26-05-2011

Deliverable DS3.3.2: eduGAIN Use-case Analysis

Deliverable DS3.3.2 v1.0

Contractual Date:	31-03-2011
Actual Date:	26-05-2011
Grant Agreement No.:	238875
Activity:	SA3
Task Item:	T3
Nature of Deliverable:	R
Dissemination Level:	PU
Lead Partner:	NORDUNET/UOGOTH
Document Code:	GN3-11-134
Authors:	Brook Schofield (TERENA), Valter Nordh (NORDUNET), Josh.Howlett (JANET), Shannon Milsom (DANTE)

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7 2007-2013) under Grant Agreement No. 238875 (GÉANT).

Abstract

This deliverable provides an overview of the work by SA3 T3 (eduGAIN) to analyse a range of use cases that support the development of the eduGAIN Policy Framework and Metadata Distribution System.

Table of Contents

Table of Figures

Table of Tables

Executive Summary

SA3 (Multi-Domain User Applications) Task 3 (eduGAIN) analysed a range of use-cases, from GN3 services and projects from related communities, to highlight the distinct needs in terms of interfederated access of these important projects and services. The goal was to create a generalised policy and a technological infrastructure to satisfy the wide range of participants envisaged in an interfederation environment. This deliverable provides an overview of the use-cases selected, how their needs contributed to the development of the eduGAIN Policy Framework and Metadata Distribution System (MDS) and concludes with proposals for future use-cases for targeting communities and services.

1 Use-case Analysis

1.1 Objectives

The eduGAIN interfederation service is intended to enable the trustworthy exchange of information between the GN3 Partners' federations. The initial goal is to enable pan-European Web Single Sign On (Web SSO) to both GÉANT services and to those provided by other communities represented by, or associated with, the GN3 Partners.

In collaboration with GN3 Service providers (SP) and users, the SA3 T3 team identified, defined and evaluated a range of use-cases to ensure that the intended outcomes of this service would meet with the expectations of a range of key projects and services. To ensure that the choice of use-cases was balanced, they were drawn from three main audiences outlined in the eduGAIN business case [BusinessCase]:

- GN3 services that need to obtain information about end users (typically for authorisation, authentication and personalisation of services).
- Federations operated by the Partners whose member organisations may provide end user information.
- Non-GN3 services associated with these federations, in particular pan-European research projects, other international research projects and commercial content providers that, while not directly associated with the Project, are nonetheless important to the GN3 community.

The requirements derived from the analysis of the selected use-cases will focus on the following areas:

- Policy and governance of an interfederation service.
- Technical requirements for service metadata distribution.

These requirements were defined within the eduGAIN project to provide supporting scenarios for the development of the eduGAIN Policy Framework and technical requirements for the Metadata Distribution System (MDS). See section 2.2.

As the uptake of the eduGAIN interfederation service is envisaged to be a gradual process, the initial group of use-cases will be supported and extended in future rounds, building on existing requirements and offering refinements as a showcase for further research and education collaboration opportunities and increasing the support for a variety of services and communities. Continued development of new use cases is an important step in the development of eduGAIN, as more services utilise the eduGAIN infrastructure.

Future rounds of use-case analysis should:

- Assign resources within the project to support specific integration efforts.
- Investigate case studies for interfederation success via eduGAIN.
- Encourage additional federations to participate in the eduGAIN service.

In addition to supporting interfederation infrastructure, the use-case analysis also aims to enable collaboration within the GN3 project and, specifically, the eduGAIN task.

Finally, the results from the analysis of the use-cases are aimed at encouraging the adoption of eduGAIN by the broader research and education community and highlight the need for the deployment of campus and country federated Authentication and Authorisation Infrastructure (AAI).

1.2 Audience

The use-cases were proposed and analysed for the following audiences:

- **eduGAIN Team (SA3 T3)**: The use-cases are especially of interest to the MDS development and Policy subtasks.
- **GÉANT Stakeholders**: The use-cases raise awareness in the GÉANT community of the scope and goals of the eduGAIN interfederation service and the importance of federation and interfederation technologies.

1.3 Selected Use-cases

A range of use-cases were proposed by the eduGAIN (SA3 T3) team with a limited number of services selected to be the focus in order to concentrate on the policy and governance of the eduGAIN interfederation service.

While requirements were contributed to the policy and metadata service subtasks, and while first results were realised, federations were able to expose their service and identity metadata to the pre-pilot and pilot stages of the eduGAIN interfederation service. This further supplied input to the analysis of these development subtasks (MDS and Policy).

The most recent round of use-cases provided input to the development subtasks with an emphasis on building federation participation, as the usefulness of the eduGAIN service depends on the number of federations participating in the service.

1.3.1 Round 1

The focus of the first round of the use-case analysis was to derive the requirements for:

- Policy and governance of an interfederation service.
- Technical requirements for service metadata distribution.

Table 1.1 lists Round 1 Use-case services.

Use-case	Description
CLARIN	CLARIN is an EU-project easing the sharing of language resources in European universities and research institutes with a large projected user-base across 32 countries.
Foodle	An online poll for meeting time agreement.
GN3 Project wiki	Original collaboration environment for the GN3 project.
eduroam TTS and wiki	The trouble ticket system (TTS) and wiki of the eduroam (GN3 SA3 T2) activity is a website to coordinate trouble shooting and knowledge about the eduroam service within GN3 participants.

Table 1.1: Round1 Use-cases

1.3.1.1 *CLARIN*

CLARIN (Common Language Resources and Technology Infrastructure) was selected as an example large-scale pan-European research project, which is representative of European Strategy Forum on Research Infrastructures (ESFRI), including collaborations such as Lifewatch and DARIAH. CLARIN is expected to have a user base of 50,000 from 171 institutions in 32 countries. A significant portion of the CLARIN infrastructure is maintained by the Max Planck Institute for Psycholinguistics in The Netherlands, which is working with SURFfederatie to federate its infrastructure. While CLARIN is a significantly complex project, beyond the capacity of the eduGAIN interfederation service during its pre-pilot and pilot stages, the development of the framework necessary to support such large scale projects is important for eduGAIN adoption by services with an extensive participant base, while aiming to minimise the user administration and authorisation burden of these pan-European projects.

The analysis of this use-case yielded the following requirements:

- The need to support the interfederation of a large number of Identity Providers (IdP) and Service Providers (SP) throughout Europe and the transfer of authorisation details of participants between many jurisdictions.
- For each user a globally unique identifier from an IdP and an email address.
- The capability for user authorisations to be maintained in a distributed manner by one or more Access database(s).

1.3.1.2 *Foodle*

Foodle is a production service of UNINETT that offers multiple identity federations, bilateral peering arrangements and social networking authentication mechanisms for its users. The demand for a federated polling and date/time selection tool has resulted in Foodle being bilaterally federated with 23 countries, five of which are via Kalmar Union (Iceland, Denmark, Sweden, Finland and Norway). The administrative burden of supporting connections to additional federations, especially for a free service, is evident to the administrators of Foodle. eduGAIN is a vital component in simplifying this administrative process and allowing Foodle to scale to an increasing user-base without scaling the administrative overhead.

Foodle is a web-based tool for several people to agree upon date and time for a meeting. It is also used for polls, voting and registration for events. With a worldwide target audience, Foodle benefits from interfederating via eduGAIN as fewer contracts are required. Allowing a large interfederated audience to Foodle also benefits the Multi-Domain User Applications Research (JRA3) Identity Federations (T2) activity with issues of scalability of discovery services, which ultimately benefits future eduGAIN users.

The analysis of this use-case yielded the following requirements:

- For each user a globally unique identifier from an IdP.
- A degraded service will be provided to users without the following attributes. These must be provided using a consistent name format, using the eduPerson schema where possible and the Object Identifier (OID) syntax:
  - The full name of the user.
  - The email address of the user.
  - The preferred language of the user, represented as an ISO 639-1 language code.
  - The name of the user’s institution/organisation.

1.3.1.3 *GN3 Project wiki*

The GN3 Project wiki was the main collaboration tool in the initial year of the GN3 project, providing all project members access to general management, administration and documentation with secure collaboration workspaces for project teams. There were also some users external to GN3, such as people collaborating within the Liaison and Support (NA4) community. Due to the transition to Microsoft SharePoint from TWiki as the platform for this environment, this instance of the GN3 Project wiki was not connected to the eduGAIN pre-pilot service. While work to connect to SharePoint is currently on hold, it will be pursued as a future service.

The analysis of this use-case yielded the following requirement:

- For each user a globally unique identifier from an IdP.

1.3.1.4 *eduroam TTS and wiki*

The eduroam® trouble ticketing system (TTS) and wiki are part of the GN3 European eduroam service supporting infrastructure. The aim is to foster collaboration between the Operational Team (OT) and Roaming Operators (RO). This service was originally connected to eduGAIN, as it existed within the GN2 project and it was important for the continuity of service that this use-case be maintained as part of the GN3 project. The GN3 eduroam confederation now supports all GN3 Partners and an increasing number of Associate Partners. It is important to intertwine the supporting infrastructure of eduroam with eduGAIN to encourage adoption of Identity Interfederation as an extension to the successful confederation of eduroam deployments. eduroam deployment is also growing outside the GN3 project, and access to supporting infrastructure is useful in advocating eduGAIN to a wider audience.

The analysis of this use-case yielded the following requirements:

- For each user a globally unique identifier from an IdP.

To access the eduroam wiki, a user must have attribute urn:mace:dir:attribute-def:eduPersonEntitlement value set to: urn:geant:edugain:entitlement:eduroam:wiki.

- A degraded service will be provided to users without the following attributes. These must be provided using a consistent name format, using the eduPerson schema where possible and the OID syntax:
  - The email address of the user.
  - The full name of the user.

1.3.2 Round 2

The focus of the second round of the use-case analysis was to:

- Assign resources within the project to support specific integration efforts.
- Investigate case studies for interfederation success via eduGAIN.
- Encourage additional federations to participate in eduGAIN service.

The aim of the project was to create a service (policy and technical infrastructure) that most federations could take part in, with few or any modifications to their existing policy and technical framework, encouraging them by making it easier to join the service. .The initial requirements from the first round of the use-case analyses were also carried forward with the goal of fine-tuning the policy, governance or technical requirements for instances that would exclude federations or services from participating.

Table 1.2 lists Round 2 Use-case services.

Use-case	Description
GÉANT IdP (GIdP)	IdP for GN3 Participants without access to their own identity federation.
GN3 Project SharePoint	Collaboration environment for GN3 project.
TERENA SPProxy	Service aggregator for TERENA Secretariat and interface to TNC2011 website.
perfSONAR	Performance monitoring and measurement toolkit.

Table 1.2: Round 2 Use-cases

1.3.2.1 *GIdP*

The eduGAIN interfederation service also operates the GÉANT Identity Provider for the Homeless (GIdP), tasked with issuing identities to users associated with GN3 Partners who do not yet operate an identity federation, or their federation has not joined the eduGAIN interfederation pilot service.

The GN3 project saw an increase of six new Partners and Associate Partners, each without an established identity federation. The GIdP service is required for these new organisations to demonstrate federated identity and provide access to eduGAIN enabled services.

These GIdP *guest* accounts were used in the GN2 project. With the transition of services to the GN3 eduGAIN service, it was important for the GIdP to be migrated and updated to support the MDS. Use of the GIdP is purposefully low and limited to GN3 Partners and Associate Partners as an alternate authentication and authorisation mechanism to access services and support testing alongside an existing or developing identity federation.

1.3.2.2 *GN3 Project SharePoint*

The GÉANT Intranet site is a focal point and portal for all the activities of the GN3 project. Currently all participants have access to this site via *guest accounts* within the environment that replaced the previous GN3 Project wiki. Having this crucial part of the GN3 projects supporting infrastructure integrated into the eduGAIN interfederation service highlights the importance and utility of eduGAIN to all GN3 partners.

Owing to resource limitations, it was not possible to perform a complete analysis of this use-case's requirements. It is anticipated that this will be attempted again once the necessary resources become available.

1.3.2.3 *TERENA SPProxy*

The TERENA Service Provider Proxy (SPProxy) aggregates a range of identity management endpoints including identity federations to the services offered by the TERENA Secretariat. For a service that targets the entire European research and education community, it is important to support the authentication mechanisms that are available within the community.

The analysis of this use-case yielded the following requirements:

- For each user a globally unique identifier from an IdP that can be expressed as one of the following, in order of preference:
    - SAML2 persistent name identifier.
    - eduPersonTargetedId.
    - eduPersonPrincipalName.
    - mail.
- If none of these identifiers are transmitted, the user will be refused access.
- The following attributes, if returned, will also be used to populate a locally managed database of the user’s information, as an alternative to a manual update:
    - The full name of the user, from the *givenName*/*sn* or *displayName* attributes.
    - The email address of the user, from *mail*.
    - The name of the institution/organisation, from *o*.

Integration with the eduGAIN interfederation service has vastly simplified the maintenance of previously bilateral peering arrangements that proved not to be scalable, as they were manually maintained by the TERENA Secretariat and the peered federation. This has increased the benefit of joining the Dutch identity federation (SURFfederatie, particularly for the TERENA Secretariat and SURFnet connected organisations that wish to participate in eduGAIN in the future. The TERENA Networking Conference (TNC2011) website is the most visible of the TERENA services connected via the SPProxy. This service does not offer guest accounts for any user. Instead, it offers the option to use a federated account (via eduGAIN or legacy bilateral peering arrangements), social networking account or OpenID. With almost 200 submissions for presentations and over 500 attendees, this service will expose a wide range of NRENs and campuses to a federated environment and the possibility of using eduGAIN.

1.3.2.4 *perfSONAR*

perfSONAR (PERFormance Service Oriented Network monitoring Architecture) enables performance monitoring measurement data exchange between networks, making it easier to solve problems between multiple hosts in an interconnected multi-domain environment. Weathermaps, looking-glasses, IP performance metric (IPPM) measurements and many other monitoring applications have already been implemented as decentralised services using the perfSONAR framework. Providing a common authentication and authorisation infrastructure for this multi-domain environment aligns with the eduGAIN interfederation service.

However, the analysis of this use-case concluded that eduGAIN was not, on the basis of existing technologies, an appropriate solution in the short-term. This is because applying eduGAIN’s Web SSO technology to applications that do not use web protocols is currently problematic. It seems more likely that a solution based on Public Key Infrastructure (PKI) would be more appropriate, at least in the short-term. This work is now continuing in SA3 T1. In the medium to long term, it is possible that the Moonshot work in JRA3 T2 will provide a more comprehensive solution.

2 Development of eduGAIN

The selected use-cases were used as a basis for creating the framework around the policy and technology requirement of the eduGAIN interfederation service. The significant outputs of the eduGAIN task team are the eduGAIN Policy Framework and MDS.

2.1 The eduGAIN Policy Framework

The inclusion of eduGAIN as an interfederation service by an existing national federation requires no changes to its policy. When a federation exposes a service to eduGAIN, certain additional practices (such as technical, attribute or privacy profiles) may be applied. The eduGAIN Policy Framework provides a foundation for establishing trust between the various participant federations and improving technical interoperability guidance between systems, this interoperability is achieved by the definition of the following profiles:

- **Metadata profile** (REQUIRED), which defines the contents of the SAML 2.0 metadata elements exchanged by the eduGAIN service.
- **Web SSO profile** (OPTIONAL), which defines the SAML 2.0 protocol message exchange for Web SSO in eduGAIN. In the beginning, Web SSO will be the primary use of eduGAIN.
- **Attribute profile** (RECOMMENDED), which facilitates interoperability by defining the syntax and semantics for end users' attributes exchanged between entities in eduGAIN. Some of the attributes are recommended, which means that the IdPs should populate them with appropriate values for the end users.
- **Data protection good practice profile** (OPTIONAL), which introduces policies and practices that adapt the European law on data protection to the attribute exchange in the eduGAIN interfederation service.

For a service, which is primarily a Web SSO interfederation, it is surprising and initially confusing, to many that the Web SSO profile is optional. The need for this profile to be optional is to support interfederation as a conduit for a range of services, notably the perfSONAR use-case, which will not use Web SSO for authentication. This provides a flexible layer for interfederation experimentation with a range of services without the need to renegotiate interfederation agreements between eduGAIN participants.

The Attribute Profile [AttrProfile] initially drew from the GN2 eduGAIN services that were expected to transition to the production eduGAIN interfederation service, particularly the eduroam tools and the Foodle service which already had a wide range of bilateral agreements with IdPs and a defined set of attributes that were commonly exchanged. The existing attribute semantics of federations need not be changed as the recommended attribute profile is designed in a way that minimises incompatibility with existing federations.

As CLARIN had yet to build an interfederation environment, their initial attribute requirements were extensive, dealing with specific use-cases relating to user identification and resource allocation. While a challenging area, an increasing number of services desired broad attribute exchange and a range of opaque and personally identifiable information.

To assist IdPs and SPs with the complications of the data protection directive of the European Union, the data protection good practice profile [GoodPractice] was included in the eduGAIN policy development. This profile is optional because interfederation entities may decide to use any other framework or set of agreements to support the exchange of information.

The eduGAIN policy subtask has done its best to design a set of good practices that may help IdPs and SPs to comply with their duties under the EU Data Protection Directive (95/46/EC). See [GoodPractice]. It is expected that federation metadata aggregation and IdP attribute release tools will evolve to further assist in negotiating data exchange between entities.

Further justification of the eduGAIN Policy Framework can be found in the *Introduction to the eduGAIN Policy Framework* [Introduction].

2.2 Metadata Distribution Service

The metadata, Web SSO and data protection good practice profiles defined as part of the eduGAIN Policy Framework created technical implications for the compliance and validation engine of the Metadata Distribution Service (MDS). See the eduGAIN Resources page on the website [Policy].

The use-cases included a range of IdP and SP implementations that would expose their metadata to the MDS for aggregation. The use-cases on their own were insufficient to build robust consumption, validation and aggregation components of the MDS. Effort was directed toward interfederating services (see section 2.3).

The need for further enhancements was recognised during the pre-pilot and pilot stages of the eduGAIN interfederation service where a range of federation metadata aggregation tools exposed real world metadata to the MDS. While no use-cases supporting SAML 1.1 metadata were accepted by the eduGAIN task force, there were federations participating in the pre-pilot stage of the eduGAIN interfederation service exposing these legacy formats. This wide range of federation infrastructure participation highlighted practical issues of entities being published multiple times from distinct sources as well as divergent metadata expiry times.

MDS development benefitted greatly from the inclusion of use-case services. It established boundaries for the eduGAIN service support in its first instance, and helped their host federations in developing a robust and scalable platform for wider eduGAIN deployment in the future.

2.3 Connected Services

As a result of the supported eduGAIN use-cases, effort was directed toward interfederating the following services:

- Foodle (via Feide).
- eduroam TTS and wiki (via AAI@HR).
- GIdP.
- TERENA SPProxy (via SURFfederatie).

These four services will be available at the start of the production stage of the eduGAIN project.

This required the connection of three National Identity Federations to eduGAIN to support these services, with the added benefit of paving the way for organisations connected to those federations to offer their user-base and services to an interfederated environment.

Work is still underway to support interfederated access to the GN3 Project SharePoint via eduGAIN, and to support the diverse range of countries identified as part of the CLARIN use case.

3 Conclusion

The use-case analysis helped to formulate the creation of the eduGAIN Policy Framework and the Metadata Distribution Service, and resulted in the successful pilot of the eduGAIN interfederation service with participating federations from Croatia, Czech Republic, Finland, Germany, Greece, Hungary, Norway, Poland, Spain, Sweden, Switzerland, The Netherlands and Turkey.

In addition to the use-cases discussed here, future use-cases will be developed into case studies to highlight the benefit of the eduGAIN interfederation service, targeting:

- Countries developing a national identity federation.
- Identity federations with a concentration of collaboration services and projects.
- International partners and federations outside of the GN3 project.
- Multi-national collaboration projects supporting large distributed user communities.

This will drive the adoption of eduGAIN during the production stage of the service (Year 3), encouraging bilateral peering arrangements between federations to be deprecated and providing real-world usage figures on the user-base and array of services available.

See www.edugain.org for a current map of participating federations. As federations from participating countries join eduGAIN, their regions will be highlighted in green.

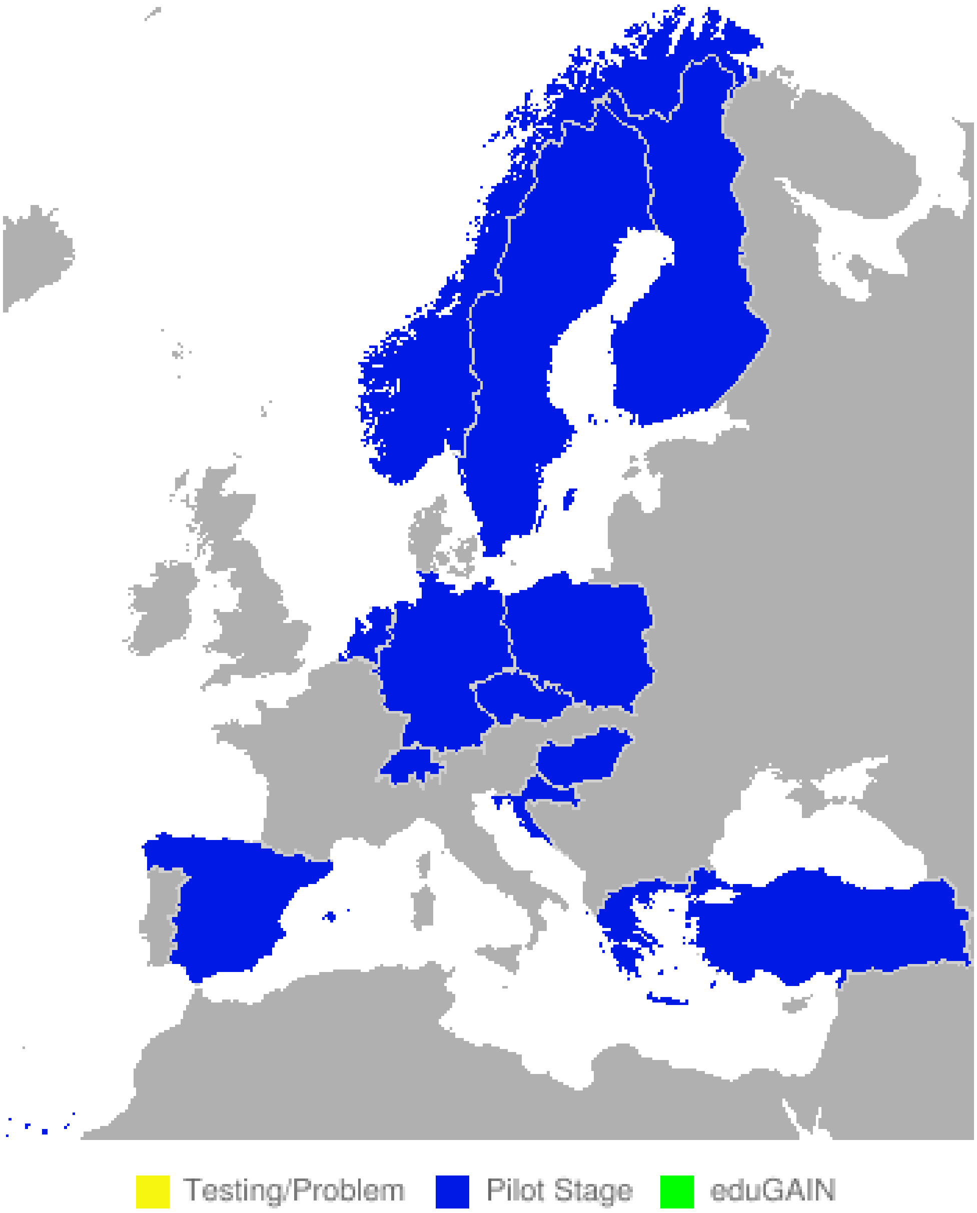

Figure 3.1: Participating federations in eduGAIN pilot

References

[BusinessCase] *eduGAIN Business Case*
J. Howlett, V. Nordh, O. Kreiter, W. Singer
https://intranet.geant.net/sites/Management/Coordination/Services/Documents/GN3-10-111v6%20eduGAIN%20Business%20case%20-%20final.doc

[Policy] eduGAIN policy
http://www.geant.net/service/edugain/resources/Pages/home.aspx

[GoodPractice] *eduGAIN Policy Framework: Data Protection Good Practice Profile*
M. Linden, A Cormack
http://www.geant.net/service/edugain/resources/Documents/eduGAIN%20Data%20protection%20good%20practice%20profile.pdf

[Introduction] *Introduction to the eduGAIN Policy Framework*
M. Linden
http://www.geant.net/service/edugain/resources/Documents/Introduction%20to%20the%20eduGAIN%20policy%20framework.pdf

[WebSSO] *SAML 2.0 WebSSO Profile*
V. Nordh, M. Linden
http://www.geant.net/service/edugain/resources/Documents/eduGAIN%20SAML%202.0%20WebSSO%20Profile.pdf

[AttrProfile] *Attribute Profile*
M. Linden
http://www.geant.net/service/edugain/resources/Documents/eduGAIN%20Attribute%20profile.pdf

[Metadata] *Metadata Profile*
T. Lenggenhager
http://www.geant.net/service/edugain/resources/Documents/eduGAIN%20Data%20protection%20good%20practice%20profile.pdf

Glossary

AAI	Authentication and Authorisation Infrastructure
CLARIN	Common Language Resources and Technology Infrastructure
DARIAH	Digital Research Infrastructure for the Arts and Humanities
eduGAIN	GÉANT Authorisation Infrastructure for Research and Education
eduroam®	Roaming confederation aiming to provide mutual roaming network access to its members
ESFRI	European Strategy Forum on Research Infrastructures
Foodle	Web-based tool to arrange meetings, run polls and register for events
GIdP	GÉANT2 Identity Provider
GN2	GÉANT2
GN3	GÉANT3
IdP	Identity Provider
IPPM	IP Performance Metric
MDS	Metadata Distribution Service
OID	Object Identifier
OT	Operational Team
perfSONAR	PERFormance Service Oriented Network monitoring Architecture
PKI	Public Key Infrastructure
RO	Roaming Operators
SAML	Security Assertion Markup Language
SP	Service Provider
SPProxy	Service Provider Proxy
TTS	Trouble Ticket System
Web SSO	Web Single Sign On