IVOM: Interoperabilität und Integration der VO-Management Technologien im D-Grid

Work Package 2: VO-Management Requirements from a Community Perspective

Authors:
Peter Gietz (DAASI International GmbH, Tübingen)
Martin Haase (DAASI International GmbH, Tübingen)
Hans Pfeiffenberger (AWI, Bremerhaven)
Michael Schiffers (Ludwig Maximilian University, Munich)

Version 0.6

Contents:

Abstract

The IVOM project [1] requires in work package 2 the specification of both functional and non-functional requirements for the VO-management in D-Grid from a community perspective. These requirements will be the driving forces behind the conceptual work in AP3. This paper focuses on the specification of these requirements.

Notational Conventions

This document complies with RFC 2119. Thus, the keywords “MUST”, “MUST NOT”, “REQUIRED”, “SHALL”, “SHALL NOT”, SHOULD”, “SHOULD NOT”, “RECOMMENDED”, “MAY”, and “OPTIONAL” in this document are to be interpreted as described in [RFC2119].

1 Introduction

Work on the VO-Management framework (see [VO]) has shown the different approaches to VO-Management in the D-Grid communities. Depending on the respective culture of scientific cooperation within the communities, significant differences in structure, complexitiy, duration, and size of scientific projects have been identified [IVOM]. These differences not only induce questions on the number of VOs, their size, and their mode of operation, but also questions on the membership to multiple VOs, on inter-VO cooperations, and on adequate authentication and authorization mechanisms.

Accordingly, the objectives of this work package are twofold:

- to specify a set of functional requirements for the integration of the communities’ VO management systems in Shibboleth-based federations. This includes a specification of user/manager perspectives and their respective role-to-capabilities mappings.
- to specify a set of accompanying non-functional requirements. This includes an estimation of quantitive metrics (size, number, duration) with respect to community VOs.

This analysis is based on interviews with the communities conducted during the first quarter of 2007. We are aware that – in the course of time – some of the requirements may change; some may become obsolete, while others have to be added. Hence, this document represents the *current snapshot* of the requirements driving the conceptual work in work package 3. In addition, it also represents the criteria for cooperating with those international projects which have been reviewed in work package 1 (see [IVOM1]) and which may contribute to the objectives of the IVOM project.

2 Methodology

This requirements specification adheres to the following methodology:

1. We first describe the current situation as the point of departure
2. We then characterize the target situation by specifying the
    a. Functional requirements for an authorization framework and the
    b. Non-functional requirements

    The requirements are derived using different views: a user’s view, a manager’s (or administrator’s) view, and a Resource Provider’s (RP) view. The requirements themselves are specified as use cases.

3 Description of the Current Situation

The current situation has been revealed during the work on [VO] as per October 2006. It can be summarized by the properties of *heterogeneity* and *small-scale.* The small-scaleness is depicted in the following table while a short summary of the heterogeneous authorization and authentication methods is given thereafter.

The work on [VO] has resulted in the following information as per October 2006 and this work in a preliminary additional assessment per June 2007:

Community	**Number of VOs (Status per 10/2006)**	**06/2007**	**Number of VOs per Community (Forecast as per 10/2006)**	**06/2007**
Astro	1		2	
HEP	> 10		> 10	
C3	0	1	-	O(10)
MediGrid	0		> 12	
InGrid	0		> 1	
TextGrid	0		> 1	
WISENT	unknown		unknown	

Although the sizes of these VOs differ remarkably, a realistic size range could not be specified. The same holds for estimating a VO's duration.

Depending on their history and their provenance, the communities use different authorization and authentication mechanisms:

- Some communities favour the VO Membership Service (VOMS) known from the EGEE project.
  Example: HEPGrid
- Other communities extend this approach by a registration component VO Management Registration Service (VOMRS).
  Example: AstroGrid-D
- Others use (or plan) Shibboleth infrastructures.
  Examples: MediGrid, InGrid, TextGrid and C3

For authentication purposes most communities rely on X.509 certificates (e.g., AstroGrid-D, HEPGrid, and MediGrid).

For any further information we refer to [VO].

4 Virtual Organizations

The following requirements are based on an understanding of Virtual Organizations which has been outlined in [VO] and which is repeated here for convenience.

The concept of Virtual Organizations is central to Grids and hence prominently anchored in the OGSA specification [OGSA]. VOs consist of individuals and/or technical resources or services

„owned“ by autonomous real organizations (legal entities) and “provided” (in a broader sense) by these entities to the VO with the objective to cooperatively achieve a common goal.

Although central, however, the term „Virtual Organization“ has not been precisely defined yet and the D-Grid communities may have a different understanding of what a VO is (see [VO]). In the context of this project we define a VO as a

> *permanent or temporary collection of geographically distributed individuals, groups of individuals, organizational units or entire organizations sharing parts of their physical or logical resources and services, their knowledge, their skills, and their information in such a way that the commonly defined goals may be achieved.*

This definition is mainly geared to the definitions proposed by the *TrustCom*-project [trustcom], earlier definitions by Foster, Kesselman, Tuecke [anatomy], and the definition presented in the Open Grid Services Architecture Glossary of Terms (Version 1.5) [glossary].

This definition is broad and admits -- amongst others – the following characteristics:

- a VO contains at least one member
- a VO may contain a complete real organization
- a VO may be a sub-VO of another VO which itself is then called a super-VO
- VOs may build VO-chains
- A VO does not “own” real resources but may administer them
- A VO is not a legal entity

5 Requirements Specification

The target system will be characterized by overcoming the restrictions posed by the heterogeneity and the small-scaleness of the current approaches. This chapter summarizes the functional and non-functional requirements for the integration of heterogeneous VO Management Systems and Shibboleth-based federations in large.

The requirement specification follows the standard approach by identifying the roles (actors) and the use cases they are involved in. However, before exploring the requirements in more detail, we discuss for some communities their specific views on VO-management issues. Please note that we elicit in the text the respective community specific requirements (being underlined and by a corresponding identifier which we will refer to when describing the uses cases.) which will be aggregated into generic uses cases thereafter.

5.1 Community Views on VO-Management

5.1.1 Earth Sciences

The following descriptions were provided by Hans Pfeiffenberger from the Alfred Wegener Institute in Bremerhaven.

In the following specification of requirements, the phrase “their VO” or similar will be used, frequently, as shorthand expression for any kind of groupings – as they exist per april 2007 – of people, working together towards a common objective, frequently using some common tools or resources. The definition of those groupings may seem clear to those involved, but in most cases it is in fact loose or at least not strictly laid down in print. Each use of “VO”, here, must therefore be treated as tentative or subject to later introspection and revision. In most cases however, these VOs describe the members or partners of research projects in any stage of the funding lifecycle, even in the consortium building stage (and groupings or specific parts of projects).

In contrast, the term community is used in its customary meaning to denote a fuzzy group of people, not a enumerable one. In our understanding a community is constituted by all those who work on a more or less specific scientific subject ("climate research") or in a discipline ("biogeochemistry"). By this definition, a community cannot be restricted to the users of a limited infrastructure or resource (there is a HEP, but not a LHC community).

Requirements are formulated in narrative form for three examples. The second and third cases expand on the first and discuss additional or specific elements. The examples C3Grid, EIFEX and IPY depict a progression in the dimensions of international and multidisciplinary composition and (current) size of VOs.

5.1.1.1 C3Grid

The project "Collaborative Climate Community Data and Processing Grid" (C3Grid), - which can be regarded to be a VO - comprises today, upon closer examination, an assembly of a small number (~30) of persons from at least three communities: Grid-developers, operators of computing and data centres and power users **(ES0)** from the earth sciences. The result of this project is expected to evolve into a permanent infrastructure, serving a number of projects (future VOs). Its userbase is expected to grow by about 300 within the duration of the project – most of them normal users from the climate community in Germany and possibly a small number of other nations, by then.

The typical requirements in terms of access control which need to be satisfied within the Grid can be inferred from the general concerns typical at least for some disciplines in the Earth Sciences, as expressed in the second example ("EIFEX"). Some more specific conditions, imposed by some providers of data, stand out, here:

The European Centre for Medium-Range Weather Forecasts (ECMWF) provides some large, important datasets (ERA40 "reanalyses") considered of economic value, but free of charge for purely scientific use **(ES1)**. In practice, mirror copies of the data are held, e.g., by the Deutsches Klima Rechenzentrum (DKRZ). This in turn will ask today for individual users' physical signature, confirming the scientific use **(ES2)**. It remains to be resolved, which other kind of assertion by whom would be acceptable to ECMWF in the long term, e.g., whether a VO representative's signature would be accepted. As of today, the prospect of this happening is not considered imminent. Rather, to execute a job in which the ECMWF data are involved, a combination of privileges – the personal, due to the signature, and the membership in a project or VO – will need to be invoked.

This immediately leads to a user-side requirement: It was expressed early in the definition phase of C3Grid, by the representative of the climate community, that their main expectation of a Grid was not a compute grid but to get unencumbered access to all data they need and whose accessibility had been agreed upon **(ES4)**, especially in the case of data from their peers **(ES5)** (as opposed to organizations like ECMWF). This characteristic should be present – without additional burden – worldwide **(ES6)**, of course, since the climate community is not a German one, but international. Buried between "unencumbered" and "agreed to" is the implicit requirement that it must be extremely easy for the partner supplying data, to add or modify an access rule for a specific person or its collaboration, its VO **(ES7)** or whatever is – to both sides – a natural way to express the relationship between themselves and in relation to their community or project(s).

Currently, the C3Grid manages itself through the local coordinators **(ES8)** of the project partners, who may declare local persons members (or revoke this status) **(ES9)** – e.g., as they are hired or leave the partners' institution - , by means of a telephone call or email to the overall coordinator (or rather, in each case, de facto the local deputies of coordinators). There is, currently, no written, formal description of the conditions of membership or description of a workflow to establish membership. This web of trust – or at least: the web of well known contacts – has been established during the phase of preparing the C3Grid project proposal (and before) and continues to grow, especially since new members may be present at project meetings. Technically, identities and memberships are mainly established through email **(ES10)**.

5.1.1.2 EIFEX

The European Iron Fertilization Experiment (EIFEX) culminated in an expedition of the research vessel Polarstern, in January 2004, when the ship carried 54 scientists and students from 14 institutes and 3 companies, from 7 European countries and South Africa into the South Atlantic. The communities, more precisely, disciplines involved, were oceanography, biology, chemistry – or to use the modern term, denoting a new discipline: biogeochemistry. This project/expedition was a follow-up to EISENEX, in 2000, which was of similar, but not identical composition.

This dimension and complexity is typical, today of projects in the Earth Systems Science. The VO-forming and –managing processes, as well the privileged access to a specific dataset (satellite coverage imaging chlorophyll at the ocean surface), provide patterns very similar to the C3Grid example – with a stronger international and interdisciplinary dimension.

In terms of required access restrictions, both EISENEX and EIFEX asked for a website accessible only to named members of the projects **(ES11)**, where data were to be exchanged. Interviews with biologists and geologists during the projects' timeframe revealed a specific mindset towards data sharing: It was generally recognized, that sharing of data would be beneficial to all – except to themselves, in case of their "own" data. De facto, a low degree of respect for the intellectual property – or rather: authorship – of datasets was expected. Consequently, a strict control over which person or group is given access – at least prior to the first publications of the data collector – needs to be implemented, if data are to be shared at all. A typical scheme is that the author, or creator, of a dataset reserves its use to himself for – typically – at least one year. Then, it may be made available to his group (local or VO) **(ES12)**. Finally, after 3 or more years, it may or may not be made available to all. (But even then, there is at least one technically implemented example outside EIFEX, not on an anonymous basis.)

This kind of schemes – possibly appearing paranoid to some – may soften or vanish if and when proper publication and citation of datasets become a (cultural) norm in science. This may last a while – and, in case it happens, lead to the need of attaching authorship, access restriction and licensing information to and carrying those with the datasets.

5.1.1.3 IPY

The International Polar Year 2007-2008 (IPY) is a "voluntary" collaboration of organisations and projects with a number of objectives related to the topic of global change, specifically to collect a data snapshot of the polar regions. It involves 50.000 participants from 60 nations. The program is organized in 170 clusters of projects (!), where (almost) only the individual projects themselves are funded, mostly on a national basis. Many, if not most, clusters are interdisciplinary and international, as in the EIFEX example, above. Some groups of clusters can be aggregated, intellectually, to a theme (e.g., "sea level rise"). This structure of IPY is visualized by Figure 1.

The (semi-)central steering **(ES13)** that exists, is executed by a "Joint Committee", some international sub-committees, national IPY committees and an IPY program office. The term voluntary refers to the fact that, since there is (almost) no central or common funding, none of the structures can be enforced. This also means, that all resources have to be supplied either by the funded projects or by the participants' institutions, at their discretion and governed by their terms **(ES14)**.

Due to the peculiar, almost fractal financing structure, one might expect any resource to be subject to limits or regulations set at a national level (data about areas of national economic interest, personal data, export restrictions for computer usage, &c.). Luckily, with the possible exception of personal data, all IPY participants have agreed to share data "freely and timely". It remains to be seen, however, if the timely-ness will be interpreted in terms of days or weeks – as anticipated in policy – or 1 to 3 years, as in the EIFEX example.

The ECMWF reanalysis dataset (or at least, the 2007-8 part thereof), mentioned under C3Grid, is offered to the IPY as well. It is not yet clear, how the limitation of "scientific use(rs) only" **(ES15)**

could be implemented – especially, if this dataset is to become part of the persistent data legacy, which is to be available to "all".

It is expected that the virtual organization of the IPY will survive the year 2008 **(ES16)** and develop into a more closely knit polar community, as a further kind of legacy of IPY. Most probably, this community will fall back onto typical patterns of access restriction (as in EIFEX). Therefore, a technical underpinning through one or a number of VO management systems would be desirable. Its implementation would clearly begin from the top (the current IPY program office). The most important and central function of the VO management system for the polar community VO would probably be a Who-is-Who **(ES17)** and general communication system **(ES18)**.

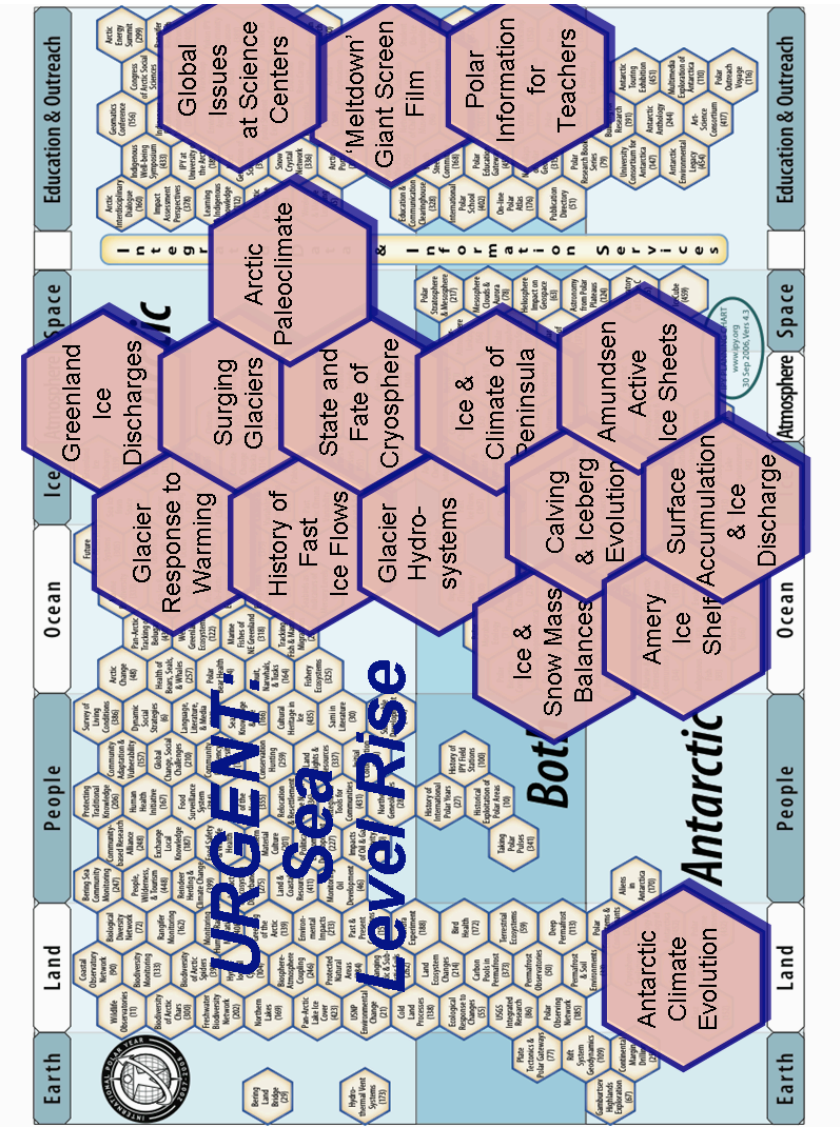

Figure 1: Organizational Diagram of the IPY [IPY]

5.1.1.4 Conclusion

It is the people, not the resources, which constitute a VO. Typical VOs in this domain would be scientific programmes, projects, communities, informal groups etc. VOs as we know them in Earth Science use resources and aggregate information, but they mostly do not maintain them and eventually do not own them:

- In most cases, mainly with temporary VOs, resources are being provided and maintained by the institutions **(ES19)** whose members take part in the VO.

- In most cases Vos have a specific mission or objectives and a limited lifetime. Therefore ownership of material (i.e. data) created in the VO eventually (if not from the beginning) will be transferred to the member institutions **(ES20)**.

- Third party resources – like the use of supercomputing facilities at regional or national centres – are either paid for by the members' organizations or are part of a grant **(ES21)** for members' (!!) projects.

  Typically, Vos cannot be granted resources because they are not legal entities. Rather, grants go to individual members' institutions which may – via a PI – arrange for the permission to use these resources for all or part of the VO membership. Thus, VO management systems need to

  - support mechanisms (or best: one simple mechanism) to express, authoritatively, information to implement restrictions on access to data, from the person level, through (sub-)VO membership **(ES22)** possibly to a global community. The information needed must be portable **(ES23)** with data and to every Grid node.

  - support the easy and swift formation of interdisciplinary, international collaborations **(ES24)**, from the phase of assembling a small group ("birds of a feather") to an operative, funded project and clusters of projects (programs). It may be necessary to express the line of work ("non-commercial").

  - support, for each VO, self-service, self-structuring, division into sub-Vos and aggregation into clusters **(ES25)**, thus enabling organic, yet structured growth of Vos, supporting a growing number of applications used and differentiation of rights expressed.

  - be interoperable with other VO management systems **(ES26)** around the world, encompassing systems used in many, if not all scientific communities.

5.1.2 TextGrid

The following descriptions were provided by Martin Haase and Peter Gietz from DAASI International in Tübingen. It should be noted that we only mark those requirements which are new compared to the requirements in chapter "Earth Sciences".

Not all resources in TextGrid (data resources for the time being, in a more advanced stage of the project service resources as well) are freely available, which is why an infrastructure for authentication and authorization is needed.
We assume different core roles a user can adopt. The *project leader* defines a project, assigns resources and users to it. Project leaders are privileged members such that they can register new users in TextGrid. The *implementor* is responsible for a certain task the project leader assigns to her, e.g. transcription of a manuscript. There may also be *casual users* that are granted read-only access to a project's resources. Each of these roles inherits the rights of the subordinate ones. Aside from these, there is also a *technical administrator* who defines e.g. TEI-schemas or XML transformation rules as needed by the project leader. The *programmer* is an external role; he wants to integrate his own tools into the TextGrid. A user can adopt many roles, even in different projects, and a certain role can be shared by many users.
Depending on the ease of how VOs can be established and managed, single projects preferably will be implemented as separate VOs. In this case there will be many small VOs, with a few superior VOs, such as the VO of all edition philologists. If this is not the case, there can be also many projects within one VO without sub-VOs.
Since TextGrid strives to attract as many users as possible and since the distribution of user certificates in the target community is to be assessed as being rather low, neither a central user administration nor a distributed Public Key Infrastructure (PKI) **(TG1)** can be considered. The

DFN-AAI, which is currently being developed, is a way to overcome these problems. It is a federation based on a Shibboleth infrastructure that allows institutions to provide their locally administrated users with resources of the entire federation in a controlled manner.
With Shibboleth, a Single Sign On (SSO) **(TG2)** solution can be implemented: Assertions (that conform to the standard SAML) about authentication status and authorization information (users' attributes) can be sent in between the individual services via a secure connection. Thus, such assertions partly fulfill the function of Kerberos tickets. While Shibboleth was originally designed for communication between web browsers and servers (where HTTP-Redirect takes a special role), enhancements were developed that do not presuppose the classical web context anymore : within the projects GridShib for Grid Computing and Lionshare for Peer2Peer-Networks, as well as within the SAML standardization and profiling for non-web scenarios.
The TextGrid architecture specifies a Service Oriented Architecture (SOA) **(TG3)** that provides individual services as modules for text analyzing, processing and publication. The modules are contacted via a user environment, a workflow editor or individual interfaces. The services themselves, in turn, contact Middleware services that provide data resources from the Grid.
The basic requirement of an AAI is that it can be incorporated into this architecture. For achieving this, the following is needed (see Figure 2):

- temporary: a central IdP for all TextGrid users. In later phases such an IdP could be used as a guest server for researchers that don't have an account at an IdP within the federation, like it is implemented in the Switch-AAI.
- at least at a later stage: a Germany-wide federation (DFN-AAI) **(TG4)** that allows access via the TextGrid users' individual home organizations by using a WAYF service.
- Authentication and authorization takes place via SAML assertions **(TG5).** These are:
    - created at the (Textgrid or Home) IdP
    - returned to the AAI component of the TextGrid user environment and are passed on to the middleware (by the AAI component and all further services like Workflow-Enactor, Tokenizer etc.)
    - received by a Policy Engine within the Middleware and are processed preferably by using GridShib SAML Tools and/or the OpenSAML library.
- Depending on the user's rights, which are encoded in the SAML assertion, the Middleware makes an access decision and enforces it.
- Files or faults are being returned

We are aware that there is an issue with keeping the Shibboleth session for SSO. The SAML profiles Shibboleth implements are not designed for web services and the HTTP session model doesn't apply to SOAP-based protocols. There is a couple of solutions to this:

- Use the myProxy-Approach and not send SAML after authentication but already a X.509-SLC or proxy.
- Develop ShibForSoap, an already seen desideratum, for which no-one found effort to do it yet (see https://spaces.internet2.edu/display/SHIB/WebServices). Within the OpenLiberty project similar attempts will be made to implement a Java library and reference implementation of the Liberty Alliance standard ID-WSF 2.0 WSC, that might produce outcome useful for this.
- Use a centralized RBAC component in the middleware that keeps track of the sessions, which could be named after the SAML assertion ID. If such a Session does not exist already, the SAML assertion is being parsed and the Session data (combined with pointers to users and roles) are being stored into the RBAC system. If it does exist, those data are used for access decisions.

Further requirements:

- Until the DFN-AAI is operational, TextGrid users have to get registered in the TextGrid IdP. The corresponding workflow has not been defined yet and is currently handled by informal communication (e-mail, telephone) and by manually entering the data into a directory service (LDAP) **(TG6)**
- The policy engine has to be implemented. PERMIS is possibly suitable for achieving this, otherwise a self-developed RBAC implementation will be adapted accordingly.

- The policies have to be administrated: Who determines which users are allowed to access which resources? With RBAC, there is a two-fold distinction:
  - The individual IdP administrators assign roles to users **(TG7)**. In a multi-layered model, some attributes are received from the home IdP but VO-specific attributes (role-assignments) from a VO-management system. It should be considered if such a VO-MS can be implemented using a Shib IdP, like it has been done in the IAM Suite of the MAMS project and in MyVocs. Similar work seems to be on the way in the U.S.
    SP admins assign resources to roles **(TG8)**. However, this could also be carried out centrally via a corresponding RBAC-system.
- The choice of the attributes that are to be included in the SAML assertions is still open/unsettled. Do "all" attributes that are specified in the attribute release policy go from the IdP to the AAI component of the user environment, which then passes them on to a service that in turn uses them for authorization at the Middleware – or should the IdP have every attribute

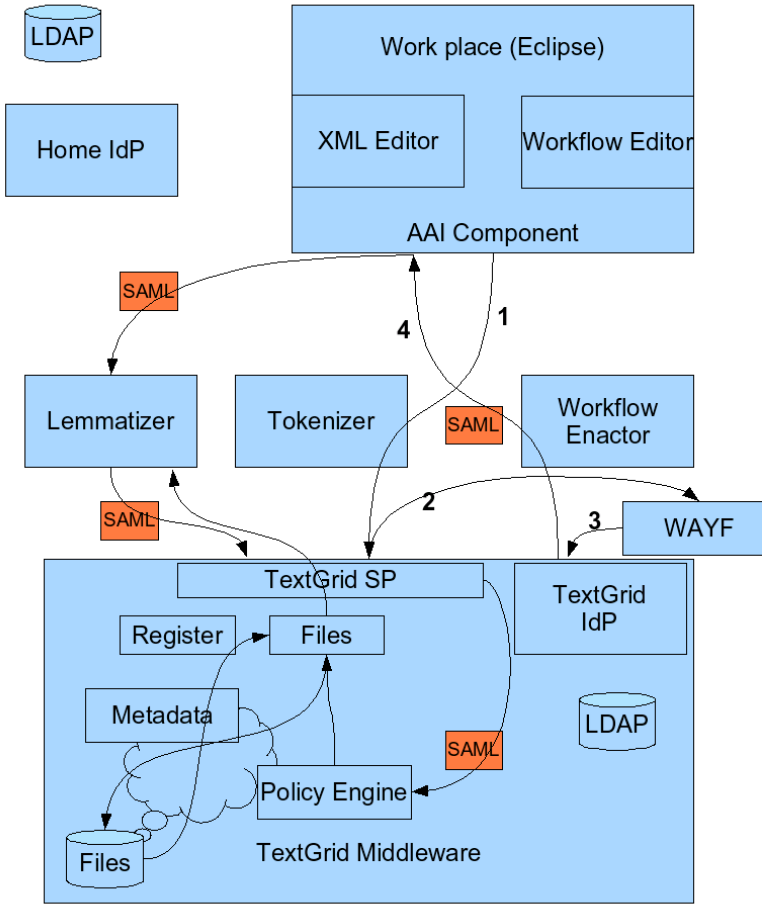

Figure 2: Integration of SAML and TextGrid

forwarding authorized by the individual user (e.g. via Autograph) – or only present the attributes to be released to the user for confirmation as a whole (e.g. via SWITCH's ARP-Viewer)? Attribute filtering can also be done on the SP-side.

- Web services **(TG9)** of the service layer have to be able to pass SAML messages through to the Middleware
- The Middleware's file services have to talk directly to the policy engine
- The SAML ticket has to be transmitted as parameter within the SOAP communication secured

by WS-Security.

Questions relevant to VO-Management

- It has to be cleared if either all TextGrid users are to be included in one VO or split up in three separate Vos (edition philologists, lexicographers, corpus linguists) – or in even more Vos, e.g. for TextGrid II (Egyptologists, Indologists etc.)
- The number of the Vos is independent from the question whether only *one* TextGrid IdP is used, or *several* Home IdPs (DFN-AAI)

Further use cases:

- A user only wants to work locally in the Eclipse-RCP/Workbench, i.e. without authentication. She should be able to use the streaming tools as long as those do not need middleware access, e.g. the Tokenizer. Read and write access is only possible locally.
- The "sitting-in-the-train" use case: working with the RCP without internet connectivity. Mechanisms for authorized replication and synchronization of data **(TG11)** to/from the local computer are needed.

5.1.3 InGrid

The following descriptions were provided by the InGrid community as a response to a review request against [VO]. Thus, the paragraphs here are fragments only. It should be noted that we only mark those requirements which are new compared to the requirements in the previous chapters.

InGrid's focus is on dynamically creating and terminating Vos **(IG1)**.

A typical InGrid scenario looks as follows: An Engineering Office receives the order to optimize a mold. The participating persons/roles are: the project manager of the Engineering Office, who creates the VO. He needs resources (hardware, software/licenses). Hence, there are the roles of a Hardware-RP and a Software RP. In addition, there are typically four user roles **(IG2)**:

- the role which accomplishes the simulation and which needs access to both input and output data. This role also needs the software for optimization and the software for casting simulation.
- the role of the Software Provider and the expert for the casting simulation, who, however, has no access to the optimization software. This role will only get the data relevant for its task.
- the role of the Software RP and expert for the optimization, who, however, has no license for the casting simulation.
- the role of the client of the Engineering Office, who only wants to analyze the results. He needs access to both input and output data and needs visualization software as well.

In some cases there may be a request for an additional Service Provider running e.g. a materials data base and for an Archive Service Provider.
There are also cases where information for billing processes is needed: the cost for licensed software may be different for academic or commercial use. Although this is regulated uniformly over the complete project, parts of the project may be more scientific oriented resulting in academic data/software usage only, including an appropriate pricing.
Also, the same user can accomplish one simulation within the same cooperation project for commercial purposes, the next for a scientific publication. How is this going to mapped to the VO-managemnt?

5.2 Roles and Actors

The roles and actors can be classified into roles stemming from technical considerations and those being derived from the community considerations above. Note that some roles may exhibit different "flavours" (e.g. a user may act as an engineer or as a scientist; an SP may provide simulation services or optimization services). However, this differentiation is of minor concern for this document.

The following roles may be identified from these sources:

- the `VO-Initiator` creates a VO, usually in his role as Principal Investigator (PI)
- the `VO-Manager` is in charge of all organizational tasks as far as Vos are concerned
- the `VO-Administrator` performs all administrative tasks as far as Vos are concerned
- the Vos have `members`
- `applicants` apply for membership
- resources are provided `by Resource Providers` (RP)

For more information regarding these roles see [VO].

5.3 Use Cases

The subsequent uses cases and their sub-cases including their attributes and constraints are abstractions from the descriptions above and, hence, MUST be considered. Some of these uses cases, however, may not be applicable for all communities. Chapter 5.4 contains a respective distribution. For cross reference purposes we list the requirements from above underneath the use cases. The order of the use cases is random and does not imply any importance.

It should also be noted that the use cases below specify *requirements* and *not* realizations. Thus, the specified procedures (primary or secondary) do *not* anticipate any implementation models. For example: specifying a member to authenticate herself as a member of a VO (use case 001) does not map directly in a GT4 scenario as there is no notion of VO, in a VOMS/VOMRS scenario, however, there is one.

The use cases are based on the following constraints or non-functional requirements:

1. The access policy MUST be able to handle the following aspects:
    a. Time (access allowed after, access not allowed before, …)
    b. Identity (access only for, access not for, …)
    c. Grouping (access for individuals, groups, organizations)
    d. Geography (national, international)
    e. Granularity (full access, partial access)
    f. Frequency (no restriction, onetime access, regular access, …)
    g. Cost (free of charge, flat fee, individual fee, degree of creditworthiness,…)
2. The following technical requirements MUST be considered:
    a. The solution SHOULD NOT rely on a specific Grid-middleware.
    b. The solution SHOULD NOT rely on specific software or software versions.
    c. Both the implementation and the deployment model MUST provide a migration path from existing community solutions to the envisioned one and from current versions to future versions.
    d. The solution MUST support a Service Oriented Architecture (SOA).

e. The User Interface MUST support web portal and command line

f. The solution MUST be policy-based.

3. The solution MUST comply to the DGI VO framework [VO].

5.3.1 Use Case 001: A VO-member wants to access a resource

5.3.1.1 Subcase: Regular resource access

Use Case ID	001_a
Cross reference to community requirements above	ES: 0,1,2,3,4,5,6,11,15,21 TG: 2,11 IG:
Short description	A member wants to access a resource provided by an RP.
Involved roles	`member, RP,`
Prerequisites	`Member` is a member of a VO and `RP` has "contributed" the resource to the VO.
Primary procedure	1. `Member` authenticates herself as member of the VO 2. RP checks member's authorization to access resource 3. RP grants access
Secondary procedures	1. `Member` may not be authorized 2. `RP` may not grant access 3. Resource may be unavailable
Non-functional requirements	• the decision process should not take longer than 1 second

5.3.1.2 Subcase: Access to a chargeable resource

Use Case ID	001_b
Cross reference to community requirements above	ES: 0,1,2,6,11,21 TG: IG:
Short description	A member wants to access a resource provided by an RP which is NOT free of charge.
Involved roles	`member, RP,`
Prerequisites	`member` is a member of a VO and the `RP` has "contributed" the resource to the VO. The `RP` has setup a price schema and accounting procedures.
Primary procedure	1. `member` authenticates herself as member of the VO 2. `RP` checks member's authorization to access resource 3. `RP` receives bill-to data 4. `RP` grants access

Secondary procedures	1. member may not be authorized 2. member may not be creditworthy 3. member may not accept proposal 4. RP may not grant access
Non-functional requirements	•

5.3.1.3 Subcase: Member belonging to a specific group wants to access a resource

Use Case ID	001_c
Cross reference to community requirements above	ES: 0,1,2,3,4,5,6,11,15,21 TG: 2,11 IG:
Short description	A member belonging to a specific group (e.g., scientist) wants to access a resource provided by an RP which is free of charge for group usage.
Involved roles	`Member, RP,`
Prerequisites	The member is a member of a VO and the RP has "contributed" the resource to the VO. The group is established and the member belongs to the group.
Primary procedure	1. `member` authenticates herself as member of the VO 2. `member` authenticates herself as belonging to group 3. `RP` checks member and group authorization to access resource 4. `RP` grants access
Secondary procedures	1. member may not be authenticated 2. member may not be authorized 3. member may not belong to group 4. group may not be authorized 5. RP may not grant access 6. member authorization may supersede group authorization
Non-functional requirements	• the decision process should not take longer than 1 second

5.3.2 Use Case 002: Member management

5.3.2.1 Subcase: Administrator or privileged user adds a new member

Use Case ID	002_a
Cross reference to community requirements above	ES: 7,8,9,10,16,20,22,24,25 TG: 6 IG: 1,2
Short description	A new member is added to the VO.
Involved roles	`applicant, VO-Administrator or privileged user, member`

Prerequisites	The VO has been established.
Primary procedure	1. applicant authenticates herself 2. admin creates new default membership based on given certificates 3. admin assigns roles to new member 4. new member is granted membership to requested VO 5. new member is granted membership to all super-Vos
Secondary procedures	1. applicant may not be authenticated 2. applicant may not be admitted (e.g., because of history or black lists) 3. super-Vos may refuse membership 4. membership request may not be accepted if remaining VO lifetime is below a given threshold 5. new member may be guest or dummy
Non-functional requirements	• membership is effective immediately

5.3.2.2 Subcase: Administrator or privileged user removes member

Use Case ID	002_b
Cross reference to community requirements above	ES: 7,8,9,10,22,24,25 TG: 6 IG: 1
Short description	A member is removed from a VO either regularly (e.g., contract termination) or irregularly (e.g., misbehaviour).
Involved roles	`Member, VO-Administrator or privileged user`
Prerequisites	Member is member of the VO.
Primary procedure	1. admin terminates all member activities 2. admin detaches member from roles 3. admin marks member as “already-been-there” 4. Record membership data for further processing (e.g., auditing, accounting)
Secondary procedures	1. person may be eligible for later renewal of membership 2. person may be eligible for keeping membership to super- and sub-Vos 3. sub-Vos and super-VO may refuse membership termination 4. member to remove may be guest or dummy
Non-functional requirements	• membership termination is effective immediately

5.3.3 Use Case 003: Member moves

Use Case ID	003
Cross reference to community	ES: 23 TG: 6

requirements above	IG:
Short description	A VO-member moves (as person) from organization A to organization B.
Involved roles	`Member, VO-Administrator,`
Prerequisites	member is member of VO.
Primary procedure	1. member informs administrator 2. admin changes member's records
Secondary procedures	1. new organization may not be a trusted site 2. new organization may not participate in D-Grid 3. new organization may not allow member to play role as hitherto
Non-functional requirements	• update is effective immediately • <<still open>>

5.3.4 Use Case 004: Role management

5.3.4.1 Subcase: Administrator assign role to member

Use Case ID	004_a
Cross reference to community requirements above	ES: 0 TG: 7,8,10 IG: 2
Short description	The VO-member is assigned a role.
Involved roles	`member, VO-Administrator,`
Prerequisites	Member is a member of a VO.
Primary procedure	1. admin checks authorization for new role 2. admin assigns new role to member
Secondary procedures	1. member may not be authorized for new role 2. group specification may collide with new role 3. member may be guest or dummy
Non-functional requirements	• role assignment is effective immediately

5.3.4.2 Subcase: Administrator removes a role from a member

Use Case ID	004_b
Cross reference to community requirements above	ES: 0 TG: 7,8,10 IG: 2

Short description	A role is detached from a VO-member.
Involved roles	`member, VO-Administrator,`
Prerequisites	Member is a member of a VO, member has assigned a specific role.
Primary procedure	1. admin blocks all member activities 2. admin detaches role from member
Secondary procedures	1. old role may be "orphaned" after assignment 2. member may not be "on role" 3. member may be guest or dummy
Non-functional requirements	• role assignment is effective immediately

5.3.4.3 Subcase: Administrator adds new role to VO

Use Case ID	004_c
Cross reference to community requirements above	ES: TG: IG:
Short description	The VO-Administrator adds a new role to the VO.
Involved roles	`member, VO-Administrator, VO-Manager`
Prerequisites	The VO exists.
Primary procedure	1. admin adds new role to VO 2. admin assigns VO-manager as default member to new role 3. admin assigns members to role
Secondary procedures	1. new role may not be authorized for resources 2. new role may not be assignable to members other than VO-Manager ("orphaned role") 3. new role may be conflicting with other roles
Non-functional requirements	• new role is effective immediately

5.3.4.4 Subcase: Administrator removes a role from VO

Use Case ID	004_d
Cross reference to community requirements above	ES: TG: IG:

Short description	The VO-Administrator removes a role from the VO.
Involved roles	`member, VO-Administrator, VO-Manager`
Prerequisites	The VO exists, the role is established and may have members assigned to it.
Primary procedure	1. admin detaches members from role 2. admin removes role
Secondary procedures	1. role may not exist 2. role may be blocked for removal (as it may be a critical or essential role like the VO-Administrator)
Non-functional requirements	• role deletion is effective immediately

5.3.5 Use Case 005: Merging VOs

Use Case ID	005
Cross reference to community requirements above	ES: 0,26 TG: IG: 1
Short description	VO A and VO B merge in VO B .
Involved roles	`VO-Manager, VO-Administrator, VO-MS`
Prerequisites	Both Vos exist.
Primary procedure	1. Archive VO A 2. Terminate VO A 3. Add members of VO A to VO B 4. Add roles of VO A to VO B 5. Add resources of VO A to VO B
Secondary procedures	1. members of A may not be authorized for B 2. members of A may already be in B 3. roles of A may be incompatible with B 4. resource access in A may be conflicting with access in B 5. members of A may be guest or dummy and not be accepted by B 6. sub-Vos of A may not be □latform□ble 7. super-Vos for A may not accept B
Non-functional requirements	• join is effective immediately • <<still open>>

5.3.6 Use Case 006: VOs as managed objects

5.3.6.1 Subcase: Formation of VOs

Use Case ID	006_a
Cross reference to community requirements above	ES: 13,17, 18,24 TG: IG: 1
Short description	A VO-Initiator creates a VO.
Involved roles	`VO-Initiator, applicant, RP, VO-Manager, VO-Administrator, VO-MS`
Prerequisites	VO-Initiator is authorized to create a VO.
Primary procedure	1. initiator invites applicants and RPs 2. initiator establishes roles (manager, administrator) 3. admin adds roles 4. admin adds members 5. admin assigns roles to members 6. admin and VO-manager deploy policies; members need to accept policies
Secondary procedures	1. All secondary procedures are related to the specifc use cases considered previously 2. Members do not accept policies
Non-functional requirements	

5.3.6.2 Subcase: Termination of VOs

Use Case ID	006_b
Cross reference to community requirements above	ES: 17,20 TG: IG: 1
Short description	A VO-Manager terminates a VO.
Involved roles	`RP, VO-Manager, VO-Administrator, VO-Member`
Prerequisites	VO-Manager is authorized to terminate a VO.
Primary procedure	1. VO-Manager stops VO 2. admin archives VO 3. admin deletes roles 4. admin deletes members 5. admin detaches resources 6. manager releases admin and manager roles
Secondary procedures	1. All secondary procedures are related to the specifc use cases considered previously
Non-functional	•

requirements	

5.4 Digest: Use Cases by Community

		Astro	HEP	C3	MediGrid	In-Grid	TextGrid	WISENT
1a	Regular resource access	X		X		X	X	
1b	Chargeable resource			X		X		
1c	Access by member of specific group	X		X		X	X	
2a	Add new member	X		X		X	X	
2b	Remove member	X		X		X	X	
3	Member moves						X	
4a	Assign role to member	X		X		X	X	
4b	Remove role from member	X		X		X	X	
4c	Add new role to VO	X						
4d	Removes a role from VO							
5	Merge VOs							
6a	Formation of VO	X				X	(x)	
6b	Termination of VO					X	(x)	

6 List of Abbreviations

IdP	Identity Provider
IVOM	Interoperabilität und Integration der VO-Management Technologien im D-Grid
RP	Resource Provider
SP	Service Provider
VO	Virtual Organization
VOMS	VO Membership Service
WAYF	Where Are You From

7 References

[anatomy] Foster, Ian ; Kesselman, Carl ; Tuecke, Steven: *The Anatomy of the Grid:Enabling Scalable Virtual Organizations*. In: International Journal of High Performance Computing Applications 15 (2001), Nr. 3, S. 200–222

[glossary] Treadwell, J.: *Open Grid Services Architecture – Glossary of Terms. Version: 1.5.* https://forge.gridforum.org/projects/ogsa-wg/document/draft-ggf-ogsa-glossary-1.5-006/en/6

[IVOM] Projektantrag *Interoperabilität und Integration der VO-Management Technologien im D-Grid*, Version 1.3 vom 27. Juni 2006

[IVOM1] P. Gietz, Ch. Grimm, R. Gröper, M. Haase, S. Makedanz, H. Pfeiffenberger, M. Schiffers: *Evaluation of International Shibboleth-Based VO Management Projects.* Report of Work Package 1 of the DGI IVOM Project, Version 1.1, May 2007

[OGSA] Foster, I.; Kishimoto, H.; Savva, A.; Berry, D.; Djaoui, A.; Grimshaw, A.; Horn, B.; Maciel, F.; Siebenlist, F.; Subramaniam,R.; Treadwell, J.; Reich, J. von: *The Open Grid Services Architecture*, *Version 1.5*. https://forge.gridforum.org/projects/ogsa-wg/document/draft-ggf-ogsa-spec-1.5/en/8

[RFC2119] S. Bradner (ed.): *Key Words for Use in RFCs to Indicate Requirements Levels*. The Internet Engineering Task Force Best Practice, March 1997. http://www.ietf.org/rfc/rfc2119.txt

[trustcom] Dimitrakos, T.; Golby, D. and Kearney, P.: *Towards a Trust and Contract Management Framework for Dynamic Virtual Organisations*. Proceedings eChallenges 2004

[VO] J.-M. Milke, M. Schiffers, W. Ziegler: *Rahmenkonzept für das Management Virtueller Organisationen im D-Grid*, November 2006. http://dgi.d-grid.de/index.php?id=118&no_cache=1&filename=VO_Rahmenkonzept_0.9a.pdf&dir=FG1/VO-Management&task=download&mountpoint=2

8 Acknowledgement

We would like to thank the D-Grid communities for their support in critically reviewing this document. Also Thanks to Nate Klingenstein, who gave some valuable comments to the TextGrid chapter.