19-11-2015

Deliverable D9.2
Market Analysis for Virtual Organisation Platform as a Service (VOPaaS)

Deliverable D9.2

Contractual Date:	30-09-2015
Actual Date:	19-11-2015
Grant Agreement No.:	691567
Activity:	SA5
Task Item:	4
Nature of Deliverable:	R (Report)
Dissemination Level:	PU (Public)
Lead Partner:	AMRES
Document Code:	GN4-1-15-76D4C
Authors:	N. van Dijk (SURFNET), M. Vermezović (AMRES), D. Pöhn (DFN/LRZ) ; A. Alper (IUCC), K. Bajnok (NIIF), A. Biancini (GARR), R. Diazmaurin (RENATER), L. Hämmerle (SWITCH), A. Harding (SWITCH), V. Nordh (SUNET), M. Prochazka (CESNET)

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 691567 (GN4-1).

Abstract
Scientific research is at the heart of every European University. Nowadays, such research is no longer an isolated activity, but has evolved into extensive collaboration between networks of researchers in multiple countries. Virtual Organisations (VOs) have emerged as the organisational form of these networks of people and resources.
The GÉANT project started a work area within SA5 - Trust and Identity Service Development with the purpose of assisting these VOs to use AAI facilities more effectively and easily for their collaborations. The proposed service portfolio GÉANT could offer for such VOs is hereinafter referred to as Virtual Organisation Platform as a Service (VOPaaS).

Table of Contents

Table of Figures

Table of Tables

Executive Summary

Scientific research is at the heart of every European University. Nowadays, such research is no longer an isolated activity, but has evolved into extensive collaboration between networks of researchers in multiple countries. With the capabilities of the internet to connect not only people but also resources, sciences have evolved into e-Science. Virtual Organisations (VOs) have emerged as the organisational form of these networks of people and resources. Broadly defined, Virtual Organisations enable groups of people to share a set of resources.

Activities aiming to support VOs to use AAI (Authentication and Authorization Infrastructure) facilities in order to share their resources were already initiated during previous GÉANT projects. Providing ready-to-use AAI infrastructure will be the next step in order to enable more VOs to share resources with the R&E end-user community by using eduGAIN. Virtual Organisation Platform as a Service (VOPaaS) is a new addition to the Federation as a Service (FaaS) product range, which is developed by SA5 task 4 as an addition to the GÉANT service portfolio and specifically aimed at Virtual Organisations. Since different VOs have varying levels of readiness for AAI uptake and therefore different requirements, the VOPaaS proposal is to deliver two levels of service: Basic and Advanced. In this way, VOs that are still in the early stages of AAI adoption can start by using the Basic services and upgrade to Advanced services as their needs and engagement grow.

1 Introduction

Virtual Organisations first emerged in areas where researchers work with resources, such as datasets and computational or experimental facilities, which are computationally intensive or require large storage. As these resources are expensive, sharing them is the only way to sustain them. In addition, scientific collaboration has grown into a global collaborative endeavour, where such resources need to be shared across the globe. The best example of this is the Conseil Européen pour la Recherche Nucléaire (CERN) Large Hadron Collider [LHC], where data is distributed to researchers globally. Other science disciplines, for example biology, and also the humanities, are now also becoming intensive data users, and sharing datasets is becoming the norm in the research field.

Authentication and authorisation infrastructures (AAI) are required to regulate who can gain access to both web and non-web based, resources. Where initially certificates were used to solve this problem within some communities, many VOs are now increasingly turning to using authentication systems based on the eduGAIN infrastructure provided by Research and Education (R&E) Identity federations for this purpose. Large VOs have already begun adopting AAI in order to meet these requirements, but not all VOs are capable of investing in this area, which often falls outside the scope of their primary aim and mission.

The GÉANT started a task “Enabling Users” in the GN3plus project with the purpose of assisting these VOs to use AAI facilities more effectively and easily for their collaborations. Based on the experience of this task, in GN4, a service offering for VOs is being developed. The proposed service GÉANT could offer for such VOs is hereinafter referred to as Virtual Organisation Platform as a Service (VOPaaS). In order to design the service offering, the work was initiated by performing the Market Analysis for VOPaaS described in this document. The objective of this analysis is to investigate the demand for services from virtual and collaborative organisations that add value to the eduGAIN inter-federation service for this market segment, as well as to make recommendations on next steps for service development in this area.

The document first looks at the use cases of several major pan-European VOs in terms of their use of federated authentication and identifies the needs and challenges they are facing. It then proposes a set of functional services that could be offered to VOs by the VOPaaS, including group management, attribute authorities, the ability to manage some attributes themselves and a functionality to enable collaborations to interact with users from outside academia.

It then analyses two service delivery models supporting a common set of requirements for advanced and long tail VOs. A model is proposed for delivering the VOPaaS platform together with the various stakeholders who form the supply chain. The benefits of providing facilities for VOs through GÉANT in terms of cost, standardisation and ease of deployment of federated AAI for VOs are examined, and cost models are analysed for different service delivery scenarios.

Finally, the market analysis sets out the conclusions and next steps by which the VO Platform can support VOs in such a way they can move from building siloed authentication and authorisation infrastructures (AAIs), as some have been doing in the past years, towards using pan-European AAI infrastructures tailored to meet the security and privacy needs of R&E.

2 Stakeholders

Collaborative efforts to enable groups to share resources via federated identity management involve many different people and organisations. Therefore, in order for GÉANT to develop and operate an effective service for collaborations, a number of stakeholders have to be considered, including:

- End users.
- Virtual Organisations.
- Home institutions.
- National Research and Education Networks.
- Identity federations.
- GÉANT Association.
- GÉANT Project.
- VOPaaS Operator, for example GÉANT Project.
- Service Providers.
- Funding Agencies.

End users include the researchers, scientists, and other collaboration partners who are members of a VO, which means they are the target end-users of VOPaaS. Typically, they need to use several services for their work. Most users have a home institution, which runs an Identity Provider (IdP) and stores user information. In order to benefit from federated identity and the VOPaaS, the user's home institution must be a member of an Identity federation and a participant of eduGAIN. A subset of end users, especially those from outside academia, may not be able to login in a federated way. Often VOs provide authentication facilities acting as a 'Guest' Identity provider to serve these users.

Virtual Organisations (VOs) are formed by a group of participants from multiple institutions, collaborating for a specific reason, e.g., a project, experiment or service. VOs may be legal entities or ERICs (European Research Infrastructure Consortiums) [ERIC], comprising several members across Europe and beyond, but may also be formed by a small group of participants, collaborating for example to write a single scientific paper. A VO enables users to share a set of resources and to use services they need for their work. The VO may have contracts with their members and various Identity federations or have informal relationships with them. A VO may run infrastructure on behalf of the collaboration effort, including Services and Identity Providers. Typically, many smaller research groups and VOs have no ability to run IT infrastructures to support their research. These groups are known as the ‘long tail’ VOs. Such VOs would already benefit from services that provide basic AAI needs. Some examples of such services are the ability to uniquely identify returning end users, a way to assign users to groups in order to delegate access rights, and a means to enable users without an identity from a home institution to login and collaborate. As a VO matures, requirements for AAI may grow so that it

may also need to maintain some additional attributes for users, support delegated group membership management, etc.

Home institutions, also known as Home Organisations (HO) in the context of Identity Federations, run an Identity Provider (IdP) for their staff and students. The IdP is typically connected to a user directory containing end user information. In the academic environment, IdPs typically support Security Assertion Markup Language (SAML). Academic IdPs are normally part of an Identity federation and can participate in eduGAIN through their Identity federation, enabling their end users to use the VOPaaS. Institutions can also run services and therefore act as Service Providers (SP).

National Research and Education Networks (NRENs) are specialised Internet service providers dedicated to supporting the needs of the research and education communities within their own country.

Identity Federations are collections of organisations operating SPs (Service Providers) and IdPs (Identity Providers) and other relevant entities that agree to interoperate under a certain rule set. In the R&E environment, they are operated by a Federation Operator who runs the processes and often provides the tools to support the operation of the Identity Federation. The Federation Operator is often but not necessarily an NREN. The majority of mature national Identity Federations in R&E are members of the inter-federation service eduGAIN.

The **GÉANT Association** was formed on 7 October 2014, when TERENA and DANTE joined forces and adopted the GÉANT name from the GÉANT Project, which continues to be a major area of the organisation's work. The GÉANT Association's Cambridge office, GEANT Limited, and its Amsterdam office, are referred to collectively as GÉANT. GÉANT is Europe's leading collaboration on network and related infrastructure and services for the benefit of research and education, It is a membership organization, owned by a core membership of 41 European NRENs, acting with and for its members to further research and education networking in Europe and globally,.

The **GÉANT Project** is a major area of GÉANT's activities that delivers innovative services to enhance the user experience, including advanced connectivity, network support and trust and identity services for NRENs, institutions and projects, and researchers and students. The GÉANT project is one of the stakeholders who will support the development and operations of VOPaaS to enable ready-to-use AAI infrastructure for VOs.

The **VOPaaS Operator** is a group of organisations in charge of infrastructure, in particular AAI facilities used by the VO. If a VO uses the VOPaaS platform, the VO Operator would be expected to be the contact point for technical and administrative issues.

Service Providers (SP) offer services for end users. These services range from general services, e.g., journals, calendar tools, etc., to targeted services such as specific datasets. In the e-Science domain, non-web applications are also important. The SP might be a member of more than one Identity federation and for international VOs should also ideally be in eduGAIN to serve users in multiple countries.

Funding Agencies, e.g., the European Commission (EC) and national funding agencies, establish certain frameworks and conditions according to which organisations may request and obtain funding. VOs often obtain funding from funding agencies for their scientific work.

3 VO Survey Methodology

A broad set of Virtual Organisations was studied in order to determine which features would be most useful for a generic platform that serves many Virtual Organisations. In the initial phase of the project, an open survey was set up allowing any VO to specify requirements and highlight the challenges they are facing. The survey was publicised to VOs at key events and on mailing lists and by direct invite to participate. In addition to this survey, the project team also conducted a study of the AAI infrastructures currently used or planned by Virtual Organisations in Europe. The study aimed to investigate both VOs of different sizes and involved in different e-science disciplines. VOs in the field of learning were also included. The requirements listed in the Federated Identity Management for research communities (FIM4R) paper [FIM4R] were also assessed. The VOPaaS project team then reviewed all the requirements to find commonalities and derive a set of generalised requirements across the various VOs.

A list of the Virtual Organisations investigated and their generalised requirements is presented in Appendix A. It is unfortunately not possible to describe all the requirements in detail here, and those set out below are not intended as a complete list. Links to the resources used are however provided for further information.

4 VO Requirements Analysis

Results of the VO survey were used to analyse and compare the VOs requirements for AAI. These were first mapped using a common vocabulary to create a list of common requirements. This list is shown in Figure 4.1, sorted by number of occurrences of particular requirement. This comprises the main result of the VO requirements survey, which included nine VOs of various sizes (see Appendix A), and is used as the main input towards the design of the VOPaaS offer.

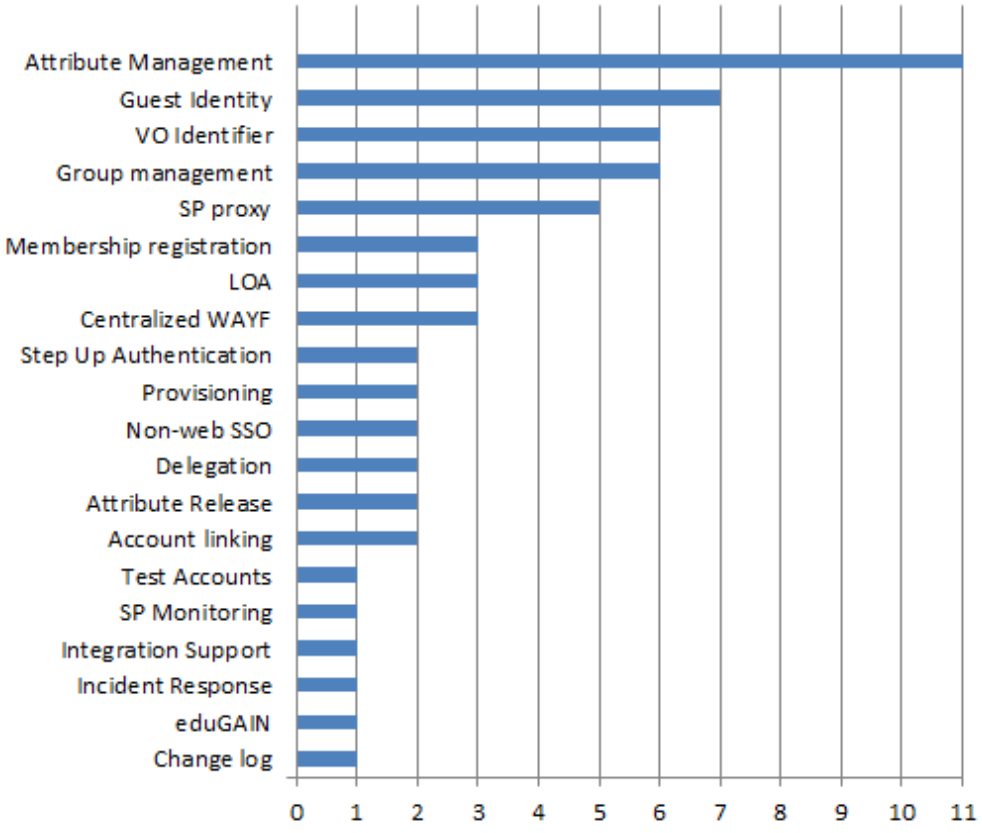

Figure 4.1: Generalised VO requirements and their frequency

5 Platform Requirements

The goal of the VOPaaS task is to identify and deliver a set of services that support Virtual Organisations. During a two-day workshop with experts from various NRENs and Identity Federations, the use cases described in Appendix A were discussed, grouped and prioritised. As a result, a shortlist of functional requirements that could be supported by a VOPaaS platform was drawn up. This section describes the capabilities that the VOPaaS offer should provide.

5.1 Platform Authentication and Authorisation

All services and interfaces should be protected against unauthorised use. To this end, the VOPaaS offer itself should be protected by an authentication and authorisation service that allows the operations team to delegate authorisation. All user interfaces of the services offered towards VOs must be web-based and available through eduGAIN and an Authentication Proxy for External Identities (extIDp) (see section 5.4).

5.2 Persistent Identifier

All of the VOs in the study are interested in using federated identity management as the basis for their AAI systems. By using federated AAI, they can allow end users to log in to VO services using their home institution account. End users may move from one institution to another, which changes their home institution identity, but may need to retain their membership status with the VO. The VOs therefore need a system to maintain a coherent set of information on their users irrespective of the authenticating systems they use. This can be done by creating a VO-specific, persistent pseudonymous identifier for VO members at the time they are onboarded. This VO-specific pseudonymous identifier is scoped as the basis for all AAI-related information in the platform, allowing a user's VO membership to be preserved.

5.3 VO Membership Registry

All of the VOs interviewed expressed a requirement for member registration. The various VOs have very different workflows for onboarding their members which were not all investigated by the study. However, earlier work carried out by the COmanage project [COmanage] suggests that a number of common strategies exist. The creation and automation of workflows for onboarding new members should be supported through the use of templates. A tool operated by the VOPaaS platform for Virtual

Organisations should allow them to use and manage these workflow templates themselves, without any interference from the VOPaaS operator. The VOPaaS operator should only need to delegate access to VO representatives, who can then work with the tools themselves.

A record should be kept for each onboarded user containing:

- Identifiers for the user obtained from the authentication sources.
- The authentication source used for authentication.
- An identifier of the VO in question.
- The workflow used for onboarding.
- The user's persistent identifier with the VO.
- The date and time the user was onboarded.

Additional implementation considerations include:

- The record should only contain the personal data that is strictly necessary for use of the VO services.
- The record should be made available using common application programming interfaces (APIs), which can be queried by services when a user logs in.
- Only services that have proper authorisation should be able to query the data of users of a specific VO.
- The service should facilitate account linking, allowing a user to connect a new authentication source to the persistent identifier.
- An out–of-band mechanism should exist to allow this account linking to take place even if the original authentication source is no longer available to the user.

5.4 External ID Provider (extIDp)

Pan-European VOs leverage eduGAIN to enable end users to authenticate using their home institution identity and use services outside their national boundaries. Unfortunately, participants of these VOs may not have an eduGAIN-ready Identity Provider. There could be three reasons for this:

1. The user's home institution has an Identity Provider in an Identity federation, but it has not yet been made available via eduGAIN.
2. The home institution may be academic in nature, but not have an Identity Provider.
3. The end user is not a member of the academic community, and therefore cannot authenticate via Identity federation and/or eduGAIN.

The first two issues are already being addressed by other GÉANT project activities. However, the third case is not addressed on a pan-European service level, and remains the number one challenge mentioned by most VOs. Some Identity Federations are providing Guest IdPs [Guest IdP], but these are often only available locally and fall outside the scope of the Federations' policies to facilitate academic collaboration, due to their perceived lower level of trust. The VOPaaS could provide a pan-European service for dealing with these kinds of end users across Europe for Identity Federations and virtual organisations alike.

To offer such a service, three challenges must first be addressed:

- Storage of personal data will need to be minimised: Running a centralised service that would potentially store data from end users all over Europe is a major challenge from a data protection perspective. One solution is to operate it as a transaction service, connecting various external identity sources such as social providers (e.g. Google, Facebook, etc.) but also Bank IDs or STORK (Secure idenTity acrOss boRders linKed) [STORK], then pass these on towards the consumed services. These transactions should be atomic and not store any personal data at any time. Attributes should only be passed on where the authentication source allows it and only where an end user has consented to it.
- Account persistence: The service should provide a single persistent pseudonymous identifier for each user for the consumed services. It should allow multiple authentication sources to be linked to this identifier as a user might log in using different sources. Checks should be performed periodically to ensure that the user is in control of the authentication sources linked. An out-of-band mechanism should also be in place to allow account recovery if the original authentication source is no longer available to the user.
- Assurance: The different external sources used for authentication vary in assurance quality. It would therefore be useful for the Service Provider requesting authentication to have relevant assurance criteria assigned to an authentication source by trusted parties. The challenge here is what definitions to use, as every VO may have different assurance definitions. The service should therefore use a well-established Level of Assurance (LoA) scheme, such as e.g. eIDAS (electronic identification and trust services) [eIDAS], as its default schema, but also allow the VOs themselves to select an alternative assurance criterion, provided it is registered in a LoA registry [LoA].

5.5 Group Management

All VOs need to manage groups used for both authorisation of their services and collaboration scenarios. A VO must be able to register groups and optionally per-group additional attributes. Such groups and attributes must be managed by the VOs themselves in a group management system. Depending on requirements, it should be possible to have users manage their groups directly, or have these managed centrally by technical persons within the VO. To allow services to learn about the groups, group memberships and group attributes, the group management system(s) should provide common, open standards-based interfaces (APIs) to allow the attributes to be queried. A mechanism must exist to allow only authorised services to consume group information of specific users.

Several products already exist that can provide the features mentioned above. Although technically similar, these products offer different approaches to how group membership information is collected, managed, and can be provisioned towards services. The VOPaaS platform should not be a generic hosting platform for collaboration tools but rather provide (a set of) selected specific tools, such as for example COmanage, HEXAA or Perun, that users understand and have a proven track record. The VOPaaS offer could leverage diversity, allowing VOs to choose the tools that best suit their needs, while at the same time providing hosted versions of the most commonly used tools to relieve VOs of the burden of installing and maintaining these products.

To support the long tail, and allow for 'ad hoc' collaborations, a simplified group management service could complement services that support VOs with more advanced use cases. The basic service should only allow the creation and management of 'flat' group membership, whereas an advanced service should also allow hierarchical groups and group-specific attributes (see Section 6 Service Delivery).

5.6 Attribute Management

All VOs typically want to manage authorisation for their services. As these services are provided in the context of the VO, attributes from the home institution alone are often not enough to allow this. As well as using groups (group membership), a VO should be able to register additional attributes, e.g., roles for the users in its collaboration, managed by an attribute management system. The users could add the additional attributes themselves, or the VO could operate an additional vetting process. Additional VO-specific attributes have to be stored in an Attribute Authority(s), and can then complement the attributes provided by the home institution. To allow services to learn about these additional attributes, the Attribute Authority (AA) should provide common, open standards-based APIs to allow them to be queried. A mechanism must exist to allow only authorised services to consume the attribute information of specific users.

Several products already exist that can provide the above features. Although technically similar, the products offer different approaches to how the attributes are collected, managed, and can be provisioned towards services. The VOPaaS offer could leverage this diversity, allowing VOs to choose the tools that best suit their needs, while at the same time providing hosted versions of the tools to take away the burden of installing and maintaining the products. A basic product offering would provide a pre-selection of tools for the less complex use cases, to remove the burden of initial choice.

5.7 Provisioning

Ultimately all relevant attribute and group information has to be provided to services. This process is called provisioning. For basic services, the VOPaaS platform should offer a limited set of default processes to allow provisioning to happen. It should also support 'just-in-time' provisioning using well-understood mechanisms that are standardised and known to be scalable, such as providing attributes as part of a SAML-based authentication.

The basic product offering would not provide capabilities for centralised deprovisioning, because of the nature of the proposed protocols. Conversely, advanced (de)provisioning mechanisms, including application-specific interfaces and non-web scenarios could be provided as part of the advanced services that the platform offers (see Section 6 Service Delivery).

5.8 Service Proxy and Attribute Aggregation

As stated earlier, all services should by default provide APIs that allow services, if authorised, to query the information stored in the platform. This enables the creation of a mesh model, where several services will interface with multiple services in the VOPaaS platform. The mesh model provides advantages in terms of privacy and security, but can also introduce various technical and

organisational challenges, depending on what software is being used by the service and the trust that exists between services and the VO.

A common way of dealing with this complexity is to introduce a (SP) proxy, which can handle authentication and attribute aggregation on behalf of connected services. Given the centralised nature of the VOPaaS platform, VOs and services must already have a certain level of trust in the operator of the platform. VOs should be able to choose whether they want to make use of a VO-specific proxy to handle transactions for them.

5.9 Collaboration Services

An AAI middleware platform cannot operate without connected services. The VOPaaS team should therefore identify a number of generic collaborative services currently provided in the NREN community and collaborate with its operators or software developers to make sure the platform can be used with these services. These are typically file sharing, wikis and other collaboration tools, and may include Foodle, Rendez-Vous, Filesender or similar services. Commercially provided services could also form part of this offering, possibly in collaboration with the GÉANT Project SA7 (Supply Chain Support) activity. As well as enabling these services for users, this would also provide a good test case to ensure that the platform can perform basic functions towards these services. If no services with a suitably pan-European access policy can be found, the VOPaaS platform should consider proving some basic collaborative services itself in conjunction with the AAI platform.

5.10 Non-Web-Based Resources

Many resources that VOs share cannot easily use web-based authentication mechanisms. This is the case for example with the commonly used SSH access to storage and compute resources. This challenge is being addressed on many fronts.

Many resource providers have adopted pragmatic solutions for this, for example by using a web-based portal, such as the portal of the We-NMR project [We-NMR]. However, as such portals are often closely bound to the resource provided, this approach is hard to generalise.

In recent years, the issue has been increasingly approached at an infrastructural level, including by using support for the SAML ECP (Enhanced Client or Proxy) profile [SAML ECP] in the Shibboleth software product [Shibboleth], and with project "Moonshot" [Moonshot]. While these technologies have been shown to work well for specific use cases, the biggest challenge is that currently neither technology is widely adopted. Other GÉANT tasks are actively working on this issue but adoption at campus and service level is expected to take some time.

With none of the above approaches for non-web based resources being widely available, there is no clear path ahead for the VOPaaS platform to provide a generalised service that could be adopted by a large majority of the use cases identified in the survey and by the FIM4R requirements. In view of this, and the fact that any non-web solution requires an infrastructure to provide authentication and authorisation, provision of non-web resources is not on the immediate roadmap.

5.11 Data Protection and Security Considerations

Offering AAI-related services involves great responsibility in terms of data protection and security. As it is expected that the majority of the platform's users will be from Europe and that it must be hosted with the EU, all of the services it offers should adhere to the GÉANT Code of Conduct adopted within eduGAIN, and to any additional EU privacy and data protection laws not reflected in the Code. Adequate security measures must also be taken to protect the uses of the data in the platform.

6 Service Delivery

The VO survey encompassed VO use cases with established AAI of different maturity levels. This highlighted the need for VOPaaS to not only provide a service for the VOs that are already working with, or towards adopting AAI infrastructures, but also for the longer tail. It is expected that the latter VOs will be less familiar with AAI concepts and therefore may prefer a less complex scenario to start with, but that their needs will grow as they understand the potential of adopting these infrastructures. The VOPaaS should therefore offer services for both VOs with modest as well as advanced AAI requirements.

Two service levels, Basic and Advanced, could therefore be delivered. A Basic service offering could combine some of the essential platform requirements described in Section 5. As such, it would only include functionalities that can be readily understood by VO users who are not familiar with this type of service. An Advanced service offering could provide services for VOs that already have better capabilities for AAI, enabling them to use all available features.

These two service delivery scenarios are not in competition but rather complementary as part of the VOPaaS offer, following a logical progression from supporting basic use cases to more advanced scenarios.

Table 6.1 shows how platform capabilities could be provided in the Basic (Section 6.1) and Advanced (Section 6.2) scenarios.

Platform capability	Basic services	Advanced services	Deployment tenancy	Data ownership
Persistent Identifier	Y	Y	Multi	VOPaaS Operator
VO Membership management	Y	Y	Multi	VOPaaS Operator
External ID provider	Y	Y	Multi	VOPaaS Operator
Basic Groups	Y	-	Multi	VOPaaS Operator
Basic Provisioning	Y	-	Multi	VOPaaS Operator
Advanced Groups	-	Y	Single	VO operator
Attribute Management	-	Y	Single	VO operator
Advanced (de-)Provisioning	-	Y	Single	VO operator
SP proxy, attribute aggregation	-	Y	Multi	VO operator

Table 6.1: Platform capabilities in Basic and Advanced service scenarios

The two scenarios also have different deployment tenancy and legal data ownership implications, which are discussed in the next sections. In both scenarios, the VOPaaS operations team would provide the underlying infrastructure and software packages.

6.1 Basic Services

Basic services should be operated by GÉANT for the benefit of all users of identity federations that have their IdPs in eduGAIN. Many eduGAIN members do not have any collaborative platform in place, or mainly focus on national collaboration. A platform offered by GÉANT would truly enable pan-European collaboration beyond borders. Furthermore, as the external ID provider could allow users from outside the Research and Education community to access connected Services, it would enable collaboration across different sectors. Collaborations could be formed through a basic group service providing easy-to-use interfaces for creating, managing and querying group memberships.

6.1.1 Workflow Example

This workflow provides an example scenario of how the proposed Basic service offering would help a VO. The scenario hypothesises three types of end users:

- A VO operator who is in charge of managing membership of VO users;
- A user called “Alice”, who is a VO user who can login using a home institution account; and
- A user called “Bob”, who is a VO user who has no home institution.

In order to complete this scenario successfully, the following tasks need to be accomplished:

1. Alice and Bob need access to a wiki service,
2. the wiki service requires authorisation in the context of the VO as only members of the VO are allowed in,
3. the VO operator has been given the authorisation to manage members in the *VO membership management service*,
4. Alice has the role of "wiki space admins" within the VO which makes her the manager of a space in the wiki service, and she has been delegated the ability to manage members of the "wiki space editors" group,
5. Bob becomes a member of the "wiki space editors" group and is able to edit the content in the wiki service.

To execute this flow, the VO operator first needs to be able to make the Alice and Bob members of the VO. To this end, the VO operator uses the *VO Membership management service* to make users members of the VO. He/she could for example perform this onboarding by sending out (email) invites to the users. In response to the invite, the users should log in to the *VO Membership management service* and provide some basic attributes.

It is assumed that Alice will log in using a federated login authenticated via her home institution. Bob has no home institution, and therefore uses an *External ID provider* to log in with his Google ID. The

VO operator evaluates the information received about the users and makes them members of the VO. As they have membership, both users are assigned a *Persistent Identifier,* which is specific for this VO.

With the *Persistent Identifier* assigned to them, both users can now log into services provided within the context of the VO, however. However, on this basis alone they are not authorised to use any functionalities. Therefore, to determine whether they have any roles within the context of the VO, the wiki services queries the *VO Membership management service* and *the Simple Groups service.*

To assign Alice the "wiki space admin" role, she needs to be added to the appropriate group. This group is managed with the *Simple Groups service* by the VO operator. After being assigned group membership of the "wiki space admins", Alice logs into the *Simple Groups service*, and this service now allows Alice to add Bob to the "wiki space editors" group.

When Bob logs into the wiki service, the wiki service queries both the *VO Membership management service* to get Bob's *Persistent VO Identifier* and the *Simple Groups service* to find that Bob has the correct role, "wiki space editors", which allows him to edit content in the wiki space. To store the roles of Alice and Bob, the wiki service may use the features for *Simple Provisioning* provided by the Basic services.

6.1.2 Stakeholder Engagement

For the Basic scenarios, use of the VOPaaS offer could be routed through the NRENs or their respective Identity Federations, as shown in Figure 6.1. This means that an NREN can deliver a service by sponsoring a VO to enable it to gain access to the VOPaaS infrastructure, while at the same time supporting the VO and helping build local expertise via normal NREN and federation channels.

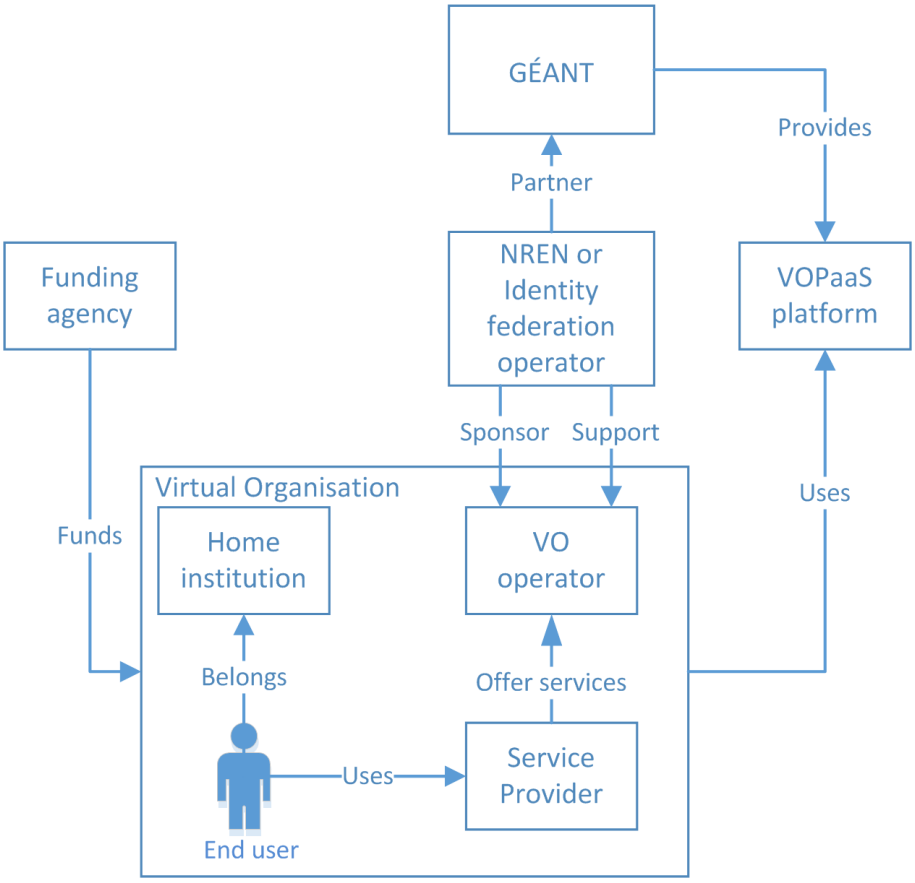

Figure 6.1: Stakeholder engagement in Basic service scenario

6.1.3 Deployment Tenancy

From a deployment perspective, all of the platform capabilities in the Basic service offer can be deployed as a multi-tenant service. A multi-tenant service can support many VOs within one functional application. Operating the Basic service as a multi-tenant service offers a number of benefits, most importantly a single point of interaction for Identity Providers.

Running the Multi-tenant services would require resources (VMs, operations team), as well as support towards Identity federations. End user support should be kept to the minimum essential, as it is preferred to respect existing NREN and federation support channels.

6.1.4 Data Ownership

It is expected that for the Basic services, the VOPaaS operator will decide on the functional capabilities of the service being offered. The VOPaaS operator will also be expected to take responsibility for the legal requirements that are placed on such a platform in regard to, for example, data protection and privacy for these services. Services deployed to support the Basic scenario should be the (legal) responsibility of the VOPaaS operator.

6.2 Advanced Services

For the Advanced use cases, it is likely that the VO will be able to act more independently. While the service could still be routed through the NRENs, direct interaction with the VOPaaS operator is also feasible. Services provided for advanced use cases are likely to contain more and VO-specific data. On a functional level, Advanced services would be managed by technically skilled staff of the VOs themselves.

6.2.1 Workflow Example

This workflow provides an example scenario of how the Advanced services would help a VO. This scenario will also make use of many of the features provided by the Basic services. The scenario hypothesises the same three end users, the VO operator, Alice, and Bob, described for the Basic workflow example (6.1.1).

In this this scenario, the following tasks need to be accomplished:

1. Alice needs access to a scientific database service
2. The scientific database service requires authorisation in the context of the VO as only members of the VO are allowed to access the service
3. The scientific database service requires permission from a data access commission to allow the VO user to access specific datasets
4. Alice wants to perform a computation on a dataset. To do so, she needs access to the web portal of a GRID compute facility where she can upload the dataset
5. As part of policy of the GRID compute facility, Alice has to provide the number of her scientific grant so that resource allocation can be charged accordingly
6. For the compute facility to actually perform the calculations, it needs to provision Alice's credentials in the form of a certificate
7. The results of the calculation is downloaded by Alice via the GRID compute facility web portal and Alice shares the results with Bob in the wiki service described in the Basic scenario.

For this scenario, it is assumed that Alice is already a member of the VO as this was accomplished with the tools provided by the Basic services. In addition, access to the scientific database service is set up using a similar process to that described to grant Alice access to the wiki service in the Basic Workflow scenario.

To proceed, Alice now first needs a copy of a specific dataset. She logs into the scientific dataset service and requests access to the dataset. Based on the *Persistent Identifier* assigned to Alice, the scientific database service queries a VO specific *Group management service* to find out if Alice is a member of the group of people who have been granted access to that specific dataset by the relevant data access committee. Such a committee would typically evaluate this type of request via an out-of-band mechanism, where the result of the decision then gets stored in a *Group management service*. For the time being, it is assumed access has been granted and additional group membership information on Alice is collected and used to allow her access to a specific scientific database. Note that in this example a *Group management service* was used, however, an *Attribute management service* would

technically have accomplished the same thing. A real-life example of such a service was developed within the Elixir Community [ELIXIR].

Alice downloads the dataset and wants to perform some computations on it. To do so, she accesses the web portal of a GRID compute facility. To be able to provide the GRID compute service with the required number of her scientific grant, Alice has to add this information to her VO specific profile. To do so, she logs into a VO specific *Attribute management service,* where she is able to link this attribute to her identity. She also grants the *Attribute management service* the right to pass the grant number to the GRID compute facility on her behalf. This does not have to be repeated for further requests.

When she logs into the portal, the portal queries the *VO Membership management* and the VO specific *Attribute management service.* The GRID compute facilities have established that Alice is a member of the VO and have been provided her grant number via the *Attribute management service*. With this being made available, Alice now has permission to upload datasets to the compute facility.

In order for access, compute cycles, and disk space to be allocated for Alice's dataset, Alice needs to provide a certificate for it. She downloads this certificate using a *provisioning service,* which has converted her SAML-based Web SSO identity into a certificate, and uploads the certificate alongside with the dataset.

Once the computation is completed, the resulting data is downloaded by Alice from the portal of the GRID compute facility. A real-life example of this type of service is used in WE-NMR project [We-NMR]. The dataset is linked to Alice´s *Persistent Identifier*, so only she has the right to download the dataset. As she wants to discuss the result with Bob, she annotates the results and puts them on the wiki service Alice and Bob share (see the Basic service workflow example at 6.1.1).

6.2.2 Stakeholder Engagement

For the Advanced services, the VOPaaS operator would deal with the VOs directly, as shown in Figure 6.2. This requires a legal entity on the side of the VO, as often found in larger, cross-national VOs, such as for example in the EU language research infrastructure CLARIN-ERIC [ERIC]. Once the legal status of the VO is established, the VOPaaS operations team grants representatives of the VO direct access to the VOPaaS services. Depending on the services used, a fee may be required and the VO could be charged directly.

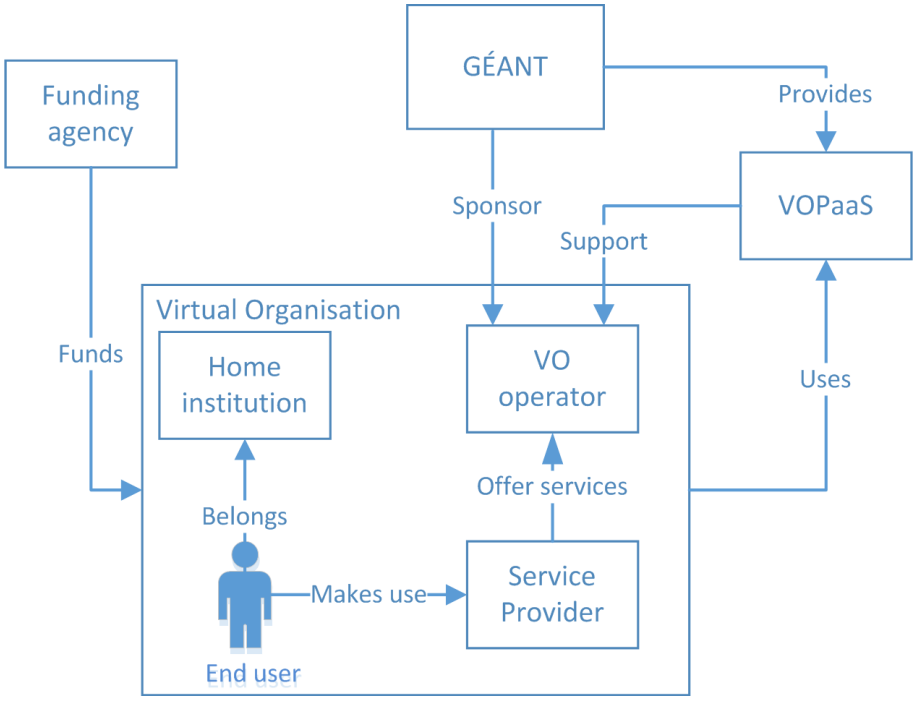

Figure 6.2: Stakeholder engagement in advanced service scenario

6.2.3 Deployment Tenancy

Advanced scenarios are expected to be better serviced using a single tenant per VO model. In such a scenario, a single instance of a service is operated by the VOPaaS platform, specifically for a VO. Modern virtualisation technologies, such as VMs and containers, can be used to reduce the operational complexity of this type of setup. While this is more complex from an operational perspective, it allows the VO to have more control of the functionality's capabilities. In addition, it makes operational sense to use VMs or containers to run a separate instance of the service as in this way, the appropriate representative of a VO would be able to authorise it to request resources from the VOPaaS operations team, which may come with a cost.

6.2.4 Data Ownership

From a legal perspective, it is recommended that the VO that takes responsibility for data protection and privacy. A similar model is currently applied by GARR for their IdP as a Service model.

6.3 Other Considerations

6.3.1 Interoperability and Open Standards

The VOPaaS platform will only use Open Source, well-known and existing products. The use of Open Source tools is required as it reduces the risk for the VOPaaS platform operations team of encountering any problems with integration, as documentation and source are available for the products at no extra cost. It also prevents lock-in by allowing VOs to move away from the

infrastructure provided by the VOPaaS and set up their own platforms with the same tools should they opt to do so. Appropriate incoming IPR checks on open source licenses will be carried out as per the GÉANT IPR policy.

6.3.2 Service Management and Maintenance

To ensure the products are maintainable and secure, installation, maintenance and software issue management requirements for the VOPaaS platform tools will need to be defined. At the same time, to ensure quality of support, the platform should help sustain the products that are in use by supporting software development teams that maintain the software, either in kind or financially. This cost could be charged to the VOs using the products in the advanced model where customisation may be required, or the GÉANT platform could employ a sponsor model so that costs only have to be partially charged where use by the VO is considered strategically important as part of a wide package.

6.3.3 Platform Branding

The collaborations that this study targets are often called “Virtual Organisations”, where “Virtual” refers to the fact that the organisation is not the "brick and mortar" institution of a researcher. For most researchers, however, there is nothing “virtual” about their engagements with these collaborative organisations. Although Virtual Organisations are most commonly associated with science projects, this study shows collaborations in the field of learning have very similar characteristics. When deploying a platform branding the platform should take this wider target market into account. One option is to name such collaborations "Collaborative Organisations" (COs) which is understood to be wider in scope than Virtual Organisations.

7 Conclusions and Next Steps

As the result of this Market Analysis, the following conclusions and recommendations for the VOPaaS offer can be derived:

- Two co-existing service offerings should be implemented within the VOPaaS platform, a Basic and an Advanced offer.
- Appropriate cost recovery mechanisms should be developed for both service offerings.
- Security and Data protection considerations should be included into the design of VOPaaS.
- Basic services should:
  - include the following services: VO Membership management, External ID provider, Basic Groups, and Basic Provisioning;
  - allow end users of eduGAIN members to be able to login;
  - include simple operations provided by GÉANT and offered to users at no additional cost.
- Advanced services should:
  - include the following services: Advanced Groups, Attribute Management, Advanced (de-)Provisioning, SP proxy, Attribute Aggregation;
  - be available to VO operators and end users authorised by the VO;
  - include operations and consultancy provided by GÉANT, with appropriate costs charged to the VO.
- Services should be deployed using a phased approach, beginning with the Basic services and progressing to the Advanced services.
- Pilots should be run for the services in close collaboration with communities, such as AARC, FIM4R, RDA, ESFRII etc., as well as with global initiatives and partners such as MAGIC and InCommon.

Further work on the VOPaaS offer should follow the general work plan outlined below:

- Design the Basic service offering (GN4 Phase 1).
- Prepare Cost Benefit Analysis for the Basic service offering, setting out who can use the service and cost models (GN4 Phase 1).
- Setup and pilot Basic service offering in collaboration with stakeholder communities (GN4 Phase 1).
- Hand over Basic service operations and support to appropriate GÉANT tasks. (begin in GN4 Phase 1, continue in GN4 Phase 2).

- Once the Basic service offer is ready for to pilot, begin designing the Advanced service offering (GN4 Phase 1).
- Prepare Cost Benefit Analysis for the Advanced service offer, setting out who can use the service and cost models (GN4 Phases1 & 2).
- Set up and pilot Advanced service offering in collaboration with stakeholder communities (GN4 Phases 1 & 2).
- Hand over Advanced service operations and support to appropriate GÉANT tasks (GN4 Phase 2).
- Promote the VOPaaS offer at events targeted at the VOPaaS stakeholder community, including AARC, FIM4R and similar groups and projects (ongoing activity, begin in GN4 Phase 1).
- Facilitate the uptake of services resulting from non-GÉANT projects to extend and enhance the VOPaaS in the future (ongoing activity, begin in GN4 Phase 1).

Appendix A VO Use Case survey

A.1 FIM4R

Identity Federations provide economic advantages, as well as convenience, for collaborative organisations and research infrastructures and their users. In order for Federated Identity Management (FIM) to be effective, the partners must have a sense of mutual trust. A number of laboratories, including national and regional research organisations, are facing the challenge of a deluge of scientific data that needs to be accessed by expanding user bases in dynamic collaborations that cross-organisational and national boundaries. Driven by these needs, representatives from a variety of research communities, including photon/neutron facilities, social science and humanities, high-energy physics, atmospheric science, bioinformatics, and fusion energy, have come together to discuss how to address these issues with the objective of defining a common policy and trust framework for Identity Management based on existing structures, Identity federations, and technologies. These communities organised a series of workshops, which resulted in the FIM4R paper.

The paper [FIM4R] discusses many aspects of federated identity management, including technical and policy issues. A summary of the requirements is presented in Annex A of the document.

A.2 Umbrella

Purpose of the VO:

Exciting opportunities are emerging for large photon / neutron facilities due to novel developments in accelerator and detector techniques. In order to optimally exploit these opportunities, the experimental environment has to be expanded from the traditional single-facility orientation to a pan-European scale by embracing remote access to experimental data or to experiments. The Umbrella system was developed as a community-wide federated authentication and authorisation system to enable access to such services.

Umbrella is supported by a series of projects, including the FP7 (Seventh Framework Programme) projects EuroFEL (Free Electronic Lasers for Europe), PaNdata (Photon and Neutron data, Europe and ODI) [PaNdata], and CRISP (Cluster of Research Infrastructures for Synergies in Physics) [CRISP].

Particularities of technical implementation:

Umbrella provides unique and persistent EU-wide end user identification and allows the implementation of a pan-European Single Sign-On mechanism. It is a hybrid system in the sense that only the minimum data needed for an unambiguous identification is stored centrally, while the remainder of data and data handling is as much as possible left to the local user offices [Umbrella].

Umbrella is based on SAML with one (1) identity provider. This central element contains just the minimum sufficient information for user identification. The remaining authentication information, and all authorisation information, remain at the local user offices of the facilities. A multi-level trust concept provides the necessary flexibility. Wherever possible, self-service elements and responsibility delegation are applied in order to minimise the administrative load and legal restrictions. In view of this self-service approach, user friendliness is an important issue.
Umbrella carried out various studies to identify the best way to set up a platform to meet their needs. These were taken as input for this study. At the time of writing Umbrella is operating a lightweight infrastructure based on the requirements defined in their studies.

Generalised requirements:

- VO specific persistent identifiers;
- ability to 'switch' from old to new Home institution;
- split between identifier and (delegated/distributed) attributes;
- central authentication with distributed authorisation;
- record users' affiliation with home institution,
- allow users to add additional attributes.

A.3 CLASSe

Purpose of the VO:

CLASSe (Cloud-ABFAB Federation Services in eduroam) is a GÉANT OpenCall project, which in its first phase investigated the use of the ABFAB (Application Bridging for Federated Access beyond Web) technology for cloud services. For its collaborative needs, the CLASSe project needs to share a number of resources between institutions in the UK and Spain.

Particularities of technical implementation:

CLASSe uses the integration of Moonshot into OpenStack as a reference implementation in the context of standardisation efforts. The suitability of the Moonshot/OpenStack prototype implementation produced by the University of Kent prior to CLASSe starting is validated for use over GÉANT, and user guides have been produced for its deployment and installation. Subsequently, CLASSe tested solutions to improve traditional Single Sign-On (SSO), and research on solutions of real SSO to provide access to cloud services without further authentication in the eduroam network (cloud-to-cloud SSO and network-to-cloud SSO). In its final stage, CLASSe researched and designed solutions to support the dynamic formation of communities of interest (CoIs), so that any users are able to form their own CoIs for access to their own private cloud services. This negates the need to have a separate Virtual Organisation Membership Service as is necessary in today's grids.

Generalised requirements:

- centrally managed group membership with delegated management;
- good integration points with existing applications;
- guests accounts used for authentication only;
- authentication (federated account, guest account) separate from VO membership management;
- implicit and explicit (internal/external), ad-hoc, and institutional groups.

A.4 DARIAH

Purpose of the VO:

DARIAH, the Digital Research Infrastructure for the Arts and Humanities, aims to enhance and support digitally enabled research and teaching across the humanities and arts [DARIAH]. DARIAH aims to develop, maintain and operate an infrastructure in support of information and communications technology (ICT)-based research practices. By working with communities of practice, DARIAH-EU brings together individual, state-of-the-art, digital Arts and Humanities activities across Europe. It strives to preserve, provide access to and disseminate research that stems from these collaborations and ensure that best practices, methodological, and technical standards are followed. The final DARIAH-EU infrastructure will comprise a connected network of tools, information, people and methodologies for investigating, exploring and supporting research across the broad spectrum of the Digital Humanities.

Generalised requirements [DARIAH Req]:

- non-WebSSO (ECP, Enhanced Client or Proxy);
- use a central LDAP (Lightweight Directory Access Protocol) for their collaboration connected to JIRA and Confluence;
- LDAP sync to local Home institutions/services;
- allow external SAML IDs and guest IdP;
- users are being provisioned into their central LDAP;
- implemented attribute aggregation;
- centralised WAYF (Where Are You From?);
- guest IdP.

A.5 CERN

Purpose of the VO:

Founded in 1954, the CERN laboratory sits astride the Franco-Swiss border near Geneva. It was one of Europe's first joint ventures and now has 21 member states [CERN]. At CERN, the European Organization for Nuclear Research, physicists and engineers are probing the fundamental structure of the universe. They use the world's largest and most complex scientific instruments to study the basic

constituents of matter – the fundamental particles. The particles are made to collide together at close to the speed of light. The process gives the physicists clues about how the particles interact, and provides insights into the fundamental laws of nature.
The instruments used at CERN are purpose-built particle accelerators and detectors. Accelerators boost beams of particles to high energies, before the beams are made to collide with each other or with stationary targets. Detectors observe and record the results of these collisions.

Particularities of technical implementation:

CERN already has a long history of providing collaborative tools and access to scientific resources. In view of this, it has already developed various tools internally. For most of these authentication was based on x.509 certificates, but recently SAML-based Identity federation is also being increasingly used.

Generalised requirements:

- working/Mapping x.509 certificates <-> Shibboleth (SAML) credentials;
- attribute management (release, aggregation, LoA (Level of Assurance);
- guest IdP;
- proof of identity (LoA) for authentication;
- IdPs in incident response;
- the ability to step-up LoA;
- centrally managed groups and VO attributes.

A.6 CLARIN

Purpose of the VO:

The ultimate objective of CLARIN-ERIC is to advance research in humanities and social sciences by giving researchers unified single sign-on access to a platform, which integrates language-based resources and advanced tools at a European level [CLARIN]. This will be implemented through the construction and operation of a shared distributed infrastructure that aims at making language resources, technology, and expertise available to the humanities and social sciences research communities at large.

Particularities of technical implementation:

CLARIN was one of the first Virtual Organisations to attempt to fully deploy pan-European federated identity for their collaboration. As eduGAIN was then in its very early days, CLARIN created its own “SP federation” without leveraging eduGAIN. This involved the need for a contract with every Identity Federation, and in some countries with every relevant IdP. To make that process more scalable, CLARIN deals with Identity Federations contract in a centralised way on behalf of SPs that are members of their VO.

Generalised requirements:

- need for test accounts;
- ability to actively monitor connections;

- virtual Home organization / Guest IdP;
- custom tailored Discovery Service;
- release of a basic set of attributes: name, mail, identifier, home organization + affiliation;

A.7 Virtual Campus Hub

Purpose of the VO:

The Virtual Campus Hub project ran from 2011 to 2013 and was one of the first Collaborative projects in eLearning which tried (and succeeded) to leverage eduGAIN as a way to connect distributed Learning Management Systems and Institutions across Europe [VCH].

The objective of the Virtual Campus Hub (VCH) project was to develop and implement the tools and e-learning platforms needed to establish a European, and potentially worldwide, Virtual Campus network primarily for technical universities and business schools. The Virtual Campus network used the European e-infrastructure network, including GÉANT as the communication backbone. The project formulated end-user demands for high-quality services in support of a global virtual campus network based on a VCH concept. The Virtual Campus Hub was developed through pilot use of the hub elements with special emphasis on the integration of research, innovation, and education in sustainable energy.

Virtual Campus Hub received funding from the European Union's Seventh Framework Programme for research, technological development and demonstration under grant agreement No. 283746

Generalised requirements [VCH Req]**:**

- service provider proxy;
- centralised WAYF;
- guest IdPs;
- use of eduGAIN;
- centralised group management;
- a Service Provider Hub to connect Services in one central location.

A.8 GÉANT AuthZ Management System

Purpose of the VO:

GÉANT is Europe's leading collaboration on e-infrastructure and services for research and education [GÉANT]. The GÉANT project is, in essence, a VO, whose participants are spread across different organisations at different geographical locations and need access to common services. In the past, to access those services GÉANT users needed different sets of credentials. Thanks to federated identity, participants can now log in to most of these services using the credentials issued by their home organisation. This addresses the authentication issues very well but is only half the story, as Identity

Management is incomplete without proper authorisation management. Due to the size and complexity of the project, participants can have different roles and belong to different work tasks. As such, participants can have different permissions for accessing shared resources. To be able to manage this efficiently, it is necessary for the GÉANT project to perform group management in a scalable manner.

Generalised requirements:

- workflow to register members;
- (delegated) group management;
- (delegated) authorisation management;
- local groups and local authZ (authorization) for applications;
- groups within groups.

A.9 ELIXIR

Purpose of the VO:

ELIXIR is a pan-European research infrastructure for biological data [ELIXIR]. It unites life science organizations in managing and safeguarding the amount of data being generated every day by publicly funded research. ELIXIR is an ESFRI (European Strategy Forum on Research Infrastructures) research infrastructure building a secure, evolving platform for biological data, collection, storage, and management, consisting of an interlinked set of core and specialist resources.

Particularities of technical implementation:

The ELIXIR AAI Pilot [ELIXIR Pilot] for distributed authentication of the European Genome-phenome Archive was first launched in 2012. It enables governance processes and access to personal genome data through institutional logins and automated interaction with appropriate data access committees. Further pilots are planned for spring 2015, while the deployment is planned for the ELIXIR Excelerate [ELIXIR Excelerate] project. Besides the IdPs available via eduGAIN, which have already been successfully piloted and will be extended, support for identity sources such as Google and ORCID (Open Researcher and Contributor ID), is planned in the Exelerate project. This means, users should be able to link different accounts, e.g., the ELIXIR account with the eduGAIN IdP account, the Google account and the ORCID account. In order to meet biomedical data requirements, these accounts should be assigned assurance levels. One way of increasing the level of assurance is step-up-authentication.
The ELIXIR Proxy IdP will support eduGAIN IdPs and other common IdPs, and should be used as a hub to access ELIXIR services. User information can be enriched by the ELIXIR Directory, which has interfaces with authorisation management, group management, Bona fide status management, and attribute self-management.

Generalised requirements:

- persistent identifier for account linking;
- institutional IdPs and common IdPs;
- Level of Assurance;
- step-up-authentication;

- group management;
- authorisation management;
- attribute self-management;
- logs of changes of attributes.

A.10 The Long Tail of Science

As shown in the use case analyses for large research collaborations, several VOs have already begun defining their AAI requirements and some have deployed services. However, this is not the norm throughout Europe. To be able to deploy their own infrastructure, VOs must have a high degree of internal organisation and technical capabilities, a good level of trust, and in many cases representation as a legal entity. These characteristics apply only to a small percentage of collaborations in Europe. At the most remote end of the long tail are collaborations formed for a single paper for a conference or journal. In addition, there are a wide range of H2020 and other projects, which are smaller and more decentralised than those in the typical EFSRI groupings, and which may not have a centralised legal entity or resource for infrastructure.

References

[CERN]	http://home.web.cern.ch/about
[CLARIN]	http://clarin.eu/content/mission
[COmanage]	http://www.internet2.edu/products-services/trust-identity-middleware/comanage/
[CRISP]	https://indico.cern.ch/event/177418/session/1/contribution/5/material/1/1.pdf
[DARIAH]	https://dariah.eu/about/mission.html
[DARIAH Req]	https://wiki.edugain.org/File:DARIAH-UseCase-Description_1.0.pdf
[eIDAS]	https://ec.europa.eu/dgs/connect/en/content/electronic-identification-and-trust-services-eidas-regulatory-environment-and-beyond
[ELIXIR]	https://rems.elixir-finland.org/
[ELIXIR Excelerate]	https://www.elixir-europe.org/news/elixir-accelerates-major-horizon-2020-funding
[ELIXIR FIM]	http://cds.cern.ch/record/1442597/files/CERN-OPEN-2012-006.pdf?subformat=pdfa&version=1
[ELIXIR Pilot]	http://www.elixir-europe.org/webinar/update-outcomes-elixir-aai-pilot-launched-2012
[ERIC]	https://ec.europa.eu/research/infrastructures/index_en.cfm?pg=eric
[FIM4R Doc]	https://cdsweb.cern.ch/record/1442597
[GÉANT]	http://www.geant.org/
[Guest IdP]	https://refeds.org/wp-content/uploads/2015/06/refeds-survey-2015.pptx
[LHC]	http://en.wikipedia.org/wiki/Large_Hadron_Collider
[LoA]	http://levelofassurance.org/registry.html
[Moonshot]	https://wiki.moonshot.ja.net/display/HOME/Home
[**PaNdata**]	http://pan-data.eu/
[SAML ECP]	http://docs.oasis-open.org/security/saml/Post2.0/saml-ecp/v2.0/cs01/saml-ecp-v2.0-cs01.pdf
[Shibboleth]	https://wiki.shibboleth.net/confluence/display/SHIB2/IdP+ECP+Extension
[STORK]	https://www.eid-stork.eu
[Umbrella]	https://umbrellaid.org/euu/umbrellaforusers
[VCH]	http://www.virtualcampushub.eu/About-VCH
[VHC Req]	http://www.virtualcampushub.eu/~/media/Sites/virtualcampushub/Deliverables/d5,-d-,3_virtual%20campus%20hub%20technology.ashx?la=da
[We-NMR]	http://www.wenmr.eu/wenmr/about-we-nmr-project

Glossary

AA	Attribute Authority
AAI	Authentication and Authorisation Infrastructure
eduGAIN	educational Global Authentication Infrastructure
eduroam	A global service that provides secure roaming connectivity.
CLARIN	Common Language Resources and Technology Infrastructure
CoCo	GÉANT Data Protection Code of Conduct
DAC	Data Access Committee
DARIAH	Digital Research Infrastructure for the Arts and Humanities
EBI	European Bioinformatics Institute
EC	European Commission
EGA	European Genome-phenome Archive
ESA	European Space Agency
EO	Earth Observation
ESFRI	European Strategy Forum on Research Infrastructures
FaaS	Federation as a Service
FIM	Federated Identity Management
FIM4R	Federated Identity Management for Research – a forum for AAI providers, e-Infrastructures and users
FIMig	Federated Identity Management Interest Group of the RDA
FP7	Seventh Framework Programme for Research and Technological Development
GÉANT	The collaboration of European NRENs, delivering e-infrastructure and services to research and education.
GN3plus	The administrative name for the previous phase of the GÉANT project, which ran from 1 April 2013 to 30 April 2015.
GN4-1	GN4, Phase 1, also known as SGA1, is a project run by the GÉANT project from 1 April 2015 to 1 April 2016 and part-funded from the EC's Horizon 2020 programme under Grant Agreement 653998.
IdM	Identity Management
IdP	Identity Provider
JRA	Joint Research Activity
LoA	Level of Assurance
MDS	Metadata Distribution Service
NA	Networking Activity
NREN	National Research and Education Network
PKI	Public Key Infrastructure
REFEDS	Research and Education Identity Federations

RDA	Research Data Alliance
REMS	Resource Entitlement Management System
SA	Service Activity
SA5	Service Activity 5 “Application Services”
SAML	Security Assertion Markup Language
SCI	Security for Collaborating Infrastructures
SSH	Social Science and Humanities.
SP	Service Provider
SSO	Single Sign-On
VC	Video Conference
VO	Virtual Organization
VOOT	Virtual Organization Orthogonal Technology