05-05-2015

Deliverable D14.1 (DJ3.0.1) Report on the Achievements and Recommendations for any Future Work

(Identity and Trust Technologies for GÉANT Services)

Deliverable D14.1 (DJ3.0.1)

Contractual Date:	27-02-2015
Actual Date:	05-05-2015
Grant Agreement No.:	605243
Activity:	JRA3
Task Item:	Tasks 1 & 2
Nature of Deliverable:	R (Report)
Dissemination Level:	PU (Public)
Lead Partner:	GÉANT Association
Document Code:	GN3PLUS14-1007-25
Authors:	L. Florio (GÉANT Association), R. Poortinga - van Wijnen (SURFnet), M. Kremers (SURFnet), A. Biancini (GARR), R. Hedberg (NORDUnet/University of Umea), A. Solberg (NORDUnet/UNINETT), D. Poehn (DFN/BADW-LRZ), D. Janny (PSNC)

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7 2007–2013) under Grant Agreement No. 605243 (GN3plus).

Abstract

This deliverable reports on the achievements of the Research Activity on Identity and Trust technologies during 2013-2015. The research topics and goals of this activity include distributed attribute management for authorisation and groups, support tools for Identity Federations and new identity technologies.

Table of Contents

Table of Figures

Executive Summary

This deliverable reports on the achievements of the Research Activity on Identity and Trust technologies during 2013-2015. The research topics and goals of this activity include distributed attribute management for authorisation and groups, support tools for Identity Federations and new identity technologies.

The main results of this JRA are two-fold:

1. Part of the effort has been directed towards engaging the relevant groups to progress the specifications of OpenId Connect (3.4), UMA (3.4) and Certificate Transparency (3.1). This work also involved improving the development of software that implements specifications as per IETF procedures. SUNET has been involved in all of the efforts mentioned above and has developed different tools that implement the specifications of OpenId Connect, UMA and Certificate Transparency.

2. Another part of the work has been directed towards researching solutions to manage groups and user affiliation in an innovative and effective way. Examples of this work are, InAcademia (2.4), Federated Authorisation (2.2) and the renewed specification for the VOOT protocol (2.1).

Section 1 introduces the Identity and Trust research activity and its goals, and reports on the outreach activities and standardisation work that formed an important part of the work of the team. Besides standard outreach within the project, the task has engaged with two Open Call projects (HEXAA and GÉANT Trust Broker) and with communities such as ORCID, OpenID Connect, Kantara Initiative and Internet2.

The achievements of the two tasks in this activity are covered in Sections 2 and 3 of this document. They describe the work carried out by the Attributes and Groups (2) and Identity and Trust (3) tasks per work item.

Section 4 outlines the overall conclusions drawn from the research activity.

1 Introduction to the Identity and Trust Research Activity

Collaboration across institutional borders is a cornerstone of research, and is one of the strengths of the R&E community. Federated access has been a major milestone in terms of facilitating the sharing of resources between the community's institutions in a secure and user-friendly manner. However, the uptake of federated access and requests from different international research collaborations for its adoption, have highlighted new use cases that are not easily addressed by the current model. Some of the most commonly encountered of these are:

- Access to services that require information on whether a user is member of a (specific) group; this information is normally unknown to the users' home institutions where the authentication takes place.
- Access to services by researchers not affiliated with any institution the service accepts identities from[1].

The Identity and Trust Research Activity aims to develop and harmonise technologies used by Identity Federations to facilitate collaboration, support virtual organisations, support sharing of resources and explore new identity and trust protocols.

The key goals of this activity are:

- To enhance existing technologies used to implement federated access to support different use cases;
- To explore future technologies that could eventually replace SAML (Security Assertion Markup Language) [SAMLOverview], the current standard used by R&E federations;
- To offer support for virtual organisations;
- To develop tools for the operators of Identity Federations and service providers to ensure that best practices are followed.

[1] This is often referred to as 'guest use' as the researcher in question is not affiliated to a known federated institution. Since in most cases these researchers are full members of the collaboration, this term is used mainly to refer to their status with respect to the institutions in question.

To achieve these goals, the work is split up into two tasks: *Attributes and Groups* and *Identity and Trust technologies*, with the former focusing on the federated access (and authorisation) support for federated access and virtual organisations and the latter focusing on future identity technologies and tools for identity federations and service providers. The tasks are described individually in chapters 2 and 3.

What the tasks have in common in order to reach their respective goals is the need to do outreach (both project internal and external), to do protocol specification and standardisation and (further) development and validation of software. These common concerns are summarised in the following sections.

1.1 Internal Outreach

1.1.1 Service Activity 5

The JRA3 research activity has worked most closely with the "Users access and applications" service activity (SA5).

Two of its work items have reached sufficient maturity to be migrated to SA5: InAcademia and Federated Authorisation.

InAcademia is a framework to validate users' affiliation (users are either students, staff or faculty members, or alumni). This information is necessary for those services that offer favourable conditions for the academic community. The legal aspects of InAcademia were investigated in collaboration with the eduGAIN team of SA5. More information about InAcademia is provided in Section 2.

Federated authorisation proposes a model to manage groups of users as well as their attributes via an attribute provider or attribute authority operated at federation level. More information on this item work is also provided in Section 2.

1.1.2 Open Calls

JRA3 has worked very closely with two projects funded via the GN3plus Open Call: HEXAA and GÉANT-TrustBroker, and the same topics from the HEXAA project will continue to be addressed in future projects, while GÉANT-TrustBroker is planned to become a specific task within that same activity with the goal of finalising its development and preparing for pilots as part of the SA5 activity in a future stage.

1.1.2.1 *HEXAA*

The Project
The HEXAA, Higher Education EXternal Attribute Providers, GN3plus Open Call project has the goal of developing a new Attribute Authority (AA) implementation to facilitate AA integration into federations and eduGAIN. This work is carried out by MTA Sztaki and NIIF INSTITUTE, both based in Hungary.

Virtual Organisations (VOs) have the problem that much of the user data they require to authenticate and authorise users to their services is VO specific. This project-specific information is not well suited to be managed by their institutional IdP. A side effect of this is that the user data is not always provided to the relaying parties in a consistent way.

There are already different VO management solutions to manage groups and to provide additional group information. HEXAA is a VO management tool that operates as a VO and Profile management interface, which is capable of registering any attribute, handling the necessary consent to release the attribute. HEXAA is implemented in a standardised way via the SAML2 Attribute Authority binding.

More information on HEXAA can be found at [HEXAA] and [HEXAA_EDUID].

The collaboration

The JRA3 activity worked together with HEXAA and other Identity and Attribute related services from within the NREN community, such as PERUN (Cesnet), SURFconext (SURFnet) and Unity IDM (University of Warsaw), in order to gain a better understanding of the world of attribute providers, and the related use cases, issues and future challenges, such as those involved in standardisation. Interoperability testing between the different attribute providers' tools, researching the problem space and an assessment of the tools were carried out. The tool assessment was performed by the JRA3 task in collaboration with the other parties mentioned above. More information on this topic is reported in Section 2 Achievements of the Attributes and Groups Task.

1.1.2.2 *GÉANT-TrustBroker*

The Project

The OpenCall project GÉANT-TrustBroker (GNTB) facilitates the user-triggered, on-demand exchange of IdP and SP metadata as a basis for SAML-based authentication and authorisation. GNTB allows users to initiate the first-time contact between SPs and IdPs to perform the required preparations for identity data exchange in a fully automated manner. Therefore, it automates the setup of IdP-SP communication through extensions to the SAML implementation Shibboleth. These extensions include the metadata of the IDPs and SPs and - for IDPs - any needed attribute conversion rules.

Implementation

A prototype was implemented to demonstrate the usage of GNTB. The implementation of GNTB involves different components:

- Central GNTB service extending the Centralized Discovery Service of Shibboleth.
- Extension of the Embedded Discovery Service of Shibboleth for use with the central GNTB service
- Extension of the Shibboleth software for IDPs and SPs

The combined result automates previously manual configuration steps. The extensions are an important part of the GÉANT-TrustBroker Open Call project, and are described in various white papers and presentations found on the project web page [TrustBroker].

Standardisation Work

The core workflow was defined based on the specification for the initial metadata exchange and further functionalities, such as the attribute conversion rule repository. The GNTB core workflow, called Dynamic Automated Metadata Exchange (DAME), has been submitted as an Internet-Draft (I-D) to the IETF. The I-D concentrates on the initiation of the metadata exchange in order to lead the user and to reduce waiting time as well as workload for the IDP/SP administrators. The initial draft 00 of DAME was submitted to the IETF on 18 June 2014. Feedback from REFEDS, JRA3 and SA5 led to an improved 02 version of the draft, including a detailed example and a comparison with related technologies such as Metadata Query Protocol, REEP and IdP proxying.

Outlook

Plans are to continue this work in future projects, specifically related to the following items:

- Improvement of the GNTB demonstrator for the pilot phase, including improvements of the IDP/SP software extension and the central GNTB service, in order to have a fully working service for the pilot phase. Closer collaboration with SA5, as well as the acquisition of pilot users is considered.
- Enhancement of GNTB to include results of REFEDS, e.g., REEP and other GN-related approaches, such as HEXAA.
- Further standardisation work on the I-D of DAME.

1.2 External Outreach

Outreach has been an important aspect of the work of the Identity and Trust activity. The outreach activities carried out were mostly in two areas: dissemination of the research work results and engagement with relevant groups.

As part of its general dissemination work, the Task team provided regular updates to the TERENA Task Force on European Middleware Coordination and Collaboration [TF-EMC2] as well as at mayor events such as NRENs conferences, the TERENA Networking Conference and Internet2 events.

As regards engagement with other groups, the team has been working closely with SA5, but has also reached out to external communities. One of the communities the research activity interacted with is ORCID, a not-for-profit initiative aiming to provide a unique persistent identifier for researchers [ORCID]. The team discussed the possibility of accessing the ORCID site using federated credentials. There was clearly a lot of interest in enabling federated access on the ORCID site and it was agreed that ORCID would join the SURFnet Identity Federation and would be exposed to eduGAIN via that route. At the time of writing, ORCID are in the process of reviewing the legal documents required for them to join the SURFnet federation.

SURFnet also discussed ORCID use cases for InAcademia. Further work on this will be carried out as part of the InAcademia pilots that are expected to start in the spring of 2015.

The team has carried out a number of tests to query the ORCID API using the OAuth protocol to retrieve the ORCID identifier associated with a user. It is considered that this would be an efficient way to use ORCID as an external third-party attribute provider.

A simple module was implemented in SimpleSAMLphp to aggregate the ORCID identifier with the attributes returned by IdP. Figure 2.1 shows the federated login window to access the ORCID login page.

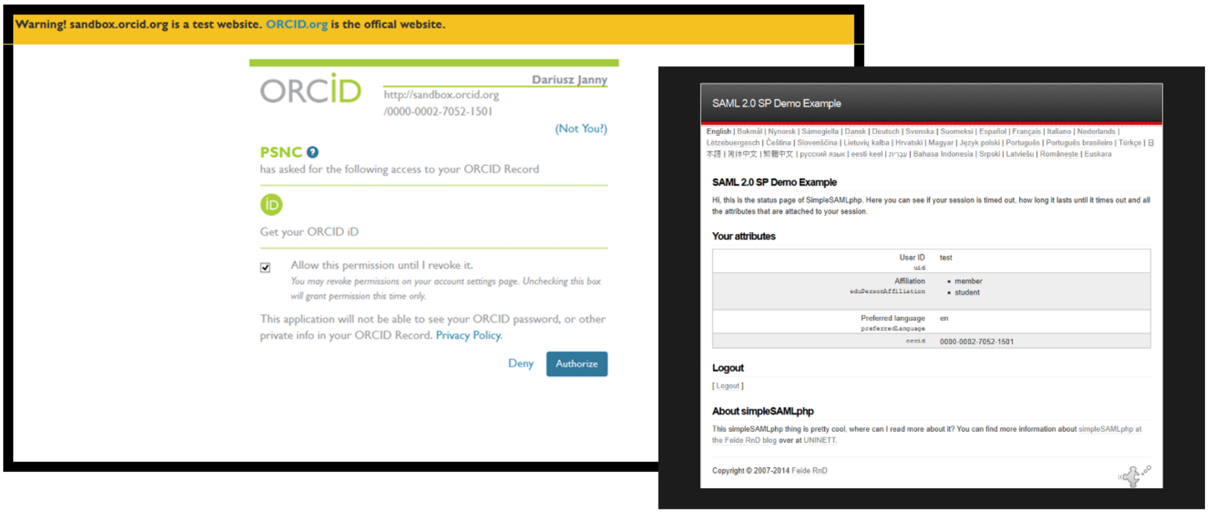

Warning! sandbox.orcid.org is a test website. ORCID.org is the offical website.

Dariusz Janny
http://sandbox.orcid.org
/0000-0002-7052-1501

PSNC
has asked for the following access to your ORCID Record

Get your ORCID iD

☑ Allow this permission until I revoke it.
You may revoke permissions on your account settings page. Unchecking this box will grant permission this time only.

This application will not be able to see your ORCID password, or other private info in your ORCID Record. Privacy Policy.

SAML 2.0 SP Demo Example

English | Bokmål | Nynorsk | Sámegiella | Dansk | Deutsch | Svenska | Suomeksi | Español | Français | Italiano | Nederlands | Lëtzebuergesch | Čeština | Slovenščina | Lietuvių kalba | Hrvatski | Magyar | Język polski | Português | Português brasileiro | Türkçe | 日本語 | 简体中文 | 繁體中文 | русский язык | eesti keel | עִבְרִית | Bahasa Indonesia | Srpski | Latviešu | Românește | Euskara

SAML 2.0 SP Demo Example

Hi, this is the status page of SimpleSAMLphp. Here you can see if your session is timed out, how long it lasts until it times out and all the attributes that are attached to your session.

Your attributes

User ID uid	test
Affiliation eduPersonAffiliation	• member • student
Preferred language preferredLanguage	en
orcid	0000-0002-7052-1501

Logout

[Logout]

About simpleSAMLphp

This simpleSAMLphp thing is pretty cool, where can I read more about it? You can find more information about simpleSAMLphp at the Feide RnD blog over at UNINETT.

Copyright © 2007-2014 Feide RnD

Figure 1.1: ORCID login and attribute aggregation

The code for this integration with ORCID is accessible in the GIT repository [SSP]. Through its work on OpenID Connect, the team engaged with the OpenID Connect Foundation, the entity responsible for its development. More information on OpenID Connect is provided in Sections 3.4 and 3.5.

1.3 Standardisation Work and Protocols

The Task’s activities in three main areas in particular have contributed to its standardisation work. These are:

- The work on **Certificate Transparency** specifications – Certificate Transparency is an open framework promoted by Google “to detect SSL certificates that have been mistakenly issued by a certificate authority or maliciously acquired from an otherwise unimpeachable certificate authority”. The SUNET team have been involved in the specs work that takes place in the IETF working group Public Notary Transparency [TRANS_WG];
- The contributions to the **User Managed Access** (UMA) specifications – The purpose of the UMA Work Group [UMA] is to develop an OAuth-based protocol that enables an individual to control the level of access services have to his/her data. The Identity and Trust Technologies task leader has been involved in the working group in charge of the specifications. This work is driven mostly within the Kantara Initiative.
- Finally the work on the **OpenID Connect** specifications – OpenId Connect is a new protocol built on OAuth2 to request and receive information about authenticated sessions and end-users. The Identity and Trust Technologies task leader has played a key role in defining and implementing the specs.

Alongside these standardisation activities, futher work was also carried out on developing and finalising the VOOT protocol and data model. Following a major update, SCIM was deemed to be a more suitable protocol base layer for VOOT than OpenSocial, on which the first iterations of VOOT were loosely based. More detailed information on VOOT can be found in Section 2.1.

1.4 Software development

The team spent considerable effort on developing new software or extensions for existing software. The main products of this work are listed by work item below. More detailed information on each item can be found in the relevant items in Sections 2 and 3.

- **Federated Authorisation**
  A VOOT connector for Grouper and a VOOT enrolment plugin for Moodle were developed, as well as a Shibboleth authenticator for MediaWiki.

- **InAcademia**
  Apart from investigating the legal aspects, almost all effort spent on this work item is dedicated to development of the InAcademia service.

- **Certificate Transparency**
  To validate the specifications from the IETF working group Public Notary Transparency, a new Certificate Transparency implementation was developed.

- **Federation Lab**
  New test tools were developed for expansion of the Federation Lab, allowing automated verification of identity protocols such as SAML2 and OpenID Connect and of associated metadata.

- **UMA and OpenID Connect**
  Implementations for both UMA and OpenID Connect were developed, which were among the first few implementations available for both protocols. The same work item also developed the official OpenID Foundation conformance test tool and started doing the same for UMA. In addition an extension for non-web authentication was created for OpenID Connect.

2 Achievements of the Attributes and Groups Task

The work in this task has progressed mostly according to the plan. The task has delivered even an additional work item (InAcademia). The main results are described below.

2.1 VOOT – Data Model and Protocol

Work on the VOOT protocol started during the GN3 project.

VOOT is a protocol and data model to dynamically fetch information about the group and roles of the currently authenticated user accessing a service or application.

In retrieving pre-defined group information about the users, VOOT differs from alternative protocols in a number of ways:

- VOOT operates on the current authenticated user, with delegated authorisation using OAuth 2.0 (three-legged).
- VOOT is designed to work in a dynamic fashion where information is fetched when needed, rather than being pre-provisioned or synchronised in a batch job.
- Because of the targeted dynamic mode of operation, typical use cases are supported using a minimal set of HTTP requests, allowing a fluid user experience without unnecessary delays.

VOOT is designed as a modern HTTP-based protocol on top of OAuth 2.0. VOOT is demonstrated to work very well with OpenID Connect, one of the emerging standards web applications use for authentication.

Since work on the VOOT data model began during the GN3 project, several underlying supporting transfer protocols have been considered over time, including OpenSocial and SCIM. Owing to a major update (2.0) of SCIM, combined with the attention and traction of the standard, SCIM has been chosen as the base layer protocol design. A deeper analysis of the updated SCIM protocol demonstrated that SCIM fails to properly support the three items mentioned above. In addition SCIM operates a very limited data model for groups that does not allow the addition of attributes to the membership relation between a user and a group. For example, a user who is a member of a 'Course' cannot have an associated student or teacher role assigned to that relation. However, even if SCIM does not perfectly solve the use cases, it is nonetheless a very well-designed protocol, with an advanced and flexible extension model.

The current VOOT version 2.0 intends to re-use as much as possible features from the SCIM data model and protocol, mainly adding better support for the focus points above. The VOOT data model defines four entities: users, memberships, groups and group types. Definition of specific group types is not part of the core specification itself, but is instead left for communities to standardise.

User
Membership
Group
Group type

Figure 2.1: Entities defined in the VOOT 2.0 data model

Although VOOT may operate well in various architectures, it works particularly well with API platforms such as SURFconext [SURFconext] and Feide Connect [Feide_Connect].

SURFconext is a production service in the Netherlands that uses an early version of VOOT for its group API.

Feide Connect is a significant API Platform that is currently being built in Norway. Feide Connect builds its group APIs entirely on the latest version of VOOT. Feide Connect aggregates groups in parallel from several sources, making the resulting information 'consumable' by services using the VOOT protocol in combination with OpenID Connect.

VOOT has been implemented several times as part of the work with the Feide Connect prototype. This has ensured that Feide Connect always used the latest version of VOOT enabling rapid validation of changes, but also means that the production-ready VOOT implementation in Feide Connect is so recent that its source code has not yet been made available, though it will be made available soon [Feide_Code]

A series of open source collaboration tools is soon expected to arrive with plugin support for VOOT in combination with OAuth 2.0 and OpenID Connect – driven by platforms such as SURFconext and Feide Connect.

Although VOOT provides a solution to a real need by collaborations and their (collaboration) tools, the future of VOOT depends on whether it will be supported in existing and future API Platform, in combination with collaboration tools with plugins provided. It is hoped that a more stable, structured approach to standardisation will derive from the growing use of VOOT, and also that specific group types will be defined outside the specification itself. Hopefully some collaboration effort between the NRENs may result in standardised data models for specific group types for higher education and research.

The latest revision of VOOT version 2.0 has also been used in a proof of concept to showcase federated authorisation, outlined in the next section.

The full specification is available at a separate web site [OPENVOOT].

2.2 Federated Authorisation

One of main goals of this research activity was to study and test solutions to implement an external system able to manage user authorisation in a cross-organisational context. The work in this area has focused on using Grouper [Grouper] as the tool to manage groups to provide federated authorisation.

The main criteria for choosing Grouper were:

- The possibility of creating groups with users coming from different institutions;
- The possibility of delegating the management of different groups to different users;
- Grouper's widespread deployment.

A proof of concept was implemented to demonstrate the use of Grouper to manage federated authorisation. The proof of concept tests the integration of Grouper with three main applications considered to be representative: MediaWiki, Moodle, and GARRbox, a custom application developed by GARR for their community.

2.2.1 Grouper for Authorisation Management across Organisational Boundaries

Grouper is an enterprise access management system designed to manage the highly distributed environments and heterogeneous information common to universities. In principle, different tools could be used to obtain similar results, assuming that the chosen group management tool has similar key functionalities.

Grouper has been used to implement an external system to manage user authorisation based on groups. Within each group, each user has specific access rights. Such groups can include users from different organisations (authenticated by different IdPs). This approach enables management of virtual organisations, that is, a group of users coming from different institutions but sharing the same access rights with regard to a specific set of applications.

Group administrators are responsible for managing groups by adding or removing users. Specific group administrators can be assigned to each group. These can be different for every (sub) group, allowing for very flexible group management arrangements.

Grouper offers a web interface that administrators can use to see group status and manage enrolments and memberships.

The Grouper management interface will operate in a Federation (or Inter-federation) as an SP. So group managers will be given access to this application and will be able to access it to manage groups through federated authentication, both the managers and the members of the group.

2.2.2 Applications integrated with Grouper

Three different applications have been integrated with Grouper to test the principles of external authorisation:

MediaWiki – structured Wiki, typically used to run a project development space, a document management system, a knowledge base, or any other groupware tool.

Moodle – a learning platform designed to provide educators, administrators and learners with a single, robust, secure and integrated system to create personalised learning environments.

GARRbox – a custom solution, developed by GARR, to provide personal cloud storage for the Italian biomedicine community.

These three applications have been chosen because they are among the most representative applications used in the R&E community. These use cases, moreover, provided the opportunity to study external authorisation by adding complexity step by step.

The first application, MediaWiki, has relatively simple authorisation needs. To effectively manage authorisation, MediaWiki, in the same way as any other wiki, only has to retrieve some additional information about users at the moment of login. The information retrieved is related to the membership of the user to specific groups. To retrieve this attribute, the approach followed was to instruct the SP (or eventually an SP proxy) to contact the right Attribute Authority, as defined in the SAML standard, to retrieve group membership information.

The scope of the technical activities of the PoC has been to integrate MediaWiki access management with groups and users defined within Grouper. In this way MediaWiki delegates the authorisation process, leveraging an external system, at a federation level. To prove the use case worked effectively MediaWiki was accessed using federated identities, and it was verified that the proper access grants are released to the users depending on group definitions inside Grouper.

Moodle, as a Learning Management System, has more complex requirements in terms of user authorisation. Moodle internally permits to identify users with different roles (students, teachers, etc.) for specific courses. So roles are not assigned statically to users, but a single user can have different roles for different courses (he can be student of one course and teacher of another, for instance). This information, however, can continue to be retrieved at the time of login from the Attribute Authority, as described for MediaWiki.

Moodle also needs to retrieve a course list and an enrolment list. To perform this operation, it can be configured to retrieve lists of groups and users from an external source. To implement this functionality, VOOT has been chosen as the protocol to describe the communication mechanism between Moodle and Grouper and to retrieve external group information. The choice of VOOT was determined by considerations that it is a standard protocol and it does not create any “lock-in” to Grouper, allowing the use of a different tool to represent group information.

The scope of the technical activities of the PoC has been to integrate Moodle courses with groups defined inside Grouper. The general idea was to describe in Grouper a structure of groups that could reflect the courses in Moodle. At this point the users could be assigned to a group (with different entitlements) to represent the role that each user will have for the specified course. To prove the use case worked effectively, Moodle was accessed using federated identities and it was verified that the proper access grants are released to the user depending on group definitions inside Grouper. Moreover, Moodle must be integrated with Grouper to obtain the list of courses (defined as groups in Grouper), the list of teachers and the list of students for every course.

Our last use case, the application GARRbox, helped identify the key aspect that should be taken into consideration when designing externalised authorisation for a custom application. GARRbox is intended to be a service similar to Dropbox. It allows a user to have disk space available over the Internet for storing data and files as a personal backup. It was decided to try and integrate GARRbox into the Grouper authorisation management process for two main reasons:

1. Since GARRbox is a custom application, it can help understand how emerging applications can be designed and modelled to be fully compliant with the delegated authorisation process introduced in this task;

2. GARRbox raises some interesting issues as regards user authorisation going beyond the cases described so far. GARRbox, in addition to groups, needs additional authorisation and other attributes for its users (such as the size in GB of their virtual disk inside the application).

GARRbox represented an interesting use case especially because of this need for additional information, which it was possible to implement into Grouper as membership attributes. This data is then passed to GARRbox – via VOOT – in user-specific attributes.

To prove the use case worked effectively, the groups and the authorisation attributes needed by the GARRbox application were externalised to Grouper. GARRbox was then integrated with Grouper to obtain the list of groups, and of the users participating to every group, as well as to obtain a set of other attributes for every user accessing the application.

2.2.3 Tools Delivered

Small pieces of code and components were developed during the research activities In order to integrate the applications described above. All code developed as part of the activity has been shared with the community, specifically:

- **VOOT connector for Grouper:** A specific connector was developed in order to implement a VOOT interface for Grouper. This connector works on the same classes and logic implemented for the Grouper WebServices but fully implements version 0.9 of the VOOT protocol (the latest available at the moment of these activities).
  The VOOT connector for Grouper has been shared with Internet2 and integrated in the Grouper official repository. The code was entered in the Grouper 2.2.1 release in November 2014. The documentation [VOOT_CONNECTOR] and code [GROUPER_VOOT] for the connector are available online.

- **MediaWiki Shibboleth authenticator:** In order to integrate MediaWiki into Grouper and enable this application to manage the isMemberOf attribute correctly by describing user membership to groups, a modification of the Shibboleth Authentication module has been developed. This modification has been released to the extension manager and is now freely available to the community. The relevant documentation page on MediaWiki has been updated accordingly [Shibboleth] and the code has been shared with the extension manager on his GitHub account [GitHub].

- **Moodle enrolment plugin for VOOT**: In order to retrieve course and enrolment information from Grouper, a specific enrolment plugin has been developed for Moodle. This plugin is available on GitHub and has been submitted as a JIRA task to the Moodle staff. The JIRA task, submitted inn November 2014 [CONTRIB-5413], a Moodle forum page [Moodle] and the code developed [Moodle-Enrol_VOOT] are available online.

2.2.4 Lessons Learnt and Possible Evolutions

This research activity enabled the Task to test how an external authorisation system could be used to federate user authorisation and to gain greater understanding of how an effective delegation can be implemented in Grouper, as well as of which technical aspects need to be taken into consideration to achieve this.

The PoC showed how different applications need to be modified to appropriately use attributes from an Attribute Authority. Security aspects were also looked at.

The main recommendation for any future work is to decouple the two main functions of user authorisation:

1. Only authorised users should be granted access to the application. This verification has a binary result (the user can either be granted access or denied access to the application). Due to the ease of this verification and its high relevance and importance, this operation could be executed directly by the SP before the application itself is contacted. In this way, no bug or other problem in the application code can affect authorisation and result in improper access.

2. Authorised users must be assigned specific permissions depending on which groups they belong to. This fine-grained grant assignment must be executed by the application after it receives the “isMemberOf” attribute from the Attribute Authority.

One aspect proved especially relevant, i.e. leveraging standard protocols to achieve the integration of applications. In the Proof-of-Concept, the Attribute Authority mechanism and the VOOT protocol were used. Both these interfaces are standard and can be supported by different authorisation systems. This means not being bound to any specific technological choice for its constituent components (Grouper for instance) and that different tools could be switched to in later phases of a production service.

Some potential issues and negative aspects of the proposed approach and solution were also identified as a result of these activities. In particular the Grouper interface was found not to be very user friendly (although the latest version, 2.2.1, did show strong improvements). This less-than-optimal user experience is a strong drawback for the solution since administrative users are confronted with this every time they use Grouper to perform group management tasks (which can be divided over quite a number of relatively non-technical people using the Grouper approach). For this reason, in order to use Grouper in a production environment, it will be very important to provide suitable training and documentation to help these users learn to operate the interface. Depending on the environment in which this solution is to be used, it may even be advisable to design and create specific new (simplified) user interfaces suited for the purpose.

Finally the PoC helped identify the main elements that are missing in order to effectively leverage the Attribute Authorities (AA) approach in real SAML federations. In particular the PoC suggested two main areas of improvement:

1. **Discovery**: currently it is not clear how SPs can discover and use AAs within a federation. At the moment an SP must be manually configured to retrieve additional attributes from one specific AA, but a discovery mechanism to identify this external authority should be designed. This discovery process must be different to the one used for authentication (the Discovery Service (DS) used to find IdPs for a user who wants to login to a service). In fact, the discovery for IdPs is driven directly by the user (who is able to choose the right IdP from a list). In the AA discovery process, user-driven discovery may not be a good option.
2. **Privacy**: currently the AA releases user information to be added as new attributes to the user session without verifying that the user requesting the information is entitled to receive them. Moreover, the user is not informed of the attributes that the AA is releasing for him/her. For authentication, different solutions have been proposed to this general problem, one being for uApprove to show the user the attributes that will be released beforehand and allowing him/her to deny the release of some of these attributes. This mechanism is not present for AA attribute release, but a similar solution could be studied and introduced.

2.3 Assessment of Existing Attributes and Groups Tools

Researchers often work together in cross-organisational and international collaborations. The services these collaborations use require specific user information (or attributes), which for the most part are not (and cannot be expected to be) provided by institutional identity management systems.

In recent years, multiple tools became available within the R&E space, such as attribute, group and workflow services, in order to solve (parts) of the problems mentioned. It is however a challenge for a VO to choose the right tool. In order to provide a more clear view of this topic, the research activity, in collaboration with HEXAA, PERUN, SURFconext and Unity IDM, has drawn up an overview of well-known existing tools in the R&E space [AA]. The information for this overview was gathered from the tools' developers themselves, as well as from feedback on the tools' usability provided by the REFEDS community. The overview provides information on the capabilities of the different products for specific topics, subdivided into a number of categories (General, Identity management, VO/Group management, standards compatibility, and privacy). A VO can use this comparison to select the right tool(s) with their own requirements and use cases in mind. The overview is a living document, allowing new products to be added or incorporating changes to the products listed [AA].

2.4 InAcademia

2.4.1 Introduction

Many commercial service providers offer discounts for students and/or staff members in academia. For these services it is critical to reliably validate whether a user is indeed a student or staff member of an institution, as this is the basis for the discount provided. Other attributes may also be useful, but could be provided by the person directly. As the discounts for students and staff are often considerable, these services are highly valuable for the users.

The information flow for validation is very different from the flow for identification as is currently common within Identity Federations. In the identification process, the service possesses no prior knowledge of a user, and the Identity provider provides all profile (attribute) information. In a validation workflow, on the other hand, the service already has a profile of the user, obtained through different means, and simply needs to be able to reliably verify some of the claims made by the person (e.g. is someone really a student? So that he/she is entitled to that special discount). In this example, it is only the fact whether someone is a student that needs to be verified, while where or what they study is irrelevant – the service simply needs the assurance that the claim of being a student has been verified by an authoritative source, without the need to reveal additional attributes.

Identity Federations in Research and Education currently support the delivery of attributes using the well-known SAML authentication interfaces of their federations. Although Service Providers only requiring validation can join a federation and use the attributes provided by an Identity provider to carry out their own assessment, there are downsides to this approach.

One of these is that a federation Service Provider needs to abide by a rather stringent federation policy put in place among other reasons to protect user privacy. However, if the Service Provider were not to receive such privacy-sensitive information in the first place, it is possible that a much lighter policy could be applied.

Another downside is the SAML protocol itself, which is considered by many to be rather cumbersome to implement and maintain. This especially where the only functionality needed is a simple assurance that a claim provided by someone is correct, a functionality that SAML does not even provide since it was designed for the purposes of authentication and attributes provision, not validation. In recent years, new protocols such as OpenID Connect have emerged which are equally suitable for secure attribute transportation, but which have a significantly lower technical impact.

Finally, many federations already support services with validation needs on a national level. This means, for example, that Microsoft services have 27 interfaces towards the national federations, which is extremely time-consuming to maintain for the service.

All these issues can for the most part be addressed by creating a single transnational validation service. This single service could be a service provider in the national federations (via eduGAIN) with an accompanying policy to be defined, which could provide a validation service for those services that require validation only. This means an attribute release profile with anonymous attributes, possibly combined with user consent, allowing a much lighter policy to be applied.

The services in need of validation could then use a much simpler technical connection to the validation service, while centralisation and high standardisation would significantly lower the cost of connecting and supporting such connections at the same time.

Using a transnational eduGAIN rather than federation-based service decreases the burden on national federations, while at the same time providing an incentive to Identity Providers to join eduGAIN.

2.4.2 Use cases

In order to design the InAcademia service itself, it is helpful to take into consideration a number of use cases that it will have to support. A number of use cases (or issues) were in fact what prompted the initial work on developing the InAcademia service, with a few more added after some consideration. These use cases are presented below.

- Microsoft offers various services for students, including Dreamspark and Microsoft IT Academy. These services include access to (almost) free software and training materials and certification. In the current situation, even if Microsoft joins eduGAIN, it would still need to deal with potentially a few thousand IdPs including metadata handling, etc., just to get ‘student status (and because of that to enjoy a free or discounted service). A simple and preferably single point of contact would significantly lower the effort needed to supply these kinds of services. (Note: it is assumed that if more privacy-invasive attributes are needed, a ‘full’ federation package will be required).

- Five NRENs are jointly procuring (pan-national) mobile services for students and staff within their national borders. The providers that will be delivering this service will need an economical cheap and easy system for student and staff validation as part of the onboarding process. InAcademia provides exactly the right type of service for this.

- In certain areas, for example in local campus towns, SMEs could offer digital services without the need to actually become fully-fledged Federation members. The upfront investment for joining the Federation, both technically as well as in terms of contractual obligations, is considered to require too much investment for many small businesses. In The Netherlands, a use case exists where local theatres and sports accommodations would like to leverage student status as part of the reservation flow for tickets and accommodation booking.

- In some EU countries, email addresses in use by academic institutions can be recognized by their extension, such as “.ac.uk” in the UK. Some Service Providers leverage this as a means to determine student status. However, this approach has inconsistent results as, even leaving aside the fact that university staff will also have the same email address extension, more and more institutions are abandoning the provision of student email accounts, instead allowing students to self-register private email addresses. In (most) other countries institutional email addresses cannot be filtered by extension, and in addition students are also allowed the use of self-registered email addresses.

- Many Institutions prefer to release as little personal information as possible to external parties. By using a gateway which ensures the provision of globally unique identifiers to connected SPs, institutions only need to release attributes to a trusted party operated by their federation.

- eduGAIN, the pan-European interfederation effort, currently suffers from a 'Catch-22' type problem, in that there are not enough SPs to make it worthwhile for IdPs to connect and vice-versa. A common strategy for mitigating this problem is to offer a service that is attractive to IdPs in order to encourage them to join eduGAIN. Given the potentially very high value it could have for their users, as well as the fact that joining eduGAIN is rather low effort for the IdP, the InAcademia service could be just such a service. This would then pave the way for any future services, whether commercial or collaborative in nature.

2.4.3 Development, status, and timeline

2.4.3.1 *Design*

The InAcademia simple validation service was developed based on the use cases, issues and requirements presented in the previous sections. The service provides a simple and secure way to validate users' affiliation with an academic institution. It is designed to scale on a pan-European level and allow connections, by acting itself as proxy, between multiple services and an Identity provider through eduGAIN. At the same time, the InAcademia service meets the requirements for privacy and data protection as defined by the eduGAIN Code of Conduct. The picture below shows an overview of the setup of InAcademia.

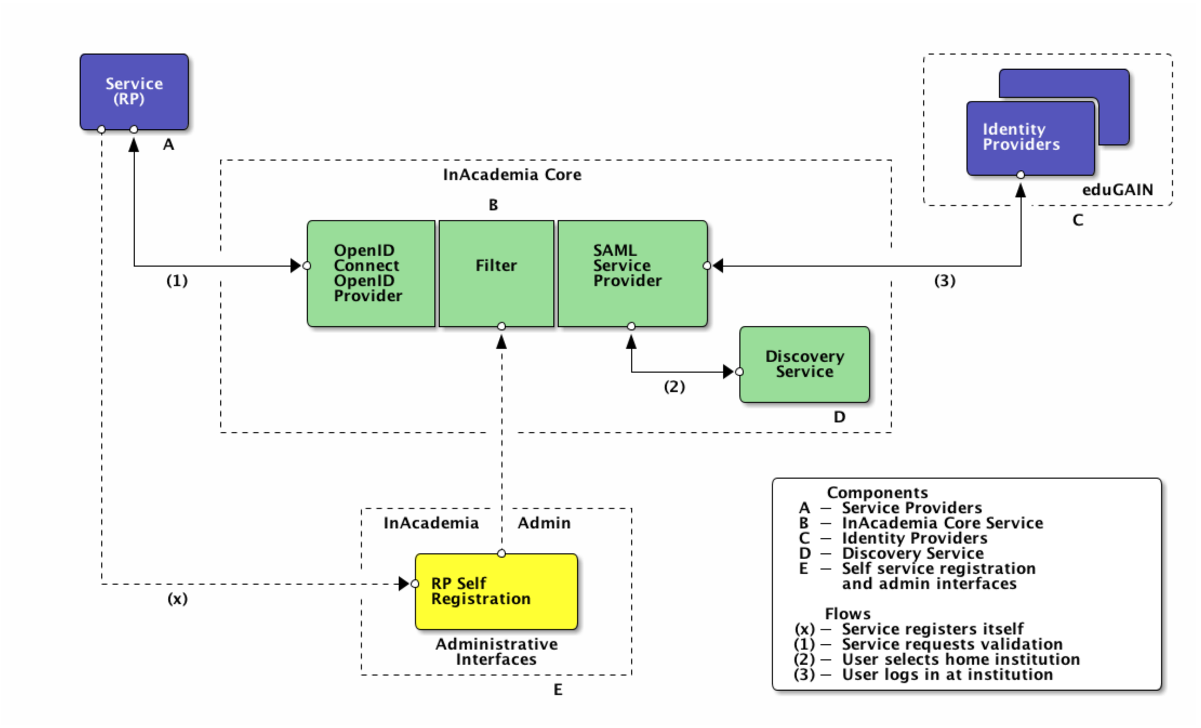
Service (RP)
A
InAcademia Core
B
OpenID Connect OpenID Provider
Filter
SAML Service Provider
Identity Providers
eduGAIN
C
(1)
(3)
Discovery Service
(2)
D
InAcademia
Admin
RP Self Registration
(x)
Administrative Interfaces
E
Components
A – Service Providers
B – InAcademia Core Service
C – Identity Providers
D – Discovery Service
E – Self service registration and admin interfaces
Flows
(x) – Service registers itself
(1) – Service requests validation
(2) – User selects home institution
(3) – User logs in at institution

Figure 2.2: InAcademia overview.

The InAcademia Core consists of a central block, acting as a proxy between IdPs from eduGAIN and SPs. Communication between InAcademia and the IdPs uses SAML2; meaning that from the point of view of the IdPs, the InAcademia service acts as a SAML2 Service Provider (SP).

Communication between the services requesting validation and InAcademia is based on OpenID Connect, meaning that from the point of view of the services the InAcademia service acts as an OpenID Provider, while the service requesting validation acts as an OpenID Connect Relying Party (RP).

A separate block is responsible for storing and providing configuration data to the filter central to the InAcademia core service. It also provides the administrative and configuration interfaces, such as a self-registration functionality for services.

The central 'Filter' block of the core service provides the business logic of the InAcademia service, with the following functionalities, among others:

- OIDC Identifier handling.
- OIDC Claims creation based on SAML2 attribute statements.
- Handling of the SAML RelayState encryption.
- Error handling and logging.
- Various multi-language end-user GUIs, including consent and error screens.

The operational setup will be geographically distributed, which will allow requests to be serviced by the 'closest' node and at the same time can also be used for failover purposes.

2.4.3.2 *Development and status*

Work on the InAcademia service started in the second half of 2014, with use case descriptions and requirements gathering, followed by high-level design. Coding of the rudimentary service and flows started soon after. A first run-through of the basic flow between a SAML IdP, core service and an RP based on this rudimentary service was carried out at the end of September 2014.

In Q4 2014, eduGAIN support was added, testing started with real IdPs and the error handling and logging was improved. A public website was launched [InAcademia] explaining the InAcademia service to SPs, IdPs, federations and users. It also provides the first technical information on how to connect to the InAcademia service. Additionally, discussions have begun on the legal aspects of InAcademia with regard to the eduGAIN Code of Conduct and the form and shape of the contractual agreements with RPs that would be needed for InAcademia.

At the time of writing (February 2015), nodes are being prepared at various locations across Europe to host the first beta service deployment of InAcademia. The first nodes are located in Iceland, Denmark, the Netherlands and Greece.

2.4.3.3 *Timeline*

So far, work on InAcademia has progressed ahead of schedule, and will continue in future projects in order to further rollout the service and make it operational, including the development of a business case and payment model and the onboarding of SPs. Version 1.0 of the service will be available at the end of GN3plus, along with a report on the legal considerations for InAcademia written in collaboration with the eduGAIN team. The final contribution in GN3plus will be an evaluation report including recommendations for continuation of the work in future GÉANT projects.

3 Achievements of the ‘Identity and Trust Technologies’ Task

Work carried out by this task progressed for the most part as planned, with the exception of issues related to account linking. Initially, the development of a centralised service to link accounts had been envisaged. However, no real use-cases ever materialised and the team discovered that many developers preferred to build this functionality into the service they developed. Therefore the account linking activity was discontinued.

3.1 Certificate Transparency

Certificate Transparency is an open framework promoted by Google “to detect SSL certificates that have been mistakenly issued by a certificate authority or maliciously acquired from an otherwise unimpeachable certificate authority”. The goal is to identify certificates that have been either mistakenly issued or issued by a compromised CA, by using public logs of certificates. With CT all TLS/SSL certificates issued are monitored at real-time in a public logbook. The RFC6962 [RFC6962] describes this problem as well as a possible solution. This document is however experimental and needs to be tested and confirmed by two different implementations before becoming a real Internet standard.

The Task’s engagement in this area has been two-fold:

- SUNET was involved in the specifications work taking place in the IETF working group Public Notary Transparency [TRANS_WG];
- Efforts were devoted by SUNET to implement the specifications and by SURFnet to test Google’s implementation of these specifications. The resulting tools are expected to be ready for preproduction testing at the end of the GN3plus project.

3.2 Expanding Federation Lab with new test tools

The initial Federation Lab (FedLab) suite of tools, developed under the GN3 project, has been further expanded to encompass new tools. FedLab is a set of tools for testing, verification and debugging of various Identity protocols that are used by Identity Federations.

FedLab uses sophisticated techniques for performing automated verification of identity protocols, such as SAML and, recently, OpenID Connect. There are now a number of tools that federations/IdPs/SPs should be encouraged to use, namely:

- SAML2test – To verify that an implementation conforms to the standard and the profile.
- Metadata Analysis – To verify the correctness of metadata. This can also be used to verify if the attributes provided by the IdP comply with the EU Data Protection Directive.
- Verify_entcat – to verify that an IdP acts according to the entity category it specifies.
- IdP monitor – to verify that the whole authentication process works, as predicted, for a user.

There are plans to run these tools under REFEDS during 2015, to gain experience and collect feedback from the community. Once this has been completed, additional development may be continued in future projects, when tools could be made increasingly operational. This work has been entirely led by the University of Umeå.

3.3 Identities Gateways

In an environment that is divided between a small number of different identity protocols there is a need to bridge between these, so that users may access internet resources that are protected using OpenID Connect using their University credentials (which are SAML2 based) and vice versa. Within JRA3 T2 work is underway on continuously improving the current any-to-any proxy. The software for this is available on GitHub [IdProxy].

3.4 New Protocols: UMA and OpenId Connect

User-Managed Access (UMA) is a profile of OAuth 2.0. UMA defines how resource owners can control access to protected resource by clients operated by arbitrary requesting parties, where the resources reside on any number of servers, and where a centralised authorisation server governs access based on resource-owner policies. Version 1.0 of the specification is presently in a public review period and is expected to be published at the end of March 2015.

The JRA3 T2 team has developed one of the few available implementations of the specs and are currently in the process of setting up the first instance of the UMA conformance test tool.

OpenID Connect is an identity protocol that builds on OAuth2.0. It allows clients to verify the identity of the end user based on the authentication performed by an Authorisation Server, as well as to obtain basic profile information about the end user in an interoperable and REST-like manner.

The team has not only contributed to the specifications but also written a complete OpenID Connect implementation and more notably the conformance test tool that will be the official OpenID Foundation conformations test tool. This test tool was used to certify the first OpenID Connect implementation at the start of April, with others to follow. The official launch will take place at the 2015 RSA conference in San Francisco (April 20-24).

3.5 Non-Web authentication using OpenID connect

Non-web federated authentication has long been an issue that has hindered the wider deployment of federated access. Although web access is becoming increasingly widespread, there are still a number of applications used in the R&E community that do not support web-based access.

There are different ways to implement federated access in the case of non-web based applications. One of these is based on the SAML Enhanced Client Profile (SAML-ECP) that is designed for clients other than browsers (i.e. desktop applications). Future projects may explore whether the wide deployment of this profile should be pursued.

Another option is the Moonshot project (or ABFAB protocol), which uses RADIUS protocol to transport SAML assertions. Although Moonshot has made significant progress, wider deployment is still rather challenging, as it assumes an extensive knowledge of both SAML and RADIUS and requires changes at the client machine of the user.

An alternative approach totally unrelated to SAML is to use the OpenID Connect extension for non-web, which is detailed in an internet draft [NONWEB-01]. The draft specifies a way in which long-lived authentication tokens that can be revoked can be requested from the OpenID provider using a specific scope parameter. These long-lived authentication tokens act as client-specific passwords to be used in combination with clients and protocols (such as IMAP) that do not readily support browser-based federated authentication. The long-lived nature of these authentication tokens reduces the frequency with which new tokens have to be requested, thereby improving the user experience. The use of client-specific tokens (different tokens for different clients and applications) that can be revoked means that tokens can be simply revoked for a specific client/service combination if needed, without affecting other clients and services of the same user, thereby simultaneously improving both user experience and security.

The team created a Proof-of-Concept [NONWEB-POC] implementation of this Internet-Draft.

4 Conclusions

Work in the activity has progressed mostly according to the yearly plan. There have been a number of useful results some of which (InAcademia and Federation Authorisation) are mature enough to be deployed for pilots in 2015. The further development and piloting of InAcademia are planned for the next project. Results and experiences of the Federated Authorisation work are the basis of the future GÉANT project intranet, for which the same approach, as well as some of the same tools, will be adopted.

Other results, which are also very relevant, have been achieved in the standardisation area (such as UMA and Certificate Transparency).

The progress made has helped identify new challenges. In particular, OpenID Connect, used to date mostly by the industry, is challenging the SAML2 standard protocol currently widely in use by Identity Federations. The development of InAcademia, which bridges between SAML2 and OpenID Connect, highlights "how to interoperate" rather than posing a threat to Identity federations. This could provide federations with an opportunity to extend their reach and usefulness. More work is needed to identify further ways in which OpenID Connect can be introduced into existing identity federations.

Federated access to non-web resources still remains a challenge. OpenID Connect has proven to be a good option, however its lack of deployment in the R&E environment may make it less appealing, further emphasizing the need to identify ways of introducing OpenID Connect into existing federations as well as to look into ways of advancing the Moonshot pilot and investigating the possibility of a wide-scale ECP deployment within the eduGAIN service area.

There is consensus that Identity Providers should and will in the future concentrate on identity vetting only, whereas attributes used for authorisation should be provided by other sources and be aggregated. Preliminary work on this area was a study into federated authorisation and a comparison of different groups and attributes management tools, both carried out in T1. The work on federated authorisation has shown this approach to be viable, taking into account the main recommendations of separating the two functions of user authorisation (basic access and setting specific user permissions) and leveraging standard protocols as much as possible to achieve integration without being bound to specific technological choices. The work also showed the need for further improvements to usability of the tools involved (such as Grouper) for the less technically proficient users, and for more research into discovery and privacy/consent related to attribute authorities. More work on this specific topic is planned for future activities, if possible in collaboration with the research and user communities that would benefit from such a solution.

References

[AA]	http://bit.ly/aa-overview
[CONTRIB-5413]	https://tracker.moodle.org/browse/CONTRIB-5413
[Feide_Code]	https://github.com/feideconnect
[Feide_Connect]	http://feideconnect.no
[GitHub]	https://github.com/kir-dev/mediawiki-shibboleth-authentication
[GROUPER_VOOT]	https://github.com/Internet2/grouper/tree/master/grouper-misc/grouper-voot
[HEXAA]	http://www.geant.net/opencall/Authentication/Pages/HEXAA.aspx
[HEXAA_EDUID]	https://hexaa.eduid.hu/
[IdProxy]	https://github.com/its-dirg/IdProxy
[InAcademia]	https://inacademia.org
[Moodle]	https://moodle.org/mod/forum/discuss.php?d=275042
[Moodle-Enrol_VOOT]	https://github.com/biancini/moodle-enrol_voot
[NONWEB-01]	https://tools.ietf.org/html/draft-sakimura-oidc-extension-nonweb-01
[NONWEB-POC]	https://github.com/its-dirg/non-web-oidc
[OPENVOOT]	http://openvoot.org
[ORCID]	http://orcid.org/
[RFC6962]	https://tools.ietf.org/html/rfc6962
[SAMLOverview]	http://www.oasis-open.org/committees/security
[Shibboleth]	https://www.mediawiki.org/wiki/Extension:Shibboleth_Authentication
[SSP]	https://git.man.poznan.pl/stash/projects/SSP/repos/simplesamlphp/browse
[SURFconext]	https://www.surf.nl/en/services-and-products/surfconext/index.html
[TF-EMC2]	https://www.terena.org/activities/tf-emc2/
[TRANS_WG]	https://datatracker.ietf.org/wg/trans/charter/
[TrustBroker]	http://www.geant.net/opencall/Authentication/Pages/TrustBroker.aspx
[UMA]	http://kantarainitiative.org/confluence/display/uma/Charter
[VOOT_CONNECTOR]	https://spaces.internet2.edu/display/Grouper/Grouper+Voot+Connector

Glossary

AA Attribute Authority. A server acting as an Attribute Authority role as defined in SAML 2.0 specification, cf. [SAMLOverview]

AAI Authentication and Authorisation Infrastructure

Authentication Authentication is the act of confirming the truth of an attribute of a single piece of data or entity (the user of an application, for instance).

Authorisation Authorisation is the function of specifying access rights to resources related to information security and computer security in general and to access control in particular.
More formally, “to authorize” is to define an access policy.

DS Discovery Service. Federation component which allows the user to select the Identity Provider to be used for authentication from a list.

eduGAIN The eduGAIN service allows Authentication and Authorisation Infrastructures to interact, enabling the sharing of data between federations and providing an interconnection framework to applications willing to provide their services, content or resources to multiple federations.

FaaS Federation-as-a-Service.

Federation Identity federation. An association of organisations that come together to exchange information as appropriate about their users and resources to enable collaborations and transactions.

IdM Identity Management systems.

IdP Identity Provider. A server acting in an Identity Provider role as defined in SAML 2.0 specifications, cf. [SAMLOverview].

NREN National Research and Education Network.

REFEDS Research and Education FEDerations The mission of REFEDs is to be the voice that articulates the mutual needs of research and education identity federations worldwide, see https://refeds.org

PoC Proof of Concept, realization of a certain method or idea to demonstrate its feasibility or a demonstration in principle, whose purpose is to verify that some concept or theory has the potential of being used. A proof of concept is usually small and may or may not be complete.

SAML Security Assertion Markup Language

SP Service Provider. Service Provider. A server acting in a Service Provider role as defined in SAML 2.0 specifications, cf. [SAMLOverview].

WebSSO Web Single Sign-On.