19-10-2018

Deliverable DJRA1.3: VO Platforms for Research Collaborations

Deliverable AARC2-DJRA1.3

Contractual Date:	31-08-2018
Actual Date:	19-10-2018
Grant Agreement No.:	730941
Work Package:	JRA1
Task Item:	JRA1.4
Lead Partner:	STFC
Document Code:	AARC2-DJRA1.3
Authors:	Jens Jensen (STFC), Nicolas Liampotis (GRNET), Christos Kanellopoulos (GÉANT), Licia Florio (GÉANT), Mischa Salle (Nikhef), Uros Stevanovic (KIT)

Abstract

In order to scale the users' use of research infrastructures, cyber- and e-infrastructures, it makes sense to introduce a "virtual organisation" (VO) that can unify users with a shared purpose or research activity. This document investigates this use of the VO and makes recommendations for the *platform* which maintains this VO information, both for the VO's own use but particularly for the VO's members' use of the infrastructure.

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 730941 (AARC2).

Table of Contents

Executive Summary 1

1 Introduction 2
- 1.1 The purpose of the VO 2
- 1.2 Definitions of Terms 3

2 VO Lifecycle and Scalability 4
- 2.1 Supporting the VO Lifecycle 4
  - 2.1.1 Setting up a VO 4
  - 2.1.2 Registering with an infrastructure 5
  - 2.1.3 Maintaining a VO 5
  - 2.1.4 Suspension of VO 5
  - 2.1.5 De-registering the VO from an infrastructure 6
  - 2.1.6 Decommissioning a VO 6
- 2.2 Lightweight VOs 6
- 2.3 Scalability 7

3 VO Operations for Infrastructures 8
- 3.1 Roles and Responsibilities in VOs 9
- 3.2 Acceptable Use Policy 10
- 3.3 Membership Management 10
  - 3.3.1 Account Provisioning 10
  - 3.3.2 Account Maintenance 11
  - 3.3.3 Account Deprovisioning 12
- 3.4 VO Use of Infrastructure Resources 13
  - 3.4.1 Resources and Accounting 13
  - 3.4.2 User Support 14
- 3.5 VO Platform Operations 15
  - 3.5.1 Operating the VO Platform 15
- 3.6 Usability 15

4 Attributes 15
- 4.1 Use of Attributes 15
- 4.2 Attribute Management 16

4.3 Attribute Authorities 16
4.4 Attribute Schemata 17

5 Summary of Recommendations 19

6 Challenges and Opportunities 20

7 Conclusions 21

Appendix A Roles and Responsibilities of a VO 22

References 24

Glossary 25

Table of Figures

Figure 1: User, VO, and infrastructure 2
Figure 2: VO Platform in the BPA 2
Figure 3: users, roles, and services 7
Figure 4: voPerson example 18
Figure 5: voPerson example 18

Table of Tables

Table 1: Summary of Requirements of VO Platforms 20
Table 2 VO Roles and Responsibilities 23

Executive Summary

e-infrastructures and research/cyber infrastructures (also referred to as infrastructures) offer a variety of services to enable scientific collaborations for their user base. These infrastructures make use of Virtual Organisation (VO) *platforms* to improve scalability and reduce the effort required to support these users. Via the VO, an infrastructure can manage a group of users (such as those working on a common project) by interacting with the VO manager (or other appointed people) rather than interacting with users individually. Similarly, resources are allocated to the VO, instead of micromanaging resource allocations to individual users. This shared resource allocation also makes it easier for the end-users to share data/execution of tasks/applications. The *platform* is both the software and the instantiated service which manages the VO; it also offers the capability to manage the membership and the roles of its users, and to communicate these memberships and roles to authorisation services in order to make access control decisions for access to the infrastructure's resources. This can help prevent the risk that an individual user consumes more than their fair share.

Clearly, such a VO platform must conform to a number of requirements: it must be easy to use, so people use it correctly and don't take shortcuts; it must operate correctly, providing trustworthy data to the infrastructure; and it must manage the users' acceptance of the relevant acceptable use policies (AuP).

While we do not evaluate any specific VO platforms in this document, we discuss this list of recommendations which, if supported by the platform, should make it easier for the VO to manage its members correctly. Not all requirements will be pertinent to every situation, and to some extent, the infrastructure must make a judgment on whether a given platform is suitable. Nevertheless, there is a lot more to the VO platform than a simple attribute database, and it is important to ensure that platforms make it easy to support the correct workflows, both in terms of usability and functionality.

1 Introduction

Virtual Organisations (VOs) are, loosely speaking, groups of users with a shared purpose (we provide a more precise definition below, in section 1.2.) The main benefit of forming VOs, for the purposes of this document, is to make the VO's use of infrastructures more efficient (and hence the users' use of the infrastructure), as the infrastructure need only deal with the VO and not its individual users.

Figure 1 shows a UML diagram with the connections between the user, VO, and infrastructure.

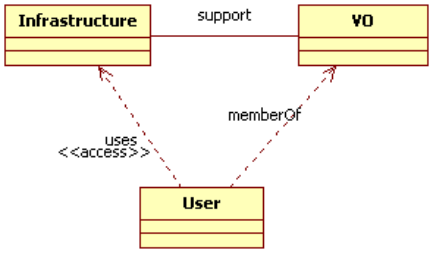

Figure 1: User, VO, and infrastructure

> Throughout this document, we assume that an infrastructure follows the AARC Blueprint Architecture [BPA] (see also Figure 2, below), which foresees an IdP-SP proxy [BPA], as well as the adoption of relevant policy frameworks the [SNCTFI], data protection frameworks ([CoCo] and [GDPR]) and relevant guidelines ([AARC-G006]).

1.1 The purpose of the VO

The purpose of a VO (from the perspective of this deliverable) is to unify a group of users of an infrastructure into a single entity. It is assumed that the users would have a shared goal/purpose, which is documented by the VO. Specifically, the VO has two main advantages:

- It scales efficiently to a larger number of users. The infrastructure can deal with representatives of the VO, rather than with users as individuals. Since a member of the VO is expected to perform work on the infrastructure that is similar to the work the other members perform, the VO structure and platform together make it possible for the infrastructure to liaise with the users through the VO as a single entity. Within the VO, typically most users will just have membership, and a minority of users will need to have special roles.
- Authorisation to use resources in the infrastructure can be managed through a user's membership in the VO. Additional restrictions based on the roles in the VO can be applied.

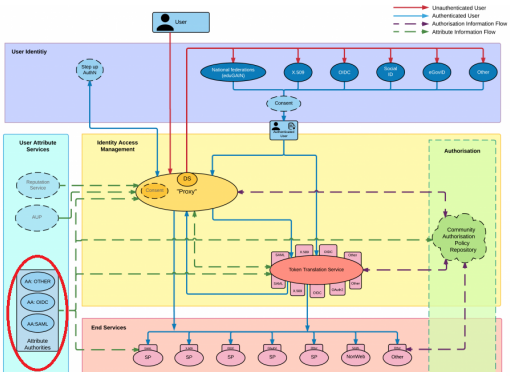

Figure 2: VO Platform in the BPA

It follows that the VO must define roles with which it interacts with the infrastructure, and there may be roles that are used internally in the VO as well. A fairly comprehensive list of such roles is given in 7Appendix A

Figure 2 shows the typical location and role of the VO Platform in the AARC Blueprint Architecture [BPA]. The typical role of the VO platform is to manage users' attributes and communicate them to the proxy where they will be processed (harmonised, translated, combined with attributes from other sources, etc.), and eventually relayed to relying parties (infrastructure authorisation services, resources, etc.) where they are typically used for access control and resource management decisions, and accounting (see 4.1.)

Fundamentally, the core message of this deliverable is that:

> The VO Platform must support the VO user's access to infrastructure resources –
> it should comply with applicable policies as well as support VOs' internal processes –
> and it should offer user friendly interfaces.

We will not look at specific VO platforms in this deliverable; rather, our interest is to identify the present and emerging requirements.

1.2 Definitions of Terms

The following definitions are based primarily on those of [AAOG], [SAML2], and the definition of Virtual Organisation is based on the "collection of users" in [SNCTFI].

An ***infrastructure*** is, in this document, a research, cyber-, or e-infrastructure.

An ***attribute*** is a named property. It consists of a *name*, and a *value*. The attribute may be defined by a *schema*, and may have a description of the *semantics* of the attribute. In particular, the schema may define the *multiplicity* of the attribute which would state whether it must occur at most once, exactly once, at least once, etc[1]. Note that the attribute itself does not specify who/what the attribute is about; this is the role of the *assertion*.

An ***attribute assertion*** is a collection of zero or more statements made by an *attribute authority*, usually about a single entity, the *subject*. In this case, one of the attributes typically assert the name or identifier of the subject, which usually requires that the subject can authenticate to both the attribute authority and the relying party with this identifier. An attribute assertion may also contain a timestamp of the time of issuance, any other limits to its validity, etc., as well as an attribute identifying the issuer. An attribute assertion is usually signed by the issuer.

An ***attribute authority*** is the issuer of attribute assertions. As with identity providers and certification authorities, *relying parties* must trust attribute authorities[2].

A ***community*** is a group of subjects having common (or at least similar) activities, goals and purposes.

A ***relying party*** is the entity that needs to make use of, and rely on the accuracy and timeliness of, the attribute assertion, and, in particular, of the binding between the identifier of the subject and the assertions. In this document, the relying party is usually the infrastructure, e.g. the authorisation services of the infrastructure or the resources themselves.

[1] The attribute may have multiple *values* but the schema defines whether the attribute is allowed to be multi-valued or not. So multiplicity is an example of metadata about the attribute (ASM, attribute schema metadata), not about the value(s) (AVM). See section 4.2.

[2] Technically, IdPs and CAs are also attribute authorities, as they provide attributes about the identity/credentials of the user.

A ***Virtual Organisation*** describes an organisational[3] entity. This entity represents a group of users that want to collaboratively use resources for a common purpose. Before entering a VO a user will[4] be requested to sign[5] its "Acceptable use Policy" (AUP). Acceptance of a user to enter the VO, or membership, may be subject to approval based on various criteria, in order to, for example, ascertain that the user's proposed work is aligned with the goals of the VO.

A ***VO Platform*** is a software product or service which implements, or at least supports and facilitates, the VO's processes and workflows, both with respect to the infrastructure used by the VO and with respect to the VO's internal processes. The VO platform will typically have human interfaces (VO members and administrators) as well as APIs for automated access (e.g. token issuance, or for external services to generate ACLs, see 3.1.)

2 VO Lifecycle and Scalability

2.1 Supporting the VO Lifecycle

This section describes the lifecycle of the VO and how it is set up, plus the lifecycle within its active period where members are joined, assigned roles, change roles, have roles unassigned, and leave.

A VO is determined by a set of technical capabilities, organizational measures and policies (including its goals and purposes) that together form the VO set-up. Some of the capabilities provide a background for the internal maintenance of the VO while others help meet the requirements stipulated on the VO by the particular target services or infrastructures, and must be adjusted for them. It is not necessary that the VO need address all requirements by its own arrangements. Some capabilities can be provided as a service by an external provider.

2.1.1 Setting up a VO

If appropriate, users should join existing VOs instead of setting up new VOs.

RECOMMENDATION: The VO Platform – if not run specifically for a single VO – should enable users to discover existing VOs and their purpose.

Some infrastructures may have a catch-all VO. The catch-all VO can accommodate members who cannot form a VO on their own, nor join another existing VO. The advantage of the catch all is that it can still have established AUPs and processes, even if members use the infrastructure for a more diverse range of activities than a dedicated VO.

3 "Organisational" in the sense of "a group of people who work together" and typically have defined roles and rules; the word "virtual" specifically denotes that the VO need not have a physical address nor be legally registered as an organisation.

4 Infrastructure policies generally make it mandatory for VOs to have an AUP. See also [SIRTFI] section PR.

5 Signing an AUP requires an explicit affirmative action by the user. The VO platform is required to track this.

2.1.2 Registering with an infrastructure

When a VO starts making use of an infrastructure, there is typically a set of requirements, such as defining necessary roles and responsibilities (see for example what [EGI] requires), ensuring the VO has an AUP, and that its members have asserted compliance with the AUP (section 3.2), etc. The list of requirements may vary from one infrastructure to the next, but a typical set of requirements is discussed further in sections 3.1 and 3.3.

From the point of view of this document, the inception of a VO is when it requests resources from an infrastructure. Whether the VO had an existence before "joining" the infrastructure, or was set up specifically to join the infrastructure, is in a sense less important – the important thing is to focus on the steps that are required of the VO for it to be set up. However, if a VO has an existing VO Platform before starting to use the infrastructure, then, obviously, it would be nice if it could be used:

RECOMMENDATION: In order to promote interoperation, VO Platforms should use standard schema and semantics, and standard protocols, or at least documented interfaces.

A VO needs to establish channels to services that the VO user will use so that information about the user can be communicated properly. Usually the VO specifies a set of user attributes that are delivered to the service as part of the interaction between the user and the service. In some cases, it may be necessary to communicate the information even before the user access the services (e.g. to populate mailing lists), which is where provisioning of user data is applied.

2.1.3 Maintaining a VO

Once a VO is set up and registered with an infrastructure, the VO will need to be managed in such a way that the requirements for the VO's use of the infrastructure remain satisfied. Since this is quite a large topic, and is central to the purpose of this document, it is the subject of Chapter 3.

2.1.4 Suspension of VO

In rare cases, a VO can be *suspended*, meaning they are no longer permitted to use the infrastructure. This can happen if the VO does not meet its requirements, including, for example, a member of the VO not adhering to the VO or infrastructure AUP. An infrastructure will generally reserve the right to suspend[6] individual users from using the infrastructure, but if the VO does not act in a satisfactory way to remedy misuse, the whole VO could be suspended. Although suspended by one infrastructure, the VO can continue to use other infrastructures.

[6] "Suspend" is used here to denote a user being temporarily barred from using the infrastructure, regardless of their other permissions attributes. Historically the word "ban" is also used. Suspension is usually done by checking the user's principal, so relies on the user being represented through a single principal.

The core principle is that the infrastructure can always override the VO's authorisations of their members. Since this topic is authorisation rather than VO Platforms, we shall not cover it further in this document (but see [AARC-G002], [AARC-G006].

2.1.5 De-registering the VO from an infrastructure

De-registering a VO refers to disconnecting the VO from the infrastructure it's using. The VO may continue to exist and use other infrastructures.

In practice, one would expect that a VO is de-registered only when its users have stopped using the infrastructure as members of the VO. Theoretically, a VO's use of the infrastructure can be revoked if the basis on which the VO was granted resources is no longer valid (see 2.1.4). For example, a project might finish, or a grant might be exhausted.

Note that the de-registering the VO may not be automatic; it may need to be requested and enacted.

2.1.6 Decommissioning a VO

A VO may decide to cease its existence when it is no longer needed or if there is another reason why the VO can't operate (lack of funding, for instance). Every VO should have a procedure defined that regulates the steps for decommissioning and responsibilities for that. It should de-register from all infrastructures it is registered with, following the appropriate rules specified by the infrastructures. The infrastructure may need to keep data about the VO and its members for future reference for a defined period of time, e.g. for auditing purposes, accounting purposes, or to meet its obligations to help resolve recent security incidents.

2.2 Lightweight VOs

One of the problems with setting up and maintaining a VO in many infrastructures is that it is a fairly heavyweight (1-2 weeks in practice) process [EGI][7]. In contrast, there is sometimes a need to set up lightweight collaborations, where a few people (fewer than 10, say) come together for a short period of time (up to, say, a month), and then cease collaborating.

The reader is reminded that *lightweight* here refers to the *processes* of setting up (2.1.1), joining an infrastructure (2.1.2) and maintaining the VO (2.1.3). Although a lightweight VO would typically have few members, this is not a requirement.

Obviously the easiest way to set up a lightweight VO is by creating a subgroup of an existing VO, with the assumptions that the participants are already members and the purposes are aligned. However, these assumptions are sometimes not fulfilled and it may be necessary to have an alternative.

[7] Note that we are not singling out EGI as having particularly heavyweight procedures; rather, EGI has one of the most mature and well documented frameworks for "on-boarding" new VOs, and their processes are therefore frequently used as guidance by other infrastructure operators.

RECOMMENDATION – The VO Platform should support lightweight collaborations in addition to the traditional VOs.

A VO Platform run by the infrastructure could make it easier to set up and maintain a new VO – but convincing the infrastructure that a new VO meets all their requirements, negotiating resource allocations, and defining roles and authorisations will still take time. We are still some way from a full VOaaS. An example of such a platform is SURFconext [SURF] that provides the capabilities to self-manage teams and to manage authorisation for service providers centrally based on such teams.

2.3 Scalability

VOs can start small and later grow into large ones, so it is recommended to start with a scalable platform. Scalability could look at scaling the number of users, roles, services, service instances, service access (that require authorisation), or infrastructures. Here, we look only at the users: if we expect 10,000 users, 20 roles, 30 (different kinds of) services, and 5 infrastructures - then users are the more challenging to scale. There may be large numbers of service instances and service accesses, but scaling those are out of scope of this document.

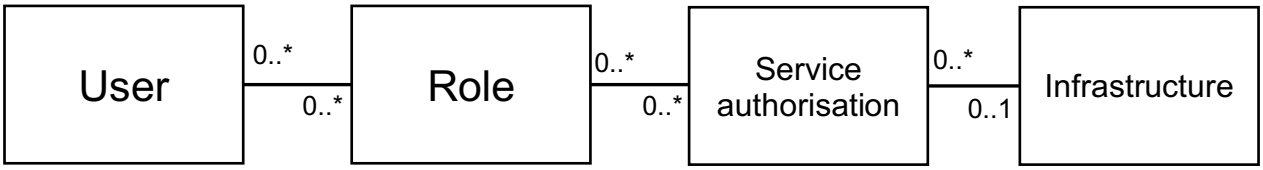

Figure 3: users, roles, and services

The first point to make, with reference to Figure 3, is that authorisation decisions may need more than the role. Consider a VO member, to whom is delegated the right/responsibility to manage software installed on an infrastructure, but only for (say) service X on infrastructure Y. We need to avoid introducing roles like "servicemgr-X-Y" because, (in our example) for 30 different types of services across five infrastructures, we have up to 150 service manager roles. Thus, the VO should introduce only the "servicemgr" role, but ensure that enough additional information is communicated to infrastructure Y for the latter to be able to make decisions regarding management of service X. This is an access control scheme where access rights are granted to users through the use of policies which combine attributes, potentially from different sources. Sources include user attributes from IdPs and VO Platforms and other attribute authorities, resource attributes, object, environment attributes etc.

RECOMMENDATION – The VO Platform should provide features that support an infrastructure implementation of scalable Role Based Access Control (RBAC) / Attribute Based Access Control (ABAC).

However, additional care should be taken to ensure that this information is not published unnecessarily, e.g. to a different infrastructure. As the VO Platform is not clairvoyant – it does not a priori know what the user intends

to do – the most sensible option is to give the users some control over which roles are asserted. Users may thus need to:

- Discover roles and request them, or have them assigned by an authorised person in the VO.
  - In particular, it should be clear what the purpose and rights of the role are.
  - If an approval process is required, someone who is authorised to make it should be notified, reminded, and either approve or reject the request (and rejection will require an explanation of why it was rejected.)
- View and select active roles.
- Delegate a role to another person, possibly limited to a time window.
- "Unrequest" a role, i.e. request to be removed from the role.

RECOMMENDATION The VO Platform should support user-friendly role management.

This user/role management is from the user's and VO administrator's view. From the VO's point of view, it needs to make sure that, for certain roles, there is always someone in that role. Thus, a smarter role management system would manage the availability of people with the right skills, and keep some people in rota, or as secondary role-holders (if the primary cannot solve the problem), or backups (if the primary role-holders are not available.) The best examples are probably in the user support role, where helpdesks and support rotas ensure that there is always someone available to respond to requests in a timely fashion – such a model could be replicated to other VO responsibilities.

As we have seen, the VO is introduced to ease the scalability of the number of users of the infrastructure, by unifying them into a common group. As the VO itself scales (i.e. gets more members, who are in turn users of the infrastructure), it becomes necessary to either reduce the number of manual processes (requiring human intervention) in the management of the VO, or to delegate these management tasks to more people. Potentially, some roles could be authorised by policy.

RECOMMENDATION – The VO Platform should support scalable VO operations through, as appropriate, delegation, automation, RBAC, and by ensuring that any constraints are met.

3 VO Operations for Infrastructures

Since we look at VOs in the context of their use of infrastructures, we need to discuss the requirements on the VO arising from this use, and how these requirements can be met. While this discussion is rooted in the current best practices, we also aim to look ahead to the future.

3.1 Roles and Responsibilities in VOs

As mentioned in section 2.3, roles are important in the scalable management of authorisation. The VO Platform must support users' discovering roles, applying for roles, getting approved (or not), getting notified, and ensuring that the process is logged in an auditable form. This is true both when the VO is set up (2.1.1) when roles are initially populated, and in its ongoing operations when they are maintained.

Role based access control is used to make authorisation more scalable (see 2.3); however, attributes provided by the IdP may also be useful for authorisation [AARC-G006]. For example, access to data can be authorised based on organisational membership (published by the IdP, if it's run by the organisation.) Access to the service can also be identity-based, with the principal (as published by the IdP) listed directly in the access control list (ACL) – this would be necessary to give access to a service which doesn't support role based access control (RBAC). In this case, the list of identities should still be obtained (and updated sufficiently frequently) through the VO platform.

RECOMMENDATION – The VO Platform must support RBAC (e.g. if special APIs are needed, to allow services to query a user's role in the pull model for authorisation, cf. [AARC2-G006])

RECOMMENDATION – The VO Platform should provide an API to enable non-RBAC services (or services with only a few users) to build and update their ACLs.

In our discussion below, we focus mainly on roles, partly for the sake of simplicity, and partly to focus on attributes managed by the VO Platform (and because a full discussion of authorisation is out of scope of this deliverable.)

- The roles need not be mutually exclusive. A VO Manager is usually also a Member of the VO[8]. The VO Manager would in principle be responsible for evaluating the membership applications but in practice would delegate this task to others (see 2.3), particularly for a distributed VO where local knowledge may be needed.
- A table in 7Appendix A shows our extrapolation of roles and responsibilities of current VOs to a large VO:
    - We do not require, nor do we expect, that a given VO will implement all these roles. The roles are included only for the purpose of illustrating the discussion of the management of roles and responsibilities.
    - The focus of the table is mainly the *responsibilities* expected of a VO, but a responsibility is likely in practice to be assigned to a role. For example, a VO must be able to suspend a misbehaving user (if the infrastructure cannot do so directly, e.g. if the VO uses a robot certificate), but the role likely to do so would be a security officer.
- Not all roles/responsibilities need to be supported by the VO platform itself. However, doing so may improve scalability – and/or security – because there will be fewer locations where information has to be maintained.
- Some of the responsibility for the VO's users' use of the infrastructure may be shared between the VO and the infrastructure. For example, operational security of the infrastructure is likely to rest primarily

[8] Conceivably, for a small VO, the infrastructure might manage the VO on behalf of the VO.

with the infrastructure, although users would have a responsibility not to compromise this. In contrast, responsibility for the applications rests primarily with the VO, although the infrastructure will have a responsibility in turn to run them securely. The responsibility for the resolution of security incidents involving the VO's users or applications is generally shared between the VO and the infrastructure, as they will likely need to collaborate in the resolution.

3.2 Acceptable Use Policy

All infrastructures have AUPs, and so do the NRENs connecting them. (Since specific AUPs are not the focus of this document, we may assume that all relevant AUPs have been combined into a single one. In this respect, it is worth mentioning AARC2 Policy work on the alignment of AUPs [AARC-AuP]; the work is being finalised at the time of writing). It is necessary to ensure that all users of the infrastructure are aware of the contents of the AUPs and have signed them (electronically) recently.

Note that we do not specify how this "signature" should be implemented. It could be a click "ok" on a web page after ticking a box saying they have read it and understood it, sending a mail to a relevant person (role) in the VO, digitally/electronically signing with a user X.509 certificate, or printing out a piece of paper, signing it, and mailing it to the relevant person/role. It may make sense for a VO Platform to support several of these methods, in order to meet the requirements of different VOs.

RECOMMENDATION – The VO Platform must support management of AUPs – tracking user's acceptance and refreshing their acceptance at required intervals.

3.3 Membership Management

Membership is a role but it is also a special type of role as it is in general the first role (and possibly only role) a user gets, and the least privileged role which grants access to resources. We may reasonably assume that every person in the VO who has access to resources in the infrastructure must be have the membership role.

The user may need to obtain a new credential in order to join the VO and/or use the infrastructure, typically if the user's existing credentials are not of a sufficient level of assurance [AARC-MJRA1.2], and step-up [AARC-G029] is not available or suitable. It is RECOMMENDED that this process be made as clear as possible, possibly through the VO Platform. Similarly, the platform should support the acceptance of all the required policies (section 3.2), not just that of the VO.

3.3.1 Account Provisioning

The VO Platform must support any traceability requirement associated with granting membership to a person – or, to be precise, associated with granting authorisation (through membership) to resources provided by the infrastructure. In particular, since users must accept AUPs before joining theVO, the platform must make sure this acceptance is recorded in an auditable form.

RECOMMENDATION – The VO Platform must support traceability, particularly if the user's identity is hidden from the infrastructure. (See also 3.4.1 and the footnotes to that section for information about hiding the user's identity.)

3.3.2 Account Maintenance

As a part of the GDPR [GDPR, AARC-G042], data controllers have an obligation to ensure that the data held about users is relevant and adequate (so, for example, not require the user to provide their date of birth when registering.)

Another of the requirements of the GDPR is that the data be kept up to date. This activity is not necessarily the user's top priority, so some means need to be found to ensure that the user is encouraged to keep the data up to date, or finding other means of maintaining the data.

Obviously, account maintenance must comply with the GDPR. Specifically, there must be a means for the user to see (or request) personal data that is held by the VO platform, and to request that inaccurate information be fixed, to see how the data will be used, to consent to its use, to see the relevant legal basis for processing of their personal data, and to request deletion of the account (see next section). Data must be secured in the platform in order to minimise the risk of leakage. Also the GÉANT Code of Conduct [CoCo] describes the necessary measures when processing personal data in order to provide a service.

RECOMMENDATION – The VO Platform must facilitate compliance with the GDPR

More specifically:

- Support roles for GDPR (7Appendix A), cf. G042: for example, there must be a data protection officer role – or some other facility – that can respond to user queries about which personal data the platform holds about them, which may require privileged access to the platform's data.
- Support scalable attribute release policies.
    - In particular, as the VO Platform is an infrastructure constituent [SNCTFI], and must support the processes of [SNCTFI] (including the GÉANT [CoCo], and [SIRTFI])
- Support limiting the data storage period in compliance with applicable policies.
- Provide mechanisms that allow users to update or remove their information in order to keep it accurate and up to date (section 3.3.2).

RECOMMENDATION – The VO Platform should facilitate compliance with the GÉANT Code of Conduct [CoCo].

Note that attribute release is necessary in order for the infrastructure to make authorisation decisions and grant the user access to resources ([CoCo] defines authorisation attributes as "NECESSARY" [CoCoSP], so it follows

that the VO Platform must release all attributes required for authorisation[9]). Micromanaging the release of every attribute to every resource is likely to be more of an annoyance than a help, so the VO Platform should make it easy to implement a transparent (to the user) release policy that is consistent with the proposed use of the infrastructure[10]. In particular, it should be clear if attributes are released to third parties. Usability is also important: if the user is likely to blindly click "yes" in a browser redirect, with no clear visibility of what will be released and what will happen to the data, then the implementation of the policy is not effective.

3.3.3 Account Deprovisioning

Yet another requirement of the [GDPR] is that the data be held only for as long as is necessary. Once the user has left, after a suitable period of time (e.g. three months[11]), the data should be deleted (note that this is different from inactivity.) Moreover, the user can request that the data be deleted. Note that some data may need to be kept for a defined period of time in order to comply with traceability and accounting requirements in the infrastructure.

Deprovisioning of accounts refers to the case where a user's account with a VO is *deactivated* or *deleted.* Deactivation retains the account in the sense that it can be reactivated, possibly after additional checks, but users cannot obtain authorisation through the VO.

Potential reasons for account deprovisioning are:

- The user has not used the account for a specified period of time (e.g. 12 months.)
    - The reason for deprovisioning here is to not keep data longer than necessary or the need to free up resources allocated to the user.
- The user's (mandatory) data is inaccurate (e.g. email bounces permanently), i.e. the user has failed their obligation to maintain their data, see 3.3.2.
- The user has violated the VO's AUP[12].
- The user has "left" – i.e. they are no longer with the home organisation, and have not moved to a related organisation.
- The user has requested the deprovisioning (as required by the GDPR).

After an account (as identified by some identifier) has been deprovisioned, it may be:

- Reactivated, if it was only deactivated. In this case, additional checks may be needed to ensure that the account data is still current.
- Reopened, if it was deleted. This is the case where the user has had their account deleted but come back at a later time and requests it back[13]. In this case, the VO has no memory of the user, and should be careful to ensure that membership and authorisations are granted to the same user – services may not have deleted their authorisations associated with the original account.

[9] As mentioned in section 2.3, the VO Platform does not necessarily know in advance which attributes the user will need for authorisation (but see also [CoCoSP] and [AARC2-DJRA1.2] section 4).
[10] In particular, the REFEDS Research and Scholarship category ("R&S") is important [REFEDS-RS].
[11] The actual time would be determined by the relevant policies.
[12] If the user violates the infrastructure AUP, it is common practice to suspend the user from using the infrastructure, rather than ask the VO to terminate their membership. See 2.1.4.
[13] In terms of authorisations, there should be no strong case for getting their old account back; it would be better to get a new one and assign the relevant roles and other authorisations to the new account. However, there may be cases where the user can prove they had the old principal and have it reallocated. Alternatively, the VO platform operator may legitimately refuse to reopen deleted accounts.

- Reallocated. Like the previous case, but the old name/principal is requested by a different user (perhaps a different person with the same name.) As in the previous case, services may not have deleted their authorisations associated with the original account owner (in the context of the VO), and may not have been notified of the deletion and may think they are talking to the original user (through an ACL; see also [AARC-G026]). Thus, the new user may obtain rights they should not have had[14].

RECOMMENDATION – VO Platforms should include features to prevent account reallocation (through the non-reassignable identifier[15] if available).

One interesting subtopic is the deletion of user data: if an IdP is to *guarantee* not to reallocate a user's principal (say, givenname.surname@org.domain.country), how is it to do so without remembering the principal? A solution could be implemented here involving hashing the original principal, and matching the hash of any future prospective principals against a list of all past hashed principals. In this case, the original data cannot be recovered directly from the hash alone (and is hence not a problem with data protection regulation), and there is enough assurance (at least for cryptographic hashes offering second preimage resistance) that two distinct principals will never map to the same hash.

Although authorisation is not the topic of this deliverable (see [AARC-G006], [AARC2-DRJA1.2]), a distinction between two ways of authorising users is pertinent: In the first case, the VO platform issues timestamped attributes or tokens which are used "downstream" by authorisation services (usually through a proxy). In this case, the attribute or token is accepted only if it is freshly issued. A similar case arises when the token issued needs to be validated against the VO platform, as in OAuth2. Contrast this with an authorisation service which builds a membership list from time to time by querying the principals of the members of the VO in order to compile an ACL. For the latter, "freshness" of the ACL can only be guaranteed as long as the ACL is continuously updated and there is no automated check, an ACL may grow stale without warning. However, it is sometimes necessary to use this approach.

RECOMMENDATION – A VO platform issuing authorisation tokens or attributes should provide the means for the relying party to check the freshness of the attributes.

3.4 VO Use of Infrastructure Resources

3.4.1 Resources and Accounting

One of the premises behind the introduction of VOs is that resources are allocated to the VO and are then shared as appropriate by its members. Obviously, there is a risk that a single user may consume all of a scarce resource and leave nothing to other members of the VO. Another perhaps less obvious risk is where the VO is

[14] According to AARC-G036 "The minimal technical requirement to link an external identity is the availability of a persistent, non-reassignable, and unique external identifier." So arguably account reallocation is a "cannot happen" scenario (as software engineers know, "can't happen" scenarios sometimes happen.)

[15] Sometimes the principal cannot be guaranteed to not be reassigned. In this case, the R&S category [REFEDS-RS] requires that additional attributes be available to make the combination unique.

granted rights to sensitive data through a shared agent acting on behalf of the VO or a portal, so the access control is "outsourced" to the VO. In this case, there is a risk that members might escalate privileges within the VO's control and gain access to a resource they should not have been able to access. Usually, these types of situation require collaboration between the VO and the infrastructure, in order to resolve them.

RECOMMENDATION – It may make sense to have VO Platforms integrated with quota/account management systems, so it knows whether to authorise access to a scarce resource to an individual.

Another sometimes sensitive issue is where the researcher is anonymised (e.g. in biomedical research, so is unknown to the infrastructure administrators[16]), or is anonymised through the use of "robot" credentials [ROB]. In both cases, the VO needs to be able to resolve who the individual user is, as the infrastructure has no way of knowing, and traceability is required by [SIRTFI].

RECOMMENDATION – The VO Platform must have logging capabilities of users' actions, so that the resource usage and its allocation can be monitored, and, if appropriate, moderated.

3.4.2 User Support

As mentioned earlier, the responsibility for support for the end user's (i.e. VO member's) work that uses infrastructure resources is also shared between the VO (or specifically, the VO's holders of the support role) and infrastructure support. Most likely, the user will not bother with this distinction and just request support. Consequently, infrastructures have helpdesks which enable assigning support tickets to the relevant parties. It follows that some of the supporters would need to (a) have domain-specific knowledge, in order to support the work, (b) be able to support the work, and (c) be members of the VO. Membership of the VO is helpful also for infrastructure support staff, so they can try to replicate the user's problems.

RECOMMENDATION – The VO platform should make it easy for the end user to find the right documentation and support.

[16] In biomed, there is occasionally a requirement to hide the researcher's real-life identity from the infrastructure, typically used for controversial research which could lead to threats to the researcher if it was known to the public. In the commercial world, sensitivity can also arise from commercial interests, e.g. researcher X working on protein Y – if this information were known to a competitor, they, too, would take an interest in protein Y.

3.5 VO Platform Operations

3.5.1 Operating the VO Platform

The VO holds data about its users and metadata about the data. This data has to live in a database somewhere, with an appropriate implementation of access controls and logging/auditing, etc. The purpose of this section is to describe the requirements on operating such a platform. Current guidelines on operating an attribute authority [AAOG] were documented by the IGTF, by comparison with an online CA. Since then, SIRTFI, CoCo, SNCTFI, and GDPR have been added as requirements on infrastructure participants and operations, respectively, which will enable us to shorten the reference to the requirements. The AARC policy team is currently working on an update of [AAOG].

More work may be needed on streamlining VO processes in current platforms.

3.6 Usability

Usability becomes important in several contexts. When users first discover the VO, request membership, request roles, and access the platform in order to manage their account, etc., usability is important. When access control policies are managed, the right balance between expressiveness of the policy language and usability of the language becomes important: if it is too complicated to define the correct roles or rules, administrators will take shortcuts which might in turn lead to privilege escalation (section 3.4.1).

RECOMMENDATION – The VO Platform should be designed and tested for usability for both administrators and end users.

4 Attributes

At its core, the VO platform is all about attributes about the members of the VO. The attributes about the authenticated user help resources make access control decisions and resource allocations and accounting.

4.1 Use of Attributes

It may be useful to briefly summarise the purposes for which a relying party might consume attributes managed by the VO Platform (and relayed to it through the proxy); the points below should be seen more as a checklist:

1. Authorisation, particularly RBAC
2. Resource allocation and accounting (e.g. against a VO allocation/quota)

3. Traceability – if the user is anonymised when accessing the infrastructure (see 3.4.1), the VO Platform or the proxy must provide the required traceability; thus there must be a session attribute or pseudonymous credential which can be used to trace the user.
4. Compliance – as we shall see below (section **Error! Reference source not found.**) there is a means of communicating to the relying parties the user's acceptance of the AUP (see also 3.2 for a brief discussion of AUPs).

Note that metadata *about the VO itself* such as the name of the VO, information about the VO's purpose and membership processes, endpoints for APIs, identity of the signing credential, AUP, etc., are typically not communicated through the proxy as it would be the same for every member, and security is improved by managing this information out of band. Nevertheless, it is still the responsibility for relevant people in the VO to ensure that this information is available and is maintained.

4.2 Attribute Management

The VO platform (via the relevant proxy) is the authoritative source of attributes pertaining to the members of the VO (see Figure 2.) Some of those attributes will need no special management: timestamps, the identity of the signer for signed assertions, the name of the VO, etc., do not need people to manage them. Others, will need to be managed by people who are themselves authorised to manage them by holding a role in the VO. It is the role of the proxy [BPA] to ensure that attributes are harmonised or translated appropriately, so they can be presented and parsed correctly by relying parties [AARC-G002], [AARC-G006]. The proxy already manages attributes from the IdP and the VO platform [AARC-G006], and may also manage proxy-internal attributes (for example, making up for inconsistencies in the attributes provided by IdPs), and, potentially, the proxy could be managing attributes from yet other sources.

Note that the VO Manager that assigns a role to a user (or approves the user's request for the role) may need to see and verify the level of assurance associated with the account in the proxy, in order to comply with relevant policies. While the level of assurance of authentication and of the account is not in scope for the VO Platform – it is the responsibility of the proxy – it may be necessary for role management. Conversely, the user should be able to see what information is required to apply for membership/roles.

By the time the attribute reaches the relying party, it might be interested in where the attribute came from. An earlier EUDAT study [EUDAT] found that the most mature work was NIST's internal report 8112 [NIST-IR8112]. A follow up technical report by EUDAT is still to be published; see section 0.) Note that this is different from the assurance associated with the user authentication and the representation of the user's credential, such as the REFEDS Assurance Framework (RAF) [REFEDS-AF].

4.3 Attribute Authorities

Attributes are used by the infrastructure to manage access to resources: at a basic level, the infrastructure needs to know whether a person is a member of the VO in question (or rather, the member needs to assert membership to the infrastructure.) Some attributes are managed by the VO platform, some come from the user's IdP, and yet other attributes may come from different sources.

Since some attributes are used for access control, it follows that these attributes need to be:

1. Communicated to the resource, either directly (through push or pull, cf [AARC-G003], or indirectly ("back channel"), such as the ACLs mentioned in section 3.1.)

2. Trusted: the infrastructure needs to be able to trust that the attribute is maintained (has been checked within a reasonable time interval, has been asserted within a suitable time window, has been granted based on suitable processes, and can be revoked in a timely manner when no longer needed), and that it is asserted by a trustworthy authority (such as a VO Platform).
3. Communicated in a suitable format, in an appropriate schema (section 4.4) and with appropriate semantics, in order to ensure correct interpretation and interoperation between infrastructures.

4.4 Attribute Schemata

Note that the voPerson schema [voPerson], as of August 2018, is not finalised, so the discussion below is based on the current draft.

Traditionally inetOrgPerson [RFC2798] – itself derived from `organizationalPerson' which in turn is derived from the X.500 `person' schema (cf. [RFC2256]) – provides a very rich schema to describe a person with their organisational affiliation, such as `surname`, `title`, `preferredLanguage`, `carLicense`, `jpegPhoto`, `departmentNumber`, `employeeType`, and much more. (In practice "derived from" means a directory entry can contain attributes from all of these schemata; however, an entry need not use all attributes and in practice, only a subset is used.)

For the purposes of educational organisations, eduPerson [eduPerson] was created as a further specialisation, i.e. the schema provides additional attributes that are meaningful/useful in an academic or other educational context. However, the attributes are typically published by the home organisation and relate to the user's identity and organisational affiliation/role, and are not directly suitable for VOs as the publication of these attributes – for example, an `eduPersonEntitlement` published by the home organisation's IdP – is outside of the VO's control. It makes sense to introduce a schema to describe a person's role within a VO. However, the `eduPersonEntitlement` attribute itself is still useful as an attribute to carry authorisation information [AARC-G002]

The reader wishing to make use of the original voPerson specification is cautioned that their terminology is different from ours. What we would call a "VO" (section 1.2) in this document is called an "organization" in theirs:

> "A VO is an organization that includes members whose identity information is obtained from multiple sources, at least one of which is external to the organization. The organization may or may not be a legal entity."

Of the different attributes in the voPerson schema, some can refer to identities the user obtains from IdPs external to the VO (e.g. a CA, a home IdP, etc.), but these could potentially be managed by the use of the earlier schemata. For our purposes (i.e. the description of roles or responsibilities) only two of the attributes defined in the voPerson schema are pertinent:

`voPersonPolicyAgreement` is likely to be useful to assert a Member's acceptance of a AUP; an option (see below) provides the time when the user agreed to the AUP.

`voPersonStatus` provides information about the user's role/status inside the VO. The value of this attribute comes with a controlled vocabulary that describes the user's status in the VO, or, using an option (see below), the status of a specific role or scope.

`eduPersonEntitlement` from the eduPerson schema is commonly used to communicate authorisation attributes, such as those expressing VO/group membership information [AARC-G002].

We have referred to options for the attributes. While every attribute has a type, specified as an OID (e.g. 2.5.4.6 for `countryName` or `c`) which is defined in a schema, [RFC4512] introduces the concept of *options* for the attributes. Less frequently used in practical LDAP deployments, it is used (here) to *tag* the attribute with metadata. The options/tags used by the two attributes above are time-# (for `voPersonPolicyAgreement`) and role-* and scope-* for `voPersonStatus`. An example (Figure 5) will demonstrate how they are used:

```
…
objectClass: eduPerson
objectClass: voPerson
eduPersonPrincipalName: joe.user@organisation.ac.uk
voPersonPolicyAgreement;time-1526398035: https://www.egi.eu/terms-of-use/
voPersonStatus: active
voPersonStatus;role-securitycontact;scope-egi: pending
…
```

Figure 5: voPerson example

Notice how the options refer to roles which need to be defined elsewhere; they are not part of this schema – only the status (active, pending, etc.) are defined by the voPerson schema.

Defining tags and making sure they are consistent is a bit more involved, but here we shall only make use of existing tags rather than define new ones. With this in mind, we can now try to map voPerson to our requirements:

- `voPersonPolicyAgreement` is relevant for a VO member to signal compliance with an AUP; they might also be useful for other compliance statements. As it is multi-valued, several could be asserted – conversely, as they are not targeted to particular resources or applications, all the relevant agreements would have to be published simultaneously.
    - As a corollary, publishing this attribute means the VO's attribute authority now maintains the status of their members' compliance with the infrastructure's AUP. This means that the infrastructure does not have to ask for it, unless, for a particular member, it is absent or has expired. However, one should be careful about maintaining the information in several places: if the infrastructure runs the attribute authority for the VO, they could update the field; if the AUP is used towards several infrastructures, one would have to be the authority for the user's acceptance.
- `voPersonStatus` can be used to list the user's roles and the status of a rule (i.e. whether it has been requested, suspended, or revoked), so could potentially reveal information that the user, or the VO, might not wish to reveal. Also, as mentioned above, the roles need to be defined somewhere.
    - Roles could be targeted to particular infrastructures, if the scope-* option/tag is used for the role. The voPerson recommended use of "scope" describes the string following "scope" as "a label describing the scope or source of the value" so we might reasonably use this for the target infrastructure. Note that we could not have used in the example "scope-egi.eu", as the full stop character is not valid in this context [RFC4512].

RECOMMENDATION – the VO Platform should use standard schemata such as eduPerson and VOPerson to communicate attributes (both when exposed to the resource directly, and when communicated to a proxy.)

---

5 Summary of Recommendations

The table below summarises that list of recommendations provided by this document.

Requirement	Description	Reference (section)
1	Make it easy for users to discover existing VOs and their purpose	2.1.1
2	Use standard schemata (inetOrgPerson, eduPerson, voPerson) and their semantics, and standard protocols, or at least documented interfaces.	2.1.2, **Error! Reference source not found.**
3	Support lightweight collaborations	2.2
4	Support RBAC/ABAC, and provide features to implement scalable RBAC/ABAC, e.g. by scoping roles (if necessary) separately.	2.3, 3.1
5	Provide user-friendly role management (discovery, application, notification, etc.), both for the applicant and the people who approve/deny the request. Users should be able to see what information is required, and all the information should be available to the approver in a single place.	2.3, 4.2
6	Provide features for delegation and automation of tasks, and support role constraints (e.g. “there must be a security contact at all times”)	2.3
7	Provide APIs for services and automation	3.1
8	Support the AUP process – tracking users signing and re-confirming all relevant AUPs	3.2
9	Make the workflow needed to join a VO as clear as possible.	3.3
10	Support user traceability, particularly if the user’s identity is hidden from the infrastructure (3.4.1)	3.3.1
11	Facilitate compliance with the GDPR	3.3.2
12	Ensure compliance with SNCTFI with the VO Platform as an infrastructure constituent [SNCTFI] of all infrastructures used by the members of the VO. Facilitate compliance with the GÉANT Code of Conduct	3.3.2, 3.5.1
13	Prevent account reallocation	3.3.3
14	Provide means for relying parties to check freshness of authorisation attributes	3.3.3

15	Consider integration of VO platform with account/quota management	3.4.1
16	Ensure user's activities in the VO platform are logged, so the platform can participate, if needed, in the resolution of security incidents	3.4.1
17	Make it easy for users to discover documentation and get support	3.4.2
18	Assess VO Platform for usability, both for the end users and for administrators	0

Table 1: Summary of Requirements of VO Platforms

6 Challenges and Opportunities

- LDAP is often at the core of attribute management, partly because of its history and its use of schemas. It is used by several infrastructures and by every organisation running Microsoft Active Directory. As voPerson (section **Error! Reference source not found.**) has highlighted, we do not always use the full capabilities of the protocols of choice even if they are standards-based and ubiquitous: voPerson uses the option/tag associated with an attribute, and LDAP provides other features such as the ability to define custom SYNTAX and USAGE definitions.
- Scalability of access control. While the introduction of roles and policy-based access control (section 2.3) allows access control to scale to more users and more services, complex policies can become difficult to manage. In particular, fine-grained access control, such as controlling individual users' access to millions of objects, is a challenge.
- Can we implement lightweight VOs (section 2.2), i.e. (typically, but not necessarily) smaller collaborations where the processes are simplified?
- Implement RAF [REFEDS-AF] and/or [NIST-IR8112] in order to communicate provenance, freshness, assurance level, and scope (and consent?) to relying parties. Both of these are fairly recent developments, so there is an opportunity to use them to solve the problems of managing and communicating attribute metadata. Yet they also pose challenges because:
    - The means by which the attribute metadata is communicated to the authorisation system is not always specified – particularly for non-SAML protocols.
    - How the relying parties make use of the extra information is also relatively unexplored. For example, a relying party should accept credentials that carry the extra information as well as those that don't, and the same level of assurance may be expressed in different ways by different credentials. In comparison, X.509 certificates issued by IGTF CAs have since at least 2009 carried policy OIDs to communicate their level of assurance, but even today, these are rarely if ever used.
    - The attribute metadata will likely have to be fitted to existing VO platforms and other attribute management systems.

- It is necessary to have means to set the attribute metadata correctly, and, in order to scale, these will have to be automated and policy-driven.

7 Conclusions

VO Platforms are essential to the management of VOs. In turn, VOs themselves – defined as groups of users with shared and/or similar use of the infrastructure – enable the infrastructure to support larger numbers of users, because the users are largely managed through the VO, and the VO through the VO Platform. This document discussed some of the functional requirements on the VO Platform, highlighting in particular features that may make the use of the infrastructure safer (in the sense of minimising the risk of mistakes), and more scalable. However, while we have not evaluated any existing VO Platform, we provide a baseline for such an evaluation, as well as, we hope, useful guidance for developers and deployers of VO platforms.

Appendix A Roles and Responsibilities of a VO

The following table lists a set of roles/responsibilities of a VO. Not all VOs will have all these roles, and only a small number are mandatory (the PI, the data protection officer). See section 3.1.

Role or Function	Source	Responsibility
Member		Accept and abide by the VO AUP, use the resources only for the purposes described by the VO. Must be contactable (by email)
Principal Investigator (PI), and, possibly, co-investigators (co-PIs)		Responsible for the VO and its activities to the funding body (resp., responsibility for individual partners' finance and activities).
VO Manager		Define scope of VO, and the VO AuP. Ensuring that the VO AuP is aligned with the requirements of the Infrastructures.
Security contact, security officers	EGI, SIRTFI	Collaborate with infrastructure in case of a security incident
User membership manager		Approve or reject user membership requests, remove users who are no more entitled to be members in a timely manner.
Group manager		Assign sub-group membership attributes to VO members who request it. The role is often appointed for every group separately.
Role manager		Assign roles attributes to the VO members who request it. The role is often appointed to every managed role separately.
User suspension	SIRTFI	May be managed by infrastructure (revocation of access), or by the VO's security officer (revocation of permissions), or both
User contact	SIRTFI	May be managed by infrastructure. But the VO may need to integrate contact information in case

		the IdP is not providing reliable data for the user (depending on what is requested by the infrastructure).
Data protection	GDPR	Data protection officer (data processor, data controller). Responsible for personal data, both of VO members and the VO's data processing if this data contains PII.
Infrastructure integration		Technical support for integrating a VO's software (applications and libraries) with the middleware provided by the infrastructure. May need a superset of permissions of VO members in order to test and troubleshoot the VO's applications and workflows.
User support		Help users run VO-specific applications on the infrastructure. May need a superset of permissions of the users in order to try to replicate and resolve the users' problems. May need privileged permissions to, say, view a user's permissions.
Data policy		Determine policy for the VO's data processing (use by its members, such as what data may be processed on the infrastructure), and data publication (policy for use of data by others, e.g. as required by funding body.) May need privileged permissions enabling them to inspect other user's data under well-defined circumstances, e.g. if abuse is suspected.
Data permissions		Who can read the VO's data: process requests for access to restricted data by non-members. Requires access to user's permissions.
Data QA		Quality assurance of data published by a VO or its members. May need access permissions to datasets, and/or metadata.

Table 2 VO Roles and Responsibilities

References

[AAOG]	Attribute Authority Operations Guidelines, https://www.eugridpma.org/guidelines/aaops/
[AARC-G002]	Guidelines on expressing group/role https://aarc-project.eu/wp-content/uploads/2017/11/AARC-JRA1.4A-201710.pdf
[AARC-G003]	Guidelines for Attribute Aggregation https://aarc-project.eu/guidelines/aarc-g003/
[AARC-G006]	Guidelines for Authorisation https://aarc-project.eu/wp-content/uploads/2017/03/AARC-JRA1.4E.pdf
[AARC-G026]	(to appear - https://wiki.geant.org/display/AARC/AARC+Architecture)
[AARC-G029]	(to appear - https://wiki.geant.org/display/AARC/AARC+Architecture)
[AARC-G042]	Data Protection Impact Assessment – an initial guide for communities, https://aarc-project.eu/ guidelines/aarc-g042/
[AARC-MJRA1.2]	https://aarc-project.eu/wp-content/uploads/2016/06/MJRA1.2-Design-for-Deploying-Solutions-for-Guest-Identities.pdf
[AARC2-DJRA1.2]	Scalable, integrated authorisation models for SPs, https://aarc-project.eu/wp-content/uploads/ 2018/07/AARC2-DJRA1.2_V4-FINAL.pdf
[AARC-AuP]	https://aarc-project.eu/wp-content/uploads/2018/04/2018-04_Athens.pdf
[BPA]	https://aarc-project.eu/architecture/
[CoCo]	GÉANT Data Protection Code of Conduct https://wiki.refeds.org/display/CODE
[CoCoSP]	[CoCo] for service providers https://wiki.refeds.org/display/CODE/What+attributes+are+relevant+for+a+Service+Provider
[eduPerson]	https://www.internet2.edu/products-services/trust-identity/eduperson-eduorg/
[EGI]	https://wiki.egi.eu/wiki/Policies_and_Procedures
[EUDAT]	http://doi.org/10.23728/b2share.20c1c0c8ba254e768fbcb67724918936
[GDPR]	https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32016R0679&from=EN
[NIST-IR8112]	https://pages.nist.gov/NISTIR-8112/ (published Jan.2018)
[REFEDS-AF]	REFEDS Assurance Framework, https://refeds.org/assurance
[REFEDS-RS]	Research and Scholarship entity category https://wiki.refeds.org/display/ENT/Research+and+Scholarship
[RFC2256]	http://www.rfc-editor.org/rfc/rfc2256.txt
[RFC2798]	http://www.rfc-editor.org/rfc/rfc2798.txt
[RFC4512]	http://www.rfc-editor.org/rfc/rfc4512.txt
[ROB]	https://www.eugridpma.org/guidelines/robot/
[SAML2]	SAML2 core
[SIRTFI]	https://wiki.refeds.org/display/SIRTFI
[SNCTFI]	https://www.igtf.net/snctfi
[SURF]	https://www.surf.nl/en/services-and-products/surfconext/what-is-surfconext/surfconext-teams/index.html
[voPerson]	https://voperson.org/

Glossary

AARC	Authentication and Authorisation for Research Collaborations
ABAC	Attribute Based Access Control
ACL	Access Control List
API	Application Programming Interface
ASM	Attribute Schema Metadata (NIST, see [NIST-IR8112])
AVM	Attribute Value Metadata (NIST, see [NIST-IR8112])
AUP	Acceptable Use Policy
BPA	Blueprint Architecture (AARC, see [BPA])
CA	Certification Authority
CoC or **CoCo**	Code of Conduct (REFEDS, see [CoCo])
EGI	European Grid Initiative
GDPR	General Data Protection Regulation (European Union)
IdP	Identity Provider
IGTF	Interoperable Global Trust Federation (www.igtf.net)
LDAP	Lightweight Directory Access Protocol
NREN	National Research and Education Network
OID	Object IDentifier
RBAC	Role-based access control
REFEDS	Research and Education FEDerations (www.refeds.org)
RFC	Request For Comments (Internet Engineering Task Force)
SAML	Security Assertion Markup Language
SIRTFI	Security Incident Response Trust Framework for Federated Identity
SNCTFI	Scalable Negotiator for a Community Trust Framework in Federated Infrastructures (IGTF)
VO	Virtual Organisation
VOaaS	VO-as-a-Service
VRE	Virtual Research Environment
XACML	eXtensible Access Control Markup Language (OASIS)