and service providers and provides information to a multitude of users organized as members of virtual project teams (see e.g. [1]) or communities. These groups of users, together with the virtual[1] resources and services [2] they use and own, are called virtual organizations (VOs) [3], [4],[5].

When different providers and virtual organizations in such a Grid use different technical platforms, which are not interoperable, conflicts may occur, as soon as these platforms must work together or exchange data. Subsequently, information is not shared in efficient and comprehensive ways and expensive integration efforts become necessary.

Furthermore, dynamic changes in the set of a VO's composition (members, virtual resources and services) respectively its consistency, as they happen in 'dynamic virtual organizations' (DVO), may lead to changes in the data an entity in a VO needs to see. Information services which do not comprehensively support such changes, deliver inappropriate and insufficient results.

These problems become evident when multiple different monitoring systems occupy the same Grid environment. Actually, most Grid monitoring systems are not interoperable with each other and, furthermore, deliver provider-specific, overall monitoring data which contains information that is mostly unrelated to a VO's status and processes. They do not provide transparent, VO-specific views on the Grid infrastructure which are dynamically shaped according to a VO's composition.

In the course of this work, a Grid monitoring scenario with multiple monitoring systems is outlined in section 2. It provides a real use case and illustrates the practical relevance of the problem. In section 3, necessary functionalities for interoperable and VO-aware Grid monitoring services are analyzed. Section 4 presents functional modules, which enable a dynamically integrated as well as VO-aware delivery of monitoring data to the consuming entities located in VOs. To achieve this, we describe data transformation functionalities as well as a delivery of VO composition rules which describe a Grid's VO mapping policy. These rules enable the classification of resource monitoring data

[1] in the terminology used in that context, a dichotomy between 'real' and 'virtual' is common, see also section 6

978-0-7695-3535-7/08 $25.00 © 2008 IEEE
DOI 10.1109/eScience.2008.48

and metrics into VO-specific sets. We then use policy enforcement points where a correlation of the delivered data with the delivered VO composition policy takes place and logical views are generated in order to realize a shaping of the data. Section 6 discusses related work. The last section briefly summarizes the findings and presents open research questions in the field.

2 Scenario

In the german D-Grid initiative [6], different communities gathered with the goal to transparently solve their computational problems by using the computing resources of a shared Grid computing infrastructure. This infrastructure bridges different real and virtual resources provided by different organizations, which in specific cases use different technical platforms as middleware. For example, the Grid middlewares Globus Toolkit 4, gLite and Unicore are used, each being monitored by a middleware-specific monitoring service like MDS4, used in the Globus Toolkit 4,the BDII/MDS2 used in gLite and the CIS used in Unicore. These services use different data description schema (information models) and clients, and provide different but intersecting sets of data which are not exchanged with each other, nor provide consistent and comprehensive views for dynamically changing VOs. Nevertheless, the D-Grid is meant as one comprehensive Grid infrastructure which should everywhere have the same capability to provide monitoring information to its users (respectively roles), regardless of the platforms they use and the administrative domains they are located in.

Regarding the aspect of dynamic VO, different monitoring systems in the D-Grid use different methods to configure the composition of local resources in a VO. The local resource providers independently define platform-dependent links or attributes in their local part of the Grid monitoring system to define the target VOs to which monitoring data is sent. For example, a MDS4 instance at a provider's site is configured with an upstream to a community instance of MDS4, or monitoring data for the BDII is tagged with a parameter that outlines the VOs a monitored gLite resource belongs to. In either case, the entries must be changed by the administrators whenever changes in a VO's composition occur and a resource is used (and must be monitored) by additional or less VOs. Independently, a central database is held in the D-Grid (the Grid Resource Registration Service GRRS), where administrators manually enter the mappings of their resources into VOs to keep track of the resource-VO relationships. These entries are not automatically synchronized with the configuration of the monitoring systems. Because of this, dynamic changes of a VO's composition are not supported by the monitoring systems in appropriate time, perimeter and accuracy. Thus, VO-specific monitoring data is not always delivered correctly to the appropriate entities of a VO. This situation leads to considerable challenges for the realization of a comprehensive, VO-specific Grid monitoring to a Grid's users which are organized in VOs.

3 Problem

Large-scale heterogenous Grid monitoring systems must deliver their data deterministically from arbitrary heterogenous data sources to arbitrary users in different, dynamically changing VOs. Besides user requirements on parameters and user interfaces (see e.g. [7]), there are fundamental functional requirements on the distributed design of such a system.

In the following, the main challenges which occur when using monitoring services in such heterogenous Grids are outlined and necessary functionalities for integration and VO-specific data provisioning are shown. Contemplating about the scenario, we face two major structural elements:

- In order to enable a deterministic processing of monitoring data about the complete Grid at all platforms, a dynamically changing environment, which applies specific, non-interoperable and independent technical platforms, must be integrated.
- the monitoring data must be provided dynamically in VO-specific patterns to entities in dynamically changing VOs.

3.1 Interoperability

Regarding the first structural element, the monitoring services in the scenario have characteristics which are (amongst others) common to large scale distributed infrastructures. They

1. use different information models which describe the data syntactically and semantically.
2. use different interface standards and protocols, impeding efforts to connect them.
3. apply configurations and interconnections which may change dynamically with different participating organizations.

These characteristics often result in inefficient data delivery and expensive integration efforts.

To integrate the data and achieve interoperability in the face of these characteristics

1. the data should be integrated ad-hoc at runtime to avoid the costs of repeated larger-scale integration efforts.

2. a flexible and automated translation functionality is necessary.

3. different interfaces and protocols must be supported and a common, preferably standardized, alternative must be given.

4. To support changes, the system must be easily configurable.

3.2 VO-specific Data Provisioning

Regarding the second structural element, it must be noted that DVOs use shared and dynamically changing parts of the resources and services of a Grid. Accordingly, dynamically tuned subsets of data must be delivered to their entities aligned to the resources and services associated or allocated to them.

Issues in that context are:

1. data must be tailored and provided in VO-specific data sets at monitoring service access points.

2. rules about how the VOs in the Grid are composed must be decided upon and provided in order to define what data is VO-specific and to which VO it belongs.

3. one central or multiple distributed, clearly defined and trusted decision point(s) must be pinpointed, which act as a source for these rules.

4. the rules must be applied to the data at all service access points where data is collected.

5. as soon as VOs change dynamically, both, the rules as well as the VO-specific data sets must be updated synchronously.

4 Solution

In order to meet the issues outlined in section 3 and to provide a blueprint for a more detailed definition of processes and software design, we now describe the three major structural blocks of functionalities required. The basic functions, and thus the main components, for a VO-based provisioning of the information gathered from heterogenous information services are:

1. homogenization of the data about the Grid's state: collection and integration of monitoring data from heterogenous sources (e.g. different middleware platforms, Grid monitoring services, information services or management tools).

2. Grid-wide federated caching and distribution of this data.

3. generation of individualized data for users in specific VOs: policy based filtering according to the association or allocation of a data source with a virtual organization.

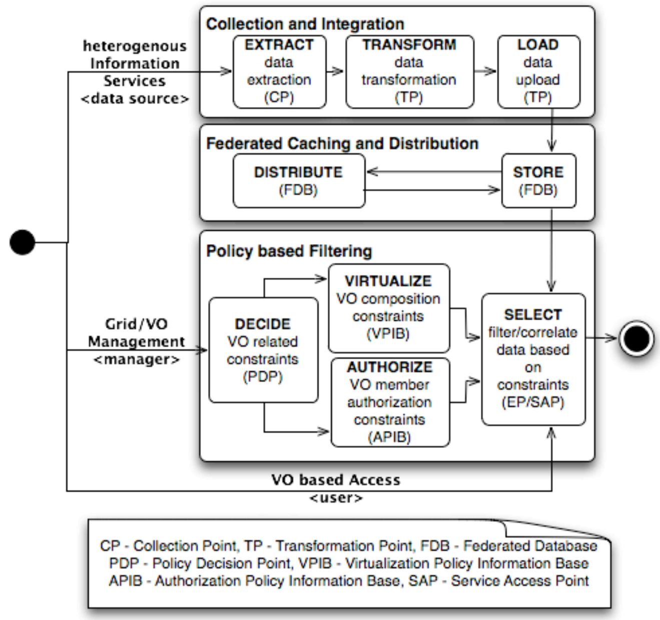

Figure 1. Components for Interoperable and VO-aware Data Provisioning

In the following, these main components as well as their subcomponents are outlined. The state machine chart in figure 1 gives an overview.

4.1 Collection and Integration

The purpose of the first component is to access the different interfaces used by different heterogenous information services and to homogenize the different data models in which monitoring data is encoded in order to assure a common system-wide language to decode the measured data about the Grid resources and services.

For its realization, we propose to realize a steady and automated (if necessary on-demand) data migration into an at runtime at least temporarily homogeneous syntactic and semantic environment. For this purpose, a flexible and automated application of an ETL process (Extract-Transform-Load, as presented for data warehousing in [8]) is used. ETL extracts, transforms and uploads heterogenously encoded data sets into an homogeneously encoded data set. When monitoring data is provided from different, heterogeneously encoded and instrumented data sources which may change often, as it occurs in our scenario, it is particularly useful to implement multiple instances of that process, each one supporting the specific environment of the data source.

The (automated) collection and integration component relies on the following steps:

Extract gets the data from its sources and should support different interface standards. Each source is accessed at a Connection Point (CP), which is usually the interface of a provider-specific information service.

Transform provides transformations to translate and map monitoring data from the data model used at the specific CP into a common data model which acts as a 'lingua franca' . We are aware that there is data that can not be mapped easily and a loss of information can occur, as syntax and semantics of different data models may not be mapped easily (e.g. semantic mappings may not be unique). This is a known problem since long, which we won't solve in general. Nevertheless, we use a pragmatic approach and define specific transformation modules (implemented as XSLTs which define the necessary mappings) for several specific monitoring services as well as a standardized common model which is used to mediate the data.

Load provides the data to the common Grid monitoring service, where the data, encoded in the common model, is kept.

4.2 Federated Caching and Distribution

The purpose of the second component is to store the homogenized data in a database using the common data model as a homogenous environment. Its functionality is to aggregate the resource and service state data gathered from heterogenous provisioning domains in a stateful (database) entity. The database entity may be setup in a single central (one for the Grid) or in a distributed and federated fashion (one at every site). In the latter case, an algorithm for data distribution is necessary. For the purpose of our scenario, a federated setup could use simple replication. This could later be optimized.

The component relies on the following steps:

Store provides methods to store the data gathered by the collection and integration component.

Distribute provides methods and protocols to distribute the cached data over the Grid's network between different provisioning domains, e.g. by replicating the monitoring data of the Grid at every resource provider's site.

4.3 Policy Based Filtering or Correlation

The purpose of the third component is to correlate or filter the data stored in the (federated) cache according to knowledge about the mapping of the monitored resources in the real organizations into their virtual counterparts at specific virtual organizations. Its functionality is to filter and correlate the homogenized state descriptions for the real organizations resources and services. This must happen according to the definition (as well as distribution) of constraints about a virtual organizations actual composition in terms of members, (virtual) services and (virtual) resources.

We define the composition of a VO (or a 'virtualization policy') as a formal set of constraints that describes which (real) entities, specifically resources and services, are mapped onto a VO, and how this should be done. For example, we use rules that just map supercomputing resources onto VOs as seen in 2. In more sophisticated cases, the constraints may also outline more complicated relationships between a RO's resource state measurements and the measurements of the associated VOs.

```
Examples (Scheme: α: A -> VOa):

HPCResource: Resource.SGIAltix -> VO.AstroGrid
HPCResource: Resource.SGIAltix -> VO.TextGrid

PBSQueue: Resource.SGIAltix.PBSa ->VO.AstroGrid
PBSQueue: Resource.SGIAltix.PBSa -> VO.TextGrid
PBSQueue: Resource.SGIAltix.PBSb -> VO.AstroGrid
```

Figure 2. Examples for VO composition rules

To enable Grid monitoring services to dynamically provide data segments according to this composition (which may dynamically change), their behaviour must become influenced. We thus define one or more grid-wide policy decision points (see [9] for policy terminology) which decide about the VO composition constraints, and additionally use policy information bases to allow the VO based Grid monitoring service to query the actual composition of a dynamic VO. After correlating and filtering the cached data as defined by the composition, the service can generate a view containing timely information about the status of the DVOs resources for its users.

The correlation and filtering component relies on the following steps, which are conducted at defined points in the distributed Grid network:

Decide One central or multiple distributed Policy Decision Points (PDP), where grid-wide policies for the actual composition of a VO as well as the memberships and authorization of members and roles are negotiated and agreed upon. Usually both, 'real' organizations, e.g. in the role of the Grid's resource providers, as well as virtual organizations (e.g. in the role of a Grid's customer organizations), are taking part in the decision making. One or more decision points should be clearly defined and available to all parties in the Grid. They could be pinpointed in a central Grid authority which closes all contracts with providers and VOs or in multilateral entities which oversee the effects of all contracts between resource providers and VOs.

Virtualize / VO-composition A Policy Information Base which provides the VO composition rules to the Grid monitoring services, i.e. the mapping of resources and services onto VOs (Virtualization Policy Information Point, VPIB). For this purpose, resource repositories or schedulers could be used.

Authorize / Members A Policy Information Point which provides the a VO membership or authorization rules, i.e. the mapping of users, administrators and other human roles, onto VOs (Authorization Policy Information Base, APIB). For this purpose, e.g., an AA-I (Authentication and Authorization Infrastructure) could be used.

Select A Policy Enforcement Point (PEP) at the Grid monitoring service, which is used to select and classify relevant measurements according to the composition rules and to provide them to authorized VO members. The PEP is ideally located at the point, where the information service is accessed by the users (service access point, SAP).

In order to reduce the complexity of the system and the possibility of policy conflicts (as described e.g. in [10]), we firstly use a a single VPIB and a single APIB to provide the rules.

5 Application

The general functionalities of the components outlined above can be applied to many environments and use cases. In the following, we comment on their relationship to IT Service Management processes and give an overview on a prototype implementation.

5.1 Management Processes

The components' functionalitites can be positioned in the light of IT Service Management processes. As the PDP is the place where the decisions about the VOs' composition as well as their memberships are made, it is directly related with the Grids', the VO's and the resource and service operators SLA and change management. The VPIB and the APIB correspond with CMDB / CMDBf items which describe the VOs' configurations as a target state and the authorizations of the VO's members. As the FDB mirrors the current state of the resources and services in the Grid, it is related to event and incident management processes as well as to monitoring in general. The selection process at the SAP/PEP corresponds with (VO-based) event correlation in the sense that they provide events correlated with the VOs' they happen in, thus realizing a VO-based filtering of the events.

5.2 D-MON: Interoperable and VO-aware Grid Monitoring

We have implemented the outlined components in the ongoing project D-MON, which targets to realize an interoperable and integrated VO-aware information service for the three different middlewares used in the D-Grid. Figure 3 depicts an architectural overview.

A mySQL database was setup to implement the stateful resource of a common information service which provides monitoring information using different interfaces for data access (at its SAP), including the OGF OGSA-conform Database Access Interface OGSA-DAI [11]. The common information service keeps the information gathered from CPs at the Unicore 6 CIS, Globus Tolkit MDS4 and the gLite-based BDII in a mySQL database which uses the GLUE 2.0 SQL schema [12].

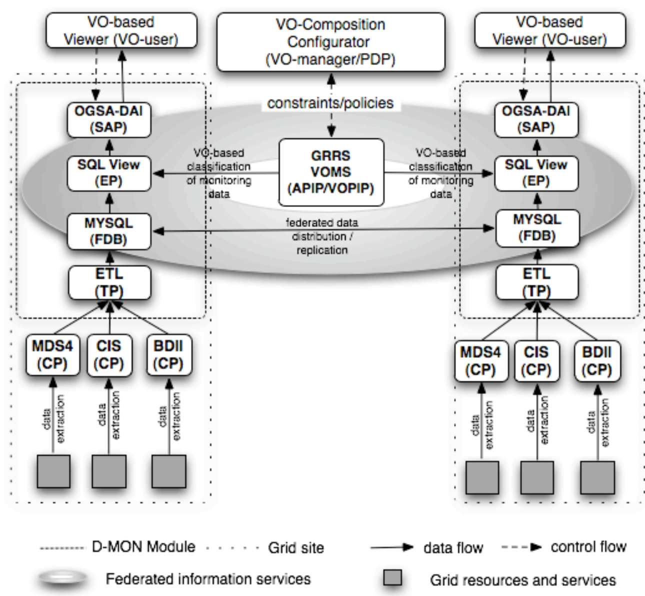

Figure 3. Implementation of the components in the D-MON system

In the ETL component, EXTRACT queries XML data from the middleware-specific information services, an XSLT-Processor TRANSFORMs them using appropriate XSLT mappings into SQL UPDATE statements for the GLUE 2.0 based SQL tables and finally LOADs them into the mySQL database. This database can be federated with replicated databases at different resource provider sites.

For the implementation of the VOPIP, we connected the D-Grid's Grid Resource Repository Service (GRRS), where we retrieve mappings configured by the Grid's VO-management. The EP located at the database uses these mappings to generate SQL views for individual VOs. We are planing to use a SAML-capable VOMS for authorization using the APIP.

Furthermore, as a VO based Viewer, we connected the Portal-Framework Gridsphere to the database in order to provide the data to users in the World Wide Web.

A prototype system is running, which provides VO-specfiic views on data about D-Grid resources to an OGSA-DAI client. The system gathers data from the three major middlewares (which use different GLUE 1.x versions to encode the data) and provides it encoded in the new OGF standard GLUE 2.0.

6 Related Work

Grids based on different middleware platforms and tools still face islands of incompatible technical infrastructure and insufficient technical integration of the VO paradigm.

While problems of Grid interoperation have been solved for some middleware constellations and specific tasks (e.g. for JobSubmission with BES [13]), there is still much ongoing work on the integration of monitoring systems in similar scenarios as outlined in chapter 2. Grid Interoperation is usually a bi-lateral activity between two Grid infrastructures [14] . Examples are the NDGF-EGEE Collaboration, which fostered interoperability between the ARC and the EGEE middlewares, integrating the ARC and the GLUE data schema and connecting accompanying monitoring services using instances of BDII. Another examples are the NAREGI project, which amongst other approaches also invented translators between CIM and GLUE, and the efforts of the GridPP and NGS [15], where the specific information systems have been coupled based on LDAP, which was common to both infrastructures, or by using a common workflow portal [16]. Initiatives like the Grid Interoperability Now collaboration (GIN-CG) at OGF try to leverage the ongoing work on Grid interoperability and to find a common minimal attribute set.

In contrast to bilateral approaches, the work at hand focuses not primarily on the translation and integration between pairs of specific data schema and interfaces. It does neither discuss the theory behind the required data integration methods (as it may be found e.g. in [17] or [18]), but outlines general functionalities necessary for such and similar use cases and exemplarily presents an implementation. It also extends the interoperability discussion to cover the involved structural issues of a Grid's composition, organization and virtual organization, and addresses the requirement to provide monitoring data to specific virtual organizations which use multiple Grids in parallel.

[19] argued that Grid information services should not in general provide users with a consistent view of global state, but should focus only on efficient delivery of state information from a single source. Indeed, a single VO member always relies on a single access point at runtime. Nevertheless, in a distributed infrastructure always multiple different access points exist and can be chosen. If there is no consistent view of a VO's (local) state at all of these points, the Grid cannot be operated by a VO in a consistent way. Thus, we additionally argue, that in order to support the distributed character of a Grid, it must be possible to provide a VO-specific delivery of consistent state information at all available access points even when there are multiple service access points of a similar type. We resolve this by proposing multiple federated (respectively clustered) information services which use a single policy decision point which describes rules for consistent VO-specific partitioned states in section 4.

VOs, as well as VO-specific data provisioning from 'real' to 'virtual' organizations, using local or VO-specific state partitions, relies on a presumed categorical separation between the terms 'real' and 'virtual'. This separation basically circumscribes the idea to classify organizational entities in 'real' and 'virtual' segments, while a 'virtual' segment spans and aggregates specific entities of 'real' segments, often using a closed computer or communication system. Alluding to [20], one could see the problem of relaying specific monitoring data from the information service of a local resource provider as a 'real' organization to specific entities organized in a 'virtual' organization, as analogous to the problem of dynamically mediating or switching data in a rule-defined and segmented way, as this is done e.g. in scenarios like virtual memory [21] or virtual LAN switching [22] (which in contrast to our scenario are homogenous environments). Correspondingly, we used policies in order to influence the behaviour (as pointed out in [23]) of the common Grid monitoring system in a rule-defined way, making its Grid monitoring service VO-aware [24]. This allows a selective tailoring of provisioned monitoring information for given customers or VOs, based on the allocation of the monitored resources and services.

Such an approach reminds on policy-based configuration management based on group memberships as it was already introduced as early as e.g. in [25]. Similar methods have been applied to Grids in [26], where Policy Decision and Enforcement points are outlined, which authorize the provisioning of Grid data to the members of a user group based on records provided by identity providers. In contrast to these approaches, the work at hand takes into account the set of (virtual) services and (virtual) resources which the group uses, owns or has allocated, as well as changes which occur in that set. It also outlines the necessary functional modules to do so. Therefore, section 4 introduced two kinds of constraints, authorization and virtualization constraints, which may be gathered from different policy information bases.

7 Conclusion and Outlook

The paper discussed critical characteristics of Grids that use heterogenous monitoring services and presented requirements as well as functional modules which are necessary to realize an interoperable, distributed and dynamic VO-based Grid monitoring. Especially in Grid infrastructures, which are usually distributed on a large scale and often change dynamically, clearly defined decision and information points are necessary in order to provide a common understanding of the composition of the Grid and its VOs'. They constitute static points of reference where real and virtual organizations agree upon the configuration of the VO's dynamic computing clouds and enables a comprehensive mapping of status data from all 'real' resources and services at the resource providers to their 'virtual' counterparts in the VO's. This accordingly allows a monitoring of resources and services based on the dynamic virtual organizations they have been allocated to.

The following research questions still remain open and require further work:

- Which policy languages are appropriate to implement more sophisticated VO-composition rules in detail?
- How could these rules look like to assure an efficient operation of a broader spectrum of virtual resources and services ?
- How can the additional complexity be understood and handled which arises when multiple Policy Decision and Information Points are used for defining the mappings?
- How can conflicts between concurrent mappings be resolved?

Acknowledgements

The author wishes to thank the members of D-MON, GLUE 2.0 and OMII Europe, the (global) Grid communities and the Munich Network Management (MNM) Team for discussion, inspiration, helpful work and support. This work has been funded in parts by the D-Grid Initiative of the German Federal Ministry of Education and Research (FKZ 01IG07010A).

References

[1] Lipnack, J., Stamps, J.: Virtual Teams: Reaching Across Space, Time and Organizations with Technology. John Wiley (1997)

[2] Dreo Rodosek, G., Hegering, H.-G., Stiller, B.: Dynamic Virtual Organizations as Enablers for Managed Invisible Grids. In: Proceedings of the 2006 IEEE/IFIP Network Operations and Management Symposium (NOMS). (2006)

[3] Mowshowitz, A.: Social Dimensions of Office Automation. In: Yovitz (Ed.), M.: Advances in Computers, Vol. 25. (1986), p. 335 - 404

[4] Lucas, H.C., Barudi, J.: The role of information technology in organization design. In: Information Systems Working Papers. (1993)

[5] Foster, I., Kesselman, C., Tuecke, S.: The Anatomy of the Grid Enabling Scalable Virtual Organizations. In: International Journal of Supercomputer Applications. (2001)

[6] Neuroth, A., Kerzel, M., Gentzsch, W.: German Grid Initiative D -Grid. Universitätsverlag Göttingen (2007)

[7] Baur, T., Felde, N.g., Reiser, H.: Konzepte zum Monitoring im Kern D-Grid (engl.: Concepts for Core D-Grid Monitoring). D-Grid Report, FG2-6. (2007)

[8] Chaudhuri, S., Dayal, U.: An Overview of Data Warehousing and OLAP Technology. In: SIGMOD Record, Vol. 26, No. 1. (1997), pp. 65-74.

[9] Westerinen,A., Schnizlein,J., Strassner,J., Scherling, M.,Quinn, B.,Perry, J.,Herzog, S.,Huynh,A., Carlson,M.: Policy Terminology. IETF Internet Draft draft-ietf-policy-terminology-00.txt (2000)

[10] Lupu, E., Sloman, M.: Policy Driven Management for Distributed Systems In: IEEE Transactions on Software Engineering - Special Issue on Inconsistency management. (1999)

[11] Antonioletti, M., Atkinson, M.P., Baxter, R., Borley, A. , Chue Hong, N.P., Collins, B., Hardman, N., Hume, A., Knox, A., Jackson, M. , Krause, A., Laws, S., Magowan, J., Paton, N.W., Pearson, D., Sugden, T., Watson, P. and Westhead, M.: The Design and Implementation of Grid Database Services in OGSA-DAI. Concurrency and Computation: Practice and Experience Vol. 17, No. 2-4. (2005), p. 357-376.

[12] Andreozzi (Ed.), S., Burke, S., Ehm, F., Field, L., Galang, G., Konya, B., Litmaath, M., Millar, P., Navarro, JP.: GLUE Secification v. 2.0 (draft 33). Open Grid Forum Specification. (2008)

[13] Marzolla, M., Andreetto, P., Venturi, V., Ferraro, A., Memon, A.S., Memon, M.S., Twedell, B, M. , Riedel, M. , Mallmann, D. , Streit, A. , van de Berghe, S., Li, V., Snelling, D. , Stamou, K. , Shah, Z.A. ,Hedman, F.: Open Standards-based Interoperability of Job Submission and Management Interfaces across the Grid Middleware Platforms gLite and UNICORE. In: Proceedings of International Interoperability and Interoperation Workshop (IGIIW). (2007), p. 592-599.

[14] Fields, L., Schulz, M.: Grid Interoperability: The Interoperations Cookbook. In: International Conference on Computing in High Energy and Nuclear Physics (CHEP '07). Journal of Physics: Conference Series 119. IOP Publishing (2008)

[15] Jensen, J., Stewart, G., Viljoen, M., Wallom, D., Young, S.: Practical Grid Interoperability: GridPP and the National Grid Service. UK e-Science All Hands Conference. (2007)

[16] Kacsuk, P., Kiss, T., Sipos, G.: Solving the grid interoperability problem by P-GRADE portal at workflow level. In: Future Generation Computer Systems Vol. 24, No. 7. Elsevier (2008), p. 744-751.

[17] Ziegler, P., Dittrich, K.: Three Decades of Data Integration - All Problems Solved? In: First International IFIP Conference on Semantics of a Networked World (ICSNW). (2004), p. 3-12.

[18] Hegering, H.G., Abeck, S., Neumair, B.: Integrated Management of Networked Systems: Concepts, Architectures, and Their Operational Application Morgan Kaufmann (1999)

[19] Czajkowski, K., Fitzgerald, S., Foster, I., Kesselman, C.: Grid Information Services for Distributed Resource Sharing In: Proceedings of the 10 th IEEE Symp. On High Performance Distributed Computing. (2001)

[20] Moshowitz, A.: Virtual Organization: Toward a Theory of Societal Transformation Stimulated by Information Technology. Quorum Books. (2002)

[21] Bensoussan, A., Clingen, C. T.: The Multics Virtual Memory: Concepts and Design. Communications of the ACM Vol. 15, No. 5. (May 1972), p. 308-318

[22] IEEE-SA Standards Board: IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks. IEEE Std 802.1Q-1998. (1998)

[23] Sloman, M.: Policy Driven Management for Distributed Systems. In: Journal of Network and Systems Management, Vol. 2, No.4. Plenum Press (1994)

[24] Baur, T., Bel Haj Saad, S.: Virtualizing Resources: Customer Oriented Cross-Domain Monitoring for Service Grids. In: Proceedings of the 10th IFIP/IEEE Symposium on Integrated Management. (2007)

[25] Gruschke, B., Heilbronner, S., Wienold, N.: Managing Groups in Dynamic Networks. In: Proceedings of the 6th International IFIP/IEEE Symposium on Integrated Management (IM99). (1999).

[26] Hommel, W., Schiffers, M.: Benutzergesteuerter Datenschutz in Grids (engl.: User-controlled data privacy in Grids). 1. DFN-Forum Kommunikationstechnologien "Verteilte Systeme im Wissenschaftsbereich". (2008)