Harmonizing the Management of Virtual Organizations Despite Heterogeneous Grid Middleware – Assessment of Two Different Approaches*

Wolfgang Kirchler[1], Michael Schiffers[2,4], Dieter Kranzlmüller[2,3,4]

[1] Technische Universität München (TUM), Germany
[2] Ludwig-Maximilians-Universität München (LMU), Germany
[3] Leibniz Supercomputing Centre, Garching (LRZ), Germany
[4] Munich Network Management Team (MNM)

Abstract

Coordinated problem solving and secure resource sharing in dynamic ensembles of organizations is critically dependent on the concept of virtual organizations (VO). However, as grid systems continue to grow in scale, exhibit greater dynamics, and become more heterogeneous, managing VOs on a grid scale becomes an increasingly difficult challenge. This is not only caused by different VO-philosophies and different middleware technologies, but it is also due to various authentication and authorization schemes. In our work we aim at harmonizing the VO management within and between these heterogeneous setups. To better understand the necessary building blocks we have realized two alternative solutions. The first one is based on an integrated approach, the second one introduces an additional abstraction layer as a proxy between VOs and the grid middleware. Both concepts have their individual advantages and disadvantages. The results of the respective assessment indicates that a combined solution may be beneficial.

1 Introduction

Since its introduction the term “grid computing” has commonly been understood as coordinated problem solving and resource sharing in dynamic, multi-institutional virtual organizations. The concept of Virtual Organizations (VO) is thus central to grids. Intuitively, VOs consist of individuals and resources “owned” and provided to the VO by autonomous real organizations (RO) – the Resource Providers (RP) or Service Providers (SP) – under certain conditions. As VOs make extensive use of these resources and services, appropriate authentication and authorization mechanisms are required. Although slightly misleading, such mechanisms are commonly denoted by “VO management”.

*This work has partially been funded by the D-Grid Initiative of the German Federal Ministry of Education and Research (PT-IN grant FKZ 01AK810B) and in other parts by the European Network of Excellence for the Management of Internet Technologies and Complex Services (IST-EMANICS grant 26854).

As an analysis of the existing literature and a survey conducted in the German D-Grid communities [1] shows, “appropriateness” in this context refers to the necessity of coping with heterogeneous middleware technologies (Globus, UNICORE, gLite), with heterogeneous authentication and authorization policies (identity based, role based, attribute based), and with attribute sets from both VOs and Shibboleth federations [2].

In principle, there are several ways to overcome these difficulties. To better understand the necessary building blocks we have investigated the benefits and the drawbacks of two of them. In a recent D-Grid project “Interoperability and Integration of VO Management Technologies in the D-Grid (IVOM)” [2] we augmented the middleware by the ability to manage various certificate types and attribute profiles (the “integrated” approach). In a separate work [3] we provided a separate VO-layer exhibiting a proxy kind of functionality for abstracting from heterogeneity (the “abstraction” approach). In this paper we briefly report on both efforts and present a short assessment with respect to the above harmonization objective.

In the following section we render the problem more precisely before we look at related work in section 3. In sections 4 and 5 we describe and assess the two solutions in more detail before we conclude our contribution in section 6.

2 Problem Statement

Accessing grid resources requires the membership to a Virtual Organization (VO) since grid resources are assigned by the respective RP to VOs only and *not* to individuals. Consequently, users have to prove both their VO membership and their rights to access the resources. A typical procedure to achieve this is to present a valid certificate. Unfortunately, though, there is no commonly agreed-upon standard for authentication and authorization in grids, neither technically nor policy-wise. Rather, it becomes apparent that most grid projects differ significantly regarding their middleware functionality for VO management support in general and their security management especially.

A typical example is D-Grid which not only supports the Globus Toolkit 4.x (GT4) middleware, but also gLite and UNICORE. Further, the D-Grid authentication mechanisms are required to support both Shibboleth and X.509, while the authorization needs to be based on Globus grid-mapfiles and the VO Membership Service (VOMS) [4]. To make matters worse, the D-Grid authentication and authorization policies need to support both pre-Web Service (WS) and WS components. For instance, the Globus Toolkit (GT4) myProxy-service [5] relies on the pre-WS mechanisms while the VOMS Policy Decision Points (PDP) and Policy Information Points (PIP) and the GridShib service [6] assume the WS flavor (see Fig. 1(a)). A similar strategy is to be recognized for gLite. In UNICORE[1] a client authenticates itself at the UNICORE gateway which passes the client requests to the Network Job Supervisor (NJS) which then queries the UNICORE User Database (UUDB) for authorization decisions (see Fig. 1(b)).

[1] UNICORE 5

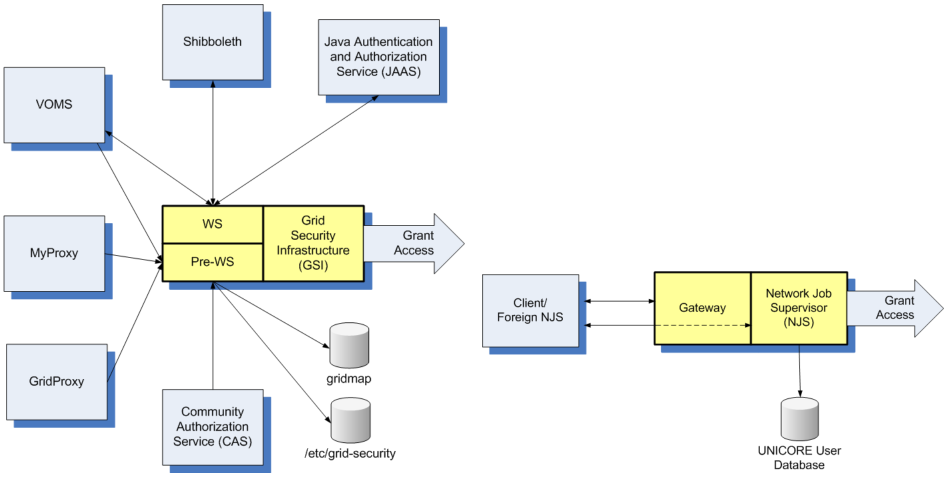
Shibboleth
Java Authentication and Authorization Service (JAAS)
VOMS
MyProxy
GridProxy
WS
Pre-WS
Grid Security Infrastructure (GSI)
Grant Access
gridmap
/etc/grid-security
Community Authorization Service (CAS)
Client/ Foreign NJS
Gateway
Network Job Supervisor (NJS)
Grant Access
UNICORE User Database

(a) Globus Toolkit 4

(b) UNICORE

Fig. 1: Authentication/Authorization Infrastructures (AAI) for Globus Toolkit 4 and UNICORE (adapted from [7])

Against this background, the "harmonization problem" addressed in this paper can be formulated as follows: *Given a multi-middleware grid infrastructure with multiple sources of identity. How can the authentication and authorization processes be unified?*

3 Related Work

There have been several efforts to solve parts of the harmonization problem. Their emphasis has mainly been on authentication and authorization in integrated X.509-oriented grid environments and Shibboleth federations. In [8] the results of an analysis with respect to the integration of Shibboleth-based and PKI-based VO management systems has been presented. It can briefly be summarized as follows:

GridShib [6] and myVocs [9] currently offer a broad set of solutions for a transparent grid and Shibboleth integration. Unfortunately they assume the Globus ecosystem. Although myVocs is restricted regarding both the attribute handling and the user/administration support, it is flexible enough to leverage the Shibboleth roles of Identity Providers and Service Providers for grids. However, both GridShib and myVocs do not support the required spectrum of middleware stacks.

VOMS [4] is a mature and stable VO management system developed as part of the gLite middleware. It has been used in production environments for several years and hence is the de-facto standard in PKI-based VO management. Furthermore, it is being actively enhanced with new features such as support for arbitrary attribute-value-pairs, which is an essential feature for any flexible VO management. However, it has to be admitted that VOMS itself does not offer the integration of Shibboleth-based campus attributes. Means would have to be found to combine VOMS with Shibboleth. VOMRS [10] offers a subset of the features of VOMS and can be used as a front-end to VOMS. Both VOMS and VOMRS do not support campus attributes yet.

Despite these promising efforts, it has to be noticed that there is currently no comprehensive solution to the harmonization problem available. This lack was the starting point for further investigating the ancillary conditions and necessary assumptions of a solution to this issue. We set up two different projects. In the D-Grid project "Interoperability and Integration of VO Management Technologies in the D-Grid (IVOM)" [2] we augmented every grid middleware and in a separate work [3] we defined a generic VO, provided a separate VO-proxy, and standardized the VO management interfaces. Both projects served as a feasibility study for a separate realization project.

4 The Integrated Approach

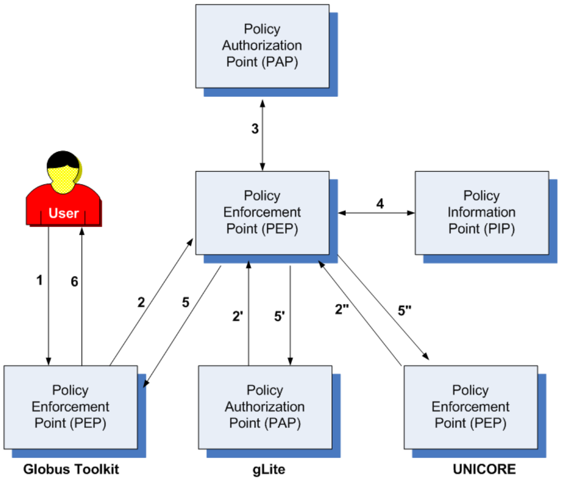
Policy Authorization Point (PAP)
3
User
Policy Enforcement Point (PEP)
4
Policy Information Point (PIP)
1
6
2
5
2'
5'
2"
5"
Policy Enforcement Point (PEP)
Policy Authorization Point (PAP)
Policy Enforcement Point (PEP)
Globus Toolkit
gLite
UNICORE

Fig. 2: Architecture for interoperable authorization (adopted from [11])

The main objective of the integrated approach is to unify the *VO management* across the GT4, the gLite, and the UNICORE middleware while at the same time providing mechanisms for an authentication and authorization scheme which not only accepts VO attributes but also campus attributes. In the D-Grid the former ones originate from the VOMS/VOMRS setups, the latter ones are provided by the Shibboleth Identity Providers. To overcome the difficulties mentioned before and their inherent scalability problems we proposed in [2] a policy based management (PBM) system with the following components *per middleware* (see also Fig. 2):

- Policy Enforcement Points (PEP) at the resources, where the authorization decisions are enforced,
- Policy Decision Points (PDP), where the authorization decisions are made based on defined policies,

- Policy Information Points (PIP) as repositories for decision information (e.g. the attributes of a user requesting access to a resource),
- Policy Authority Points (PAP), where the actual policy rules are defined.

As [2] show, implementing such an architecture across multiple grid middleware technologies can successfully be done using extensions to the Security Assertion Markup Language (SAML) protocol [12] as this turns out to be the common denominator. An example authorization process based on such an architecture is depicted in Fig. 2. It consists of the following steps:

1. A user requests access to a particular resource.
2. The PEP passes a respective request to the PDP.
3. The PDP gets access to the appropriate policy from the PAP.
4. The PDP accesses, if required, additional information like further attributes of the user, from the PIP.
5. Now the PDP can decide and returns a response to the PEP.
6. The PEP enforces the decision by either allowing or disallowing access to the resource.

As the protocol for each PEP/PDP is the same and standards based – regardless of the grid infrastructure underneath – the steps 2, 5, 2', 5', 2", and 5" in Fig. 2 are equivalent.

The major lessons we learned are:

- Depending on the deployment, all middleware systems in use have to be touched and modified according to the objectives. This can be relatively simple (as for GT4) or it can require a major effort (as for UNICORE). In all cases, though, it entails a permanent hook into the respective middleware problem and configuration management.
- The VOMS/VOMRS VO attributes specify the VO, the groups within the VO, the roles per group, and the capabilities per role. They may be encoded as either attribute certificates or as SAML assertions or both. The second case, however, requires a SAML awareness in VOMS and dedicated policy decision services since a SAML enabled VOMS does not know (key, value)-pairs and hence will not be able to distinguish between groups, roles, or capabilities.
- VO attributes are grid specific and need to be managed by a separate VO Management system irrespective of being short-lived or long-lived. Especially, portal-based grids require for authentication and authorization SAML and Shibboleth [2]. Unfortunately, however, current Shibboleth setups do not support multiple attribute authorities.
- As far as PDPs are concerned, the UNICORE implementation had to be developed from scratch [13]. For gLite to support SAML-encoded attributes the announced SAML enabled VOMS together with the gJaf authorization framework [14] is sufficient. GT4 provides with its interceptor mechanism a flexible way to interrogate PIPs and to aggregate PDP decisions using either permit-override semantics (GT 4.1) or deny-override semantics (GT 4.0). GT4 also allows for querying external PDPs using SAML callouts to authorization services.

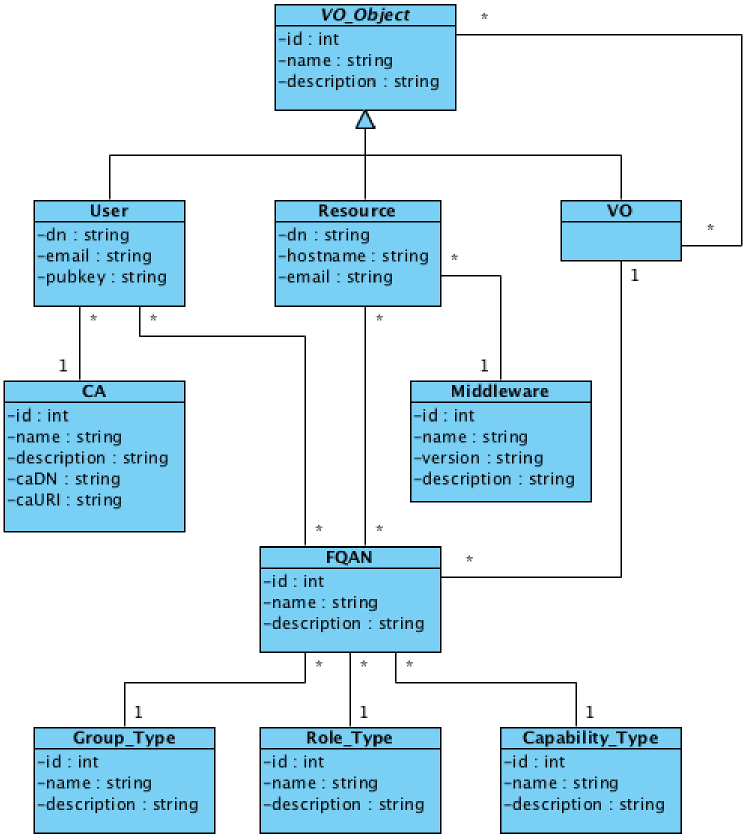

Fig. 3: VO layer data model (adopted from [3])

Apart from these technical issues it is necessary that the integrated approach requires the grid RP to belong to the respective Shibboleth federation(s) which implies the implementation of the corresponding Identity and Service Providers.

5 The Abstraction Approach

Unlike the integrated approach which mainly focuses on enhancing middleware components the abstraction approach tries to abstract with a separate VO-layer from both the middleware heterogeneity and the diverse VO management tools. For unifying the VO management *access* – as opposed to unifying the VO management itself – a generic VO structure is defined first, followed by a specification of a corresponding database scheme for persistently storing VO data, and finally the required control flows are determined. Fig. 3 depicts the VO-layer data model in more detail.

The generic VO structure leverages the work already presented in VOMS / VOMRS and covers groups, roles, and capabilities (see Fig. 3). In terms of this approach resources are considered VO members themselves which makes it easier to determine a user's access rights directly from his position in the VO structure. Because the VO structure is generic, it is flexible enough to incorporate various authentication mechanisms (with or without groups, roles, capabilities) and authorization schemes (with or without attributes). It should be noticed, however, that this approach does *not* implement a new authentication/authorization scheme. Rather, it acts like a proxy between user and middleware.

Compared to the integrated approach discussed before, this approach features an easier implementation, a greater flexibility in supporting heterogeneous VO management systems and grid middleware technologies, and a more comprehensive view on VOs – both isolated ones and overlapping ones – since it is not restricted to VO memberships of individuals only. It gains these advantages, however, at the expense of an additional overhead and a grid-wide (*not* VO-wide) central database.

6 Conclusion and Further Work

Both approaches achieve the basic unification requirement (up to a certain degree) despite the heterogeneity of the underlying building blocks. However, due to their conceptual differences, both have their individual advantages and disadvantages. While the integrated approach not only requires modifications of the grid middleware but also of policy decision and enforcement points, the abstraction solution creates additional overhead and seems to be less dependable – induced by the central database paradigm. Yet, the abstraction concept is simple to deploy and existing legacy solutions can easily be integrated. The integration concept in turn allows for a smoother integration of non-grid authorization schemes (like those deployed in Shibboleth federations) but it is less flexible regarding changes in underlying Authentication and Authorization Infrastructures (AAI).

In interesting question is now how to combine – if at all – these concepts. Basicaly, there are two promising ways to hybridize. A VOMS / VOMRS based approach would be based on the VO layer concept of the abstraction approach. The VOMS/VOMRS tables would require an extension to support resource descriptions and VOMS could then be deployed to handle all authorization requests directly. A Shibboleth based approach would also be based on the VO layer concept of the abstraction approach. As opposed to the previously described technique, VOMS stays stable. Rather, Shibboleth needs to be extended by a VO notion as expressed in the VO layer.

Experience shows, however, that any hybrid solution will be hard to push through standardization and will thus not likely to be realized short- to midterm. However, this needs to be explored in more depth in subsequent works.

References

1. Heike Neuroth, Martina Kerzel, and Wolfgang Gentzsch. *German Grid Initiative D -Grid*. Universitätsverlag Göttingen, 2007.
2. P. Gietz, C. Grimm, R. Gröper, S. Makedanz, H. Pfeiffenberger, M. Schiffers, and W. Ziegler. A Concept for Attribute–Based Authorization on D–Grid Resources. *Future Generation Computer Systems — The International Journal of Grid Computing: Theory, Methods and Application*, 25(3):275 – 280, March 2009.
3. W. Kirchler. Entwicklung einer einheitlichen Autorisierungs- und Authentifizierungsschnittstelle für heterogene Grids am Beispiel D–Grid (in German). Master's thesis, Technische Universität München, September 2008.
4. R. Alfieri, R. Cecchini, V. Ciaschini, Luca dell'Agnello, Á Frohner, K. Lőrentey, and F. Spataro. From gridmap-file to VOMS: Managing Authorization in a Grid Environment. *Future Gener. Comput. Syst.*, 21(4):549–558, 2005.
5. Jim Basney, Marty Humphrey, and Von Welch. The MyProxy Online Credential Repository: Research Articles. *Softw. Pract. Exper.*, 35(9):801–816, 2005.
6. Tom Barton, Jim Basney, Tim Freeman, Tom Scavo, Frank Siebenlist, Von Welch, Rachana Ananthakrishnan, Bill Baker, Monte Goode, and Kate Keahey. Identity Federation and Attribute-based Authorization through the Globus Toolkit, Shibboleth, GridShib, and MyProxy. In *Proceedings 5th Annual PKI Research and Development Workshop*, Gaithersburg, USA, 2006.
7. Tobias Dussa, Ursula Epting, Bartol Filipovic, Gerti Foest, Jürgen Glowka, Joachim Götze, Christian Grimm, Markus Hillenbrand, Christian Kohlschütter, Rudolf Lohner, Siegfried Makedanz, Paul Müller, Marcus Pattloch, Stefan Piger, Tobias Straub, and Jan Wiebelitz. Analyse von AA-Infrastrukturen in Grid-Middleware (in German). Report D-Grid Fachgebiet 3-4, March 2006.
8. Peter Gietz, Christian Grimm, Ralf Gröper, Martin Haase, Siegfried Makedanz, Hans Pfeiffenberger, and Michael Schiffers. Evaluation of International Shibboleth-Based VO Management Projects. Report of Work Package 1 of the D-Grid IVOM Project, May 2007.
9. Jill Gemmill and John-Paul Robinson. myVocs and GridShib: Integrated VO Management. Presentation at the Spring 2006 Internet2 Member Meeting, Arlington/USA, April 2006.
10. VOX Project Team. *VOMRS User Guide*. Fermi National Accelerator Laboratory, Chicago, USA, 2004.
11. Peter Gietz, Christian Grimm, Ralf Gröper, Martin Haase, Siegfried Makedanz, Hans Pfeiffenberger, and Michael Schiffers. A Concept for Authorization on D-Grid Resources. Report of Work Package 3 of the D-Grid IVOM Project, September 2007.
12. Nick Ragouzis, John Hughes, Rob Philpott, Eve Maler, Paul Madsen, and Tom Scavo. Security Assertion Markup Language (SAML) V2.0 Technical Overview, Committee Draft 01, 13 March 2007, March 2007.
13. Arash Faroughi, Roozbeh Faroughi, Philipp Wieder, and Wolfgang Ziegler. Attributes and VOs: Extending the UNICORE Authorisation Capabilities. In *Proceedings Euro-Par 2007, UNICORE Summit 2007*, volume 4854 of *Springer LNCS*, Rennes, France, August 2007.
14. Yuri Demchenko. gLite Java Authorisation Framework (gJAF) and Authorisation Policy Coordination. Proceedings MWSG Meeting EGEE06 Conference, Geneva, September 2006.