26-04-2016

Deliverable D9.4
Report on Harmonisation Development and Pilots

Deliverable D9.4

Contractual Date:	31-03-2016
Actual Date:	26-04-2016
Grant Agreement No.:	691567
Work Package/Activity:	9/SA5
Task Item:	Task 1
Nature of Deliverable:	R (Report)
Dissemination Level:	PU (Public)
Lead Partner:	GÉANT Association
Document Code:	GN4-1-16-152536
Authors:	Nicole Harris (GÉANT Association), Ann Harding (SWITCH), Daniela Pöhn (LRZ), Mikael Linden (CSC/FUnet), Christos Kannellopoulos (GRnet)

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

Abstract

This deliverable reports on the work of the Trust and Identity Harmonisation Task during GN4-1 to address compatibility and interoperability issues that have arisen in research and education federations, as GÉANT and the NRENs move from national federation provision to global interfederation via eduGAIN.

Table of Contents

Table of Figures

Table of Tables

Executive Summary

The Trust and Identity Harmonisation Task was formed as part of GN4-1 (as task SA5 T1) to address compatibility and interoperability issues that have arisen in research and education federations, as GÉANT and the NRENs move from national federation provision to global interfederation via eduGAIN.

Specifically, the Trust and identity Harmonisation task aims to:

- Address key challenges in federated identity that go beyond the basic scope of eduGAIN but affect the service's ability to offer new features and functionalities.
- Attain the support of eduGAIN members for common strategies and operational practices and policies to enable enhancements to the core eduGAIN service by working closely with REFEDS in the key areas of entity categories, data protection, federation practices, levels of assurance and interoperability.
- Support the future enhancement of eduGAIN by developing underlying federation capabilities, harmonising trust and identity services to better meet user needs, and pilot these capabilities with the eduGAIN community as results mature.
- Ensure that GÉANT trust and identity development takes into account global considerations in service development and the interests of the many non-European countries in eduGAIN.
- Improve and enhance the integrity and trust of federated identity services.

These areas were respectively addressed by five subtasks within the Trust and Identify Harmonisation task. Their work and results are set out in Section 2 of this document, including development, pilots and production deployments, challenges, and recommendations for further work for each area.

Throughout its year of operation, the Trust and Identity Harmonisation Task has:

- Helped increase the usage of the Research and Scholarship Entity Category by 435%, from 21 SPs and 20 IdPs at the start of the GN4-1 project to 87 and 98 respectively at the end of the project.
- Helped increase the usage of the Code of Conduct by 225%, from 63 SPs and 43 IdPs at the start of the GN4-1 project to 83 and 68 respectively at the end of the project.
- Provided a conduit and guidance for the federation community on data protection legislation.
- Worked with AARC to provide a clear baseline assurance profile for Federations.
- Revised the eduGAIN policy to enable future technologies and scalable governance.
- Helped improve the eduGAIN website and reporting approaches.
- Provided clear recommendations on approaches to statistics for federations.

- Tracked the developments that evolve STORK2.0 and eIDAS to develop a strategy on interoperability.

There are clearly many challenges still to be faced by federations and eduGAIN as they grow, but the task has laid strong groundwork that can be built upon in the areas of future services, service development and research activities. The recommendations set out in the main sections of this report have been fed into the development of future GÉANT projects, AARC, REFEDS and other AAI initiatives.

1 Introduction

SAML federations in the Research and Education community have grown very organically, building upon the varying needs of different countries at the grassroots level and heavily influenced by the existing technical capabilities and legislative environments within those countries. These federations did not start out with a vision of interoperating with other federations and mainly catered to the needs of disparate local services. This has led to an environment populated by federations that share the same basic common traits, but have significantly different deployment and service models. As a result, their development path differs greatly from that of eduroam in terms of interoperability, as eduroam operates on a more standardised single-service model for network access.

In the last two years, federations have begun to use eduGAIN in production as a major component of their core services, which has highlighted how their differing practices present problems for interfederation. This has been confirmed by various studies, such as the TERENA AAI Study [AAI Study] and the FIM4R [FIM4R] Report. Federation Operators are starting to note that small differences in metadata registration or publication, for example, can lead to problems when consuming metadata from different federations. The Trust and Identity Harmonisation task seeks to highlight and document these harmonisation issues and to offer solutions to provide a seamless interfederation experience within eduGAIN as well as in other interoperability scenarios.

To successfully tackle this issue, the Trust and Identity Harmonisation task worked in collaboration with the SA5 Enabling Users task (SA5 T5) and the SA5 eduGAIN task (SA5 T3). The harmonisation team also worked closely with the EC-funded AARC project [AARC] and with REFEDS [REFEDS] to ensure buy-in for the recommendations from key stakeholders, such as Federation Operators and e-Infrastructure projects.

2 Harmonisation Developments

As part of its work to address compatibility and interoperability issues in research and education federations, the Trust and Identity Harmonisation task set itself the following specific objectives:

- Address key challenges in federated identity that go beyond the basic scope of eduGAIN but affect the service's ability to offer new features and functionalities.
- Attain the support of eduGAIN members for common strategies and operational practices and policies to enable enhancements to the core eduGAIN service by working closely with REFEDS in the key areas of entity categories, data protection, federation practices, levels of assurance and interoperability.
- Support the future enhancement of eduGAIN by developing underlying federation capabilities, harmonising trust and identity services to better meet user needs, and pilot these capabilities with the eduGAIN community as results mature.
- Ensure that GÉANT trust and identity development takes into account global considerations in service development and the interests of the many non-European countries in eduGAIN.
- Improve and enhance the integrity and trust of federated identity services.

In order to address these issues from a variety of technical and operational angles, the task was divided into five specialist subtasks, each covering several of the objectives with a focus on its specific area of expertise, and coordinated by the task leader and activity leader so as to gain an overview of the work. The work and results of the subtasks are set out in Section 2 of this document, including development, pilots and production deployments, challenges and recommendations for further work for each area.

2.1 Entity Categories

The concept of Entity Categories was developed by the Research and Education FEDerations group [REFEDS] based on an approach first adopted by the InCommon federation [InCommon]. Entity Categories define certain criteria for federation entities i.e. identity providers, service providers and attribute authorities, for the purpose of grouping them into clusters where all entities share the same characteristics. These characteristics could be of any type and the associated tags may have a variety of responses to them. For example, the REFEDS "Hide from Discovery" tag [REFEDS HfD] is applied to Identity Providers (IdPs) and instructs Service Providers (SPs) not to include the IdP in discovery or "WAYF" interfaces. This results in a simpler user experience and is an approach that allows service to be enhanced beyond the initial baseline in a scalable way.

The Trust and Identity Harmonisation Task has focused mainly on those Entity Categories that help IdPs release attributes to SPs in a safe and consistent manner. Within GN4-1, the task worked on the existing Research and Scholarship Entity Category [REFEDS R&S] and two new proposals respectively for an Affiliation Entity Category and an Academia Entity Category.

Overall, the Entity Category subtask helped increase the usage of the Research and Scholarship Entity Category by 435% during the period of GN4-1.

2.1.1 Development

The aim of this subtask was to work with REFEDS to identify and define entity categories to help automate and reduce the workload of attribute release and management for Identity Providers towards Service Providers. This work involved:

- Identifying priority entity categories.
- Working within REFEDS to specify entity category types and content, taking into consideration global needs.
- Piloting key use cases within the GÉANT community in collaboration with Task 5 – Enabling Users.
- Driving take-up of support for entity categories within the GÉANT eduGAIN membership by bringing proposals for their adoption to the eduGAIN SG.

Entity Categories go through a comprehensive consultation process within REFEDS before they can be endorsed as a standard for the community, and the Task was not expected to complete work on any given new category within the 1-year timeframe of the GN4-1 project. Its primary focus was therefore on promotion and adoption of the existing Research and Scholarship Entity Category within the timeframe of GN4-1, and on supporting groundwork for two new definitions with the intention of passing this work on to REFEDS for consultation and standardisation.

The Entity Category subtask also collaborated with REFEDS during GN4-1 to define two new entity categories: the Academia Entity Category and the Affiliation Entity Category. "Academia" is intended as a flag to show Service Providers that any given IdP tagged with this category is from an organisation that is defined as being "academic" against a known set of criteria. This enables the SP to make authorisation decisions which may be driven by terms of use restricting data to academic entities. In this sense, it is different from the "R&S" tag, which indicates that a service is intended for academic users – the first tag identifies an organisation as academic, while the second highlights that a resource is intended for an academic audience. The "Affiliation" category is a simple flag that only shows a user to be "affiliated" to an organisation, therefore restricting Service Provider access to all personally identifiable information (PII).

2.1.2 Pilots and Production Deployments

The primary focus for pilot and production deployment was the Research and Scholarship Entity Category (R&S). Promotional activities for R&S, in collaboration with REFEDS, included a set of "New Year's Resolutions" [REFEDS Resolutions], a piece on how to improve the federation experience (upcoming promotion), and work to highlight the importance of releasing attributes to high-profile scientific organisations such as LIGO and CERN. The task also organised a training event at TNC15

[TNC15] for Federation Operators [Attribute Training] on both Research and Scholarship and the Code of Conduct that was attended by 50 participants representing 26 federations. This event proved to be very successful in terms of promotion and as of 29th February 2016 a total of 15 federations, 14 of which were at the training event, now support the Research and Scholarship Entity Category**.**

More generally, the following statistics on adoption of deployed entity categories were recorded during the course of the project:

Date	Number of SPs	Number of IdPs	Federations
May 2015	21	20	Unknown
September 2015	46	39	8
October 2015	51	43	11
November 2015	58	44	11
December 2015	61	47	13
January 2016	65	50	13
February 2016	80	87	15
March 2016	84	95	15

Table 2.1: Number of Service Providers and Identity Providers using the REFEDS Research and Scholarship Entity Category according to eduGAIN metadata.

The number of SPs using the tag is increasing at a good rate, however for R&S to be fully successful more support is needed from the Identity Provider Community by demonstrating their intention to release attributes based on this category.

2.1.3 Challenges

The biggest challenge when supporting Entity Category development is overcoming the risk-averse approach to data protection processes within Identity Provider organisations. This position is understandable, as organisations want to protect the privacy and security of their users and comply with legislation, but can be taken to the extreme to mean that no data is released to providers.

IT departments have also suffered from a lack of investment in the policy side of attribute management. Although the technology may have been installed and well operated, administrators often lack the workflow and permissions to implement processes such as an Entity Category attribute release policy. It is moreover often difficult to identify the right person within any given organisation to take a decision on the approach to be followed.

IT departments also lack the use case to support many of the scenarios that Research and Scholarship supports. Many of the resources that benefit from this approach to attribute release are used by high-end researchers, who are often scattered across multiple organisations. This means that an individual

IT department may only receive support requests from one or two users, making the business case for implementation time hard to justify.

Additionally, defining new entity categories is a lengthy process simply due to the lack of shared definitions within the field of academia. As there is no consistent understanding of terms such as “academic”, “academia”, “affiliation” or event “staff”, work is needed to support research and debate around how these terms might be used and interpreted if used in normative documentation.

2.1.4 Recommendations for Further Work

The adoption of Entity Categories is proving to be an effective, practical and automated yet risk-aware process to meet the requirements of data protection legislation. Service Providers have been keen to embrace this approach and recognise the benefits of handling attributes in this manner. Some Identity Providers have begun to engage with the Entity Category process, but significant work is needed to ensure that support for the process becomes established in federations. Federations also need support to encourage IdPs to trust entity categories, and to implement processes and automation (such as Resource Registries) which make them easier to adopt. Federations without a Resource Registry functionality face a greater challenge in the scalable deployment of entity categories, so this gap will need to be addressed at a national level, with support from GÉANT as needed.

Work to update the eduGAIN Policy framework to provide clearer guidance on Entity Category usage is recommended, alongside ongoing promotional and federation support work in conjunction with REFEDS. It is further recommended that work also continue in this sense on the two newly introduced entity categories.

2.2 Code of Conduct

European data protection laws are commonly seen as an obstacle to the cross-national release of attributes during federated login. The Home Organisations managing end users’ personal data have cited these laws as an excuse for not releasing their end users’ attributes to a relying party (RP) despite the fact that operating within a trusted R&E federated identity environment is already more secure in terms of preserving privacy than using commercial services that request user data. This problem is widely recognised among eduGAIN interfederation service operators.

The Data Protection Code of Conduct is adopted as a means to meet the requirements of the EU Data Protection Directive in federated identity management. It defines behavioural rules for Service Providers wishing to receive user attributes from the Identity Providers managed by the Home Organisations. It is assumed that Home Organisations are more willing to release attributes to Service Providers who express conformance to the Data Protection Code of Conduct.

The GÉANT project started to develop the GÉANT Data Protection Code of Conduct early in 2012, and version 1.0 was published in June 2013, including a comprehensive library of documentation and tools supporting its adoption. Overall, the Trust and Identity Harmonisation Task has helped increase the usage of the Code of Conduct by 225% during GN4-1.

2.2.1 Development

The Article 29 Working Party is an advisory body consisting of the EU member states' data protection authorities as defined in the Data Protection Directive. The working party contributes to the uniform application of the data protection directive. The directive allows trade associations and other bodies representing other categories of data controllers to draw up codes of conducts and to submit them to the working party, which determines whether the drafts submitted are in accordance with the national provisions adopted pursuant to the Directive.

During the GN3plus project, the GÉANT Data Protection Code of Conduct was submitted to the Article 29 Working Party. In June 2015, the project received an initial response from WP29, requesting that the project further elaborate some aspects of personal data processing. The project representatives met the e-government subcommittee of the working party in September 2015 and later held a further dialog with WP29 representatives. WP29 requested that the Code of Conduct be extended to provide more practical advice on how to conform to data protection laws. At present, the code mostly rephrases the wording from the directive, whereas WP29 considers that codes of conduct should bridge the gap between abstract law and practical implementation. This could be achieved by pulling more of the supporting documentation from the wiki [Code of Conduct] into a core document set.

In October 2015, support for the EU/US Safe Harbour arrangement, which had been invalidated following a decision of the European Court of Justice, was dropped from the Code of Conduct. While this had no direct impact on the European version of the Code of Conduct in deployment, it excluded a potential approach from the international version being developed via REFEDs and has made reaching consensus on a scalable approach more difficult.

2.2.2 Pilots and Production Deployments

The subtask continued to roll out the Code of Conduct within the identity federations included in the eduGAIN interfederation service. A training event for federation operators was held at the TNC15 conference [TNC15], covering an introduction to the legal and technical sides of the Code of Conduct and three case presentations from federations with different technical architectures. The subtask also monitored the adoption of the Code of Conduct within eduGAIN. The statistics in the table below are based on the observed Code of Conduct Entity Category and Entity Category support attributes in eduGAIN SAML2 metadata.

Month	Number of SPs	Number of IdPs	Federations
December 2014	38	26	11
April 2015	42	29	-
September 2015	58	42	-
December 2015	76	50	15
March 2016	80	63	15

Table 2.2: Number of Service Providers and Identity Providers using the GÉANT Data protection Code of Conduct according to eduGAIN metadata.

2.2.3 Challenges

Although the number of entities supporting it has doubled in a year, the Code of Conduct suffers from a “chicken and egg” problem. On the one hand, not enough Service Providers are committed to the Code to make the business case clear to Identity Providers, while not enough Identity Providers support it make it an effective process for services.

Encouraging Identity Providers to support the Code of Conduct, as with other entity categories based on attribute release, has proved challenging. Despite campaigns focused on its benefits, including a Code of Conduct endorsement letter signed by four ESFRIs and one research project, many IdPs have not found enough incentive to move forward with implementation. This can be put down to three key areas of concern:

1. Fear over breaking data protection legislation and consequent legal action / fines.
2. Lack of resourcing for access and identity management within IdP organisations.
3. Lack of a clear business case, particularly in research scenarios, due to small numbers of widely distributed requesting users.

The project has adopted a two-layer approach in its dissemination. The project trains the federation operators, who then translate the material into their local language and localize the message and practices, to then provide training to any Identity and Service Providers. This approach is practical and maximises use of the federated identity trust model, but also makes the federation operators key players in the chain, so that if a federation operator does not give priority to the Code of Conduct, the Identity and Service Providers in the country will not use it.

2.2.4 Recommendations for Further Work

Work should be continued with the Article 29 Working Party to keep apprised of the changes it is expecting so that these can be discussed within the GÉANT community. Initial indications are that any requested changes are feasible and likely to result in more harmonised deployment.

In December 2015, the European Commission, Parliament and Council reached an agreement on the General Data Protection Regulation which is expected to become effective in 2018. Although no major changes to the Code of Conduct are foreseen, the new regulation provides some new opportunities for attribute release to third countries that deserve further attention.

As with the deployment of Research and Scholarship, federations should be encouraged to continue to roll out support for the European Code of Conduct and should in turn seek to support their campuses in doing so with technical, organisational and national funding means as appropriate.

2.3 Federation Operation Practices

The aim of this subtask was to collaborate with REFEDS and the EC-funded AARC project to support the work started in the REFEDS Federations Operations Practices on topics such as metadata management and key operations, among others. The subtask’s activities included:

- Participation in the REFEDS Federation Operators group to further identify and refine the practices necessary to guarantee the integrity, availability and confidentiality of the federation operations service provided by the national identity federations.
- Piloting best practice with a subset of federations in the GÉANT eduGAIN community.
- Driving take-up of support for operations best practice within the GÉANT eduGAIN membership by bringing proposals for their adoption to the eduGAIN SG.

The subtask identified four outputs to support the development of Federation Operation Practices:

1. Deliver a proposal on how eduGAIN should work with Entity Categories and further attribute release practices [eduGAIN Policy Review].
2. Propose updates to the eduGAIN core policy set to reflect changes since its last review.
3. Recommend a template Metadata Registration Practice Statement [MRPS] for eduGAIN participants.
4. Review the existing GÉANT Federation Policy Template and make appropriate changes to reflect current best practice.

All of these deliverables have been completed and appropriate recommendations, outputs and change requests made to eduGAIN and REFEDS.

2.3.1 Development

All of the changes and recommendations made in the documents described above were developed by the Trust and Identity Harmonisation team and then socialised on the REFEDS, FOG and eduGAIN mailing lists as community input. The recommendations were presented at the eduGAIN Town Hall meeting in December 2015 [eduGAIN Town Hall] and the proposed Metadata Registration Practice Statement has been subject to a full REFEDS community consultation. The purpose of the Metadata Registration Practice Statement is to document the registration practices of a federation so they can be easily compared with the practices of other federations, a process that is becoming relevant as we move towards an interfederation context.

The proposed changes to the eduGAIN policy documents focused on updating eduGAIN governance references to reflect the development of eduGAIN. This includes changes to improve the attribute release advice and to allow eduGAIN to adopt a completely technology-agnostic approach to ensure that its framework can be used in the future for other technology profiles (such as Moonshot).

The recommendations on the use of entity categories and attribute release are an important part of developing a Best Common Practice for eduGAIN and directing efforts, which in its early years were by necessity more focused on growing membership, towards increasing the quality of the service. These recommendations will form part of the baseline requirements for effective eduGAIN operations, as shown in the diagram in Figure 2.1 below.

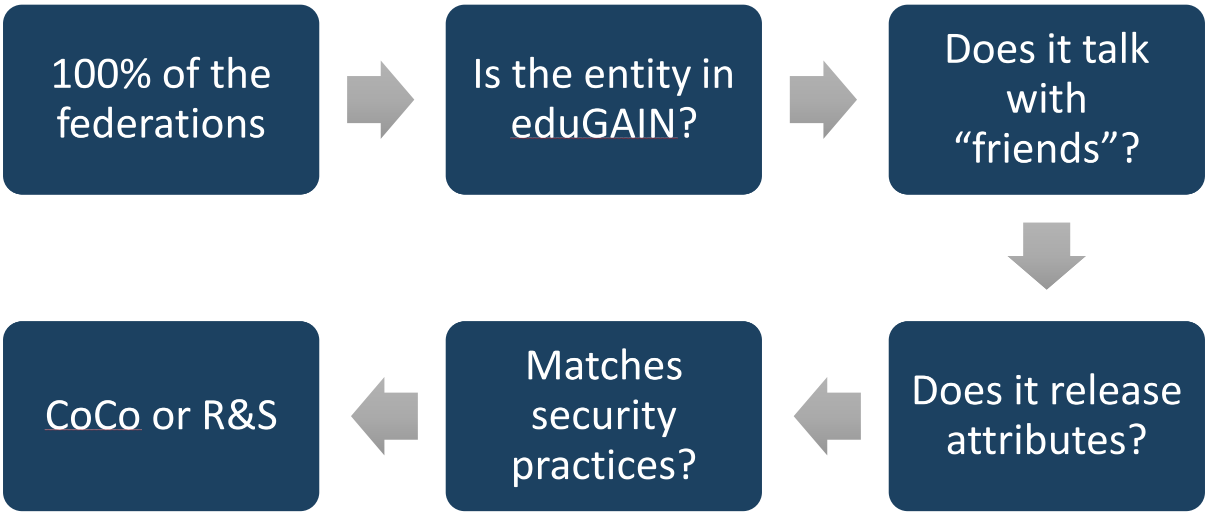
100% of the federations
Is the entity in eduGAIN?
Does it talk with “friends”?
Does it release attributes?
Matches security practices?
CoCo or R&S

Figure 2.1: proposed baseline for effective eduGAIN service delivery.

2.3.2 Recommendations for Further Work

The Trust and Identity Harmonisation team aims to achieve the adoption of the Metadata Registration Practice Statement and has started the review of the Federation Policy Template, which will be completed during the next project. The recommendations on attribute release and eduGAIN policy changes will need continued work to be completed in later projects.

2.4 Service Aspects of Assurance

The “Service Aspects of Assurance” subtask focused on the cost and service issues for Identity Providers to support increasing assurance requirements from research infrastructures. The amount of confidence that can be placed in these features is known as assurance, and is traditionally defined by being assigned a Level of Assurance (LoA). Some research communities tend to have higher requirements than federations in terms of LoA. These requirements were detailed in the FIM4R paper co-authored by various research communities’ representatives [FIM4R].

Awareness is growing in the community that a strictly hierarchical approach to assurance may not be the best suited to meeting its requirements. In traditional “levels” of assurance, each increase is seen as being an improvement on the level below. However it is rare for any given infrastructure or community to need the exact set of requirements defined within any one of those levels, which means unnecessary requirements may be placed on organisations simply because they are included in the same level as those they need.

These organisations may actually prefer to pick and choose requirements from different levels. Therefore, a better approach may be to simply refer to assurance profiles that are scoped against the specific needs of each community, allowing them to be more effectively tailored to their actual requirements. These assurance profiles may overlap in terms of hierarchical notions of strength and leave out certain sections of traditional profiles altogether. This approach is in line with the discussions of the Vectors of Trust (VoT) group within the IETF [VoTWG].

Although the FIM4R paper [FIM4R] was written in 2013, the issue of assurance remains open to this day. Some federations, e.g. InCommon (USA) and SWAMID (Sweden), have made attempts to roll out assurance profiles, which however have not produced definitive results. In order to help IdPs within eduGAIN as well as research communities, GÉANT and the Authentication and Authorisation for Research and Collaboration (AARC) project are approaching the issue from two different angles. GÉANT is addressing the aspects relating to federations and IdPs, with the objective of reducing their costs, while AARC is looking at the question from the perspective of research communities and SPs.

These two approaches, which are complementary, were presented in a White Paper [Assurance Paper], which addresses the service aspects of Levels of Assurance, outlining the findings of the surveys carried out of the federations and IdPs, as well as the results of internal dialogues with the IdPs. These results were then compared with the insights provided by AARC [MNA3.1], to draw up a set of recommendations.

2.4.1 Development

The subtask carried out an analysis of the existing state of assurance to determine feasibility of a service. The federation landscape in terms of assurance appears to be quite diverse. Federations differ in the way in which they assist IdPs. Some, such as Haka (Finland) and SURFnet (Netherlands), offer a self-assessment tool to measure the maturity of IdPs, while many others only define minimum requirements. The IdPs of the Danish federation WAYF are considered public organisations, and are required to comply with IT security rules (ISO 27001) and be audited.

The various definitions of assurance also differ. The German federation DFN-AAI established its own LoA with two classes, Basic and Advanced, focusing on three different aspects (registration and original identification, update of information, and authentication). InCommon, the US federation, adapted NIST Special Publication 800-63 [NIST SP 800-63], calling the two lowest levels Bronze and Silver. The Swedish federation SWAMID has started to roll out assurance based on the Kantara Identity Assurance Framework (IAF) that introduces Assurance Levels (ALs). SWAMID AL1 is a subset of Kantara IAF AL1.

In parallel, the AARC project [MNA3.1] conducted guided interviews with SP and research community representatives to gain a greater understanding of requirements from the perspective of the services. The survey helped to verify a proposed assurance baseline as the minimum standard for research communities, including the following requirements:

- Individual accounts (i.e. no shared accounts).
- Persistent, non-reassigned identifiers.
- Documented identity vetting, not necessarily face-to-face.
- Password authentication including some good practices.
- Departing user’s ePA changes promptly.
- Self-assessment of LoA supported with specific guidelines.
- Incident response in a later step.

Comparing the results of the questionnaire against the baseline, the following areas for improvement were identified:

- **Non re-assigned identifiers:** although persistent identifiers (such as eduPersonPrincipalName) are used, many Identity Providers currently reassign them.
- **Documented identity vetting:** although IdPs have a vetting process in place, this is not always documented.
- **Departing user's ePA changes promptly:** the time it takes to change user data is between 2 weeks and 6 months. As closing accounts depends on internal processes and some universities have alumni accounts, the eduPerson(Scoped)Affiliation should be updated within 1 month.
- **Self-assessment of assurance supported with specific guidelines:** in order to set guidelines, a template needs to be designed, which can then be used for self-assessment.
- **Incident response:** Security Incident Response Trust Framework for Federated Identity (SIRTFI), but only as a later step, since SIRTFI as well as the related minimum assurance requirements have only recently been introduced. SIRTFI enables coordination of security incident response across federated organisations.

2.4.2 Challenges

The implementation of assurance seems in general problematic. In Germany, although requirements are low, they still prove too high for many IdPs, resulting in a simplification of requirements on the part of federation operators. In the US, even though a formal LoA was introduced, five IdPs managed to obtain the Bronze certification and only one IdP received the Silver. The Swedish federation had more success while introducing Kantara AL2 for eduID.se thanks to a national requirement linked to student registration. Nevertheless, the cost for its introduction was significant, at between €20-50 per account just for eduID.

Even relatively small countries such as Switzerland have hundreds of thousands of accounts within the SWITCHaai federation. Consequently, IdPs also have costs and need manpower to reach a higher LoA, so assurance as a default feature would do away with the value proposition of federated identity for many such organisations.

A common problem that exists in today's R&E federation landscape is that it is difficult to quantify the specific degree of assurance that is offered by federations. All federations provide a baseline of assurance through the practices articulated in their policies, operational practice guides and technical requirements for participants. Federations are also expected to reach standards through participation in initiatives such as [eduGAIN]. There is however no consistent mapping of these expectations.

This problem is recognised by the community and parallel efforts to better describe and document the baseline practices of federations is underway, with a proposed Metadata Registration Practice Template being prepared for eduGAIN participants. Improved mapping of the baseline behaviour of federations will in turn inform a better understanding of the baseline expectations for IdPs.Recommendations for further work
From the above analysis, the following recommendations for assurance can be derived:

- Work should be carried out in collaboration with AARC before it ends in 2017, to finalise and document a baseline assurance profile for IdPs, mapped to the requirements identified by

AARC. Work on best practice in some areas (e.g. password authentication practice) is still required.

- The creation of a self-assessment template / tool should be investigated, with the best recommended being a GÉANT web tool, combined with SIRTFI and other monitoring/testing tools, including recommendations and best practices. This would help address documentation requirements from the AARC minimum set. Collated information would then make it easier for research providers to see where their requirements are met.
- The status quo of reassignment of identifiers should be addressed by working with REFEDS to survey the problem space and possible solutions to this embedded practice by organisations. This may include changing general recommendations on identifiers in current common documentation.
- Work with REFEDS and AARC on maturing the SIRTFI approach to incident response should be continued.
- Peer (pairwise) auditing should be implemented for those IdPs that need to document their approach against the GÉANT assurance profile in order to verify compliance, with lower costs than external audits. The results of these audits can be displayed in the web tool described above.
- Second-factor authentication should be considered: GÉANT could offer this as a service, provide a best practice profile within eduGAIN or procure a Duo-type solution for the community sectors that need it. This would require a separate business case.
- GÉANT should develop federation maturity reports aimed at managers, helping to improve the maturity of federations at a general level and to support the federations in attaining the manpower and funding to reach these goals.
- eduGAIN should work to ensure that all federation operators subscribe to a Metadata Registration Practice Statement [MRPS] that complies with the recommended standard template.

2.5 Interoperability

The aim of this subtask was to develop sustainable approaches to interoperability with e-Gov, social media identity frameworks and other lifelong identity initiatives, especially those of NRENs. The work focused on:

- Completion of pilot use cases with STORK2.0 and a roadmap towards sustainable long-term implementation based on experience gained during GN3plus.
- Engagement with standards areas, such as Kantara on behalf of eduGAIN.
- Development and use of results from GN3 JRA3 to improve interoperability within eduGAIN as appropriate, e.g. Federation Lab [Fed Lab].

Secure idenTity acrOss boRders linKed [STORK] and its evolution STORK 2.0 [STORK2] were EU co-funded projects under the ICT Policy Support Programme of the Competitiveness and Innovation Framework Programme (CIP). The aim of the STORK project was to establish a European eID Interoperability Platform pilot to allow citizens to establish new e-relations across borders just by presenting their national eID. STORK 2.0 aimed to contribute to the realisation of a single European electronic identification and authentication area.

The STORK 2.0 consortium was composed of 19 different countries (Austria, Belgium, Czech Republic, Estonia, France, Greece, Iceland, Italy, Lithuania, Luxembourg, Netherlands, Portugal, Slovakia, Slovenia, Spain, Sweden, Switzerland, Turkey, and the United Kingdom) and provided a distributed pilot infrastructure to enable cross-border interoperability of existing national electronic identity schemes. A few public authorities usually performed the role of Identity providers in STORK2.0 in each country (from 1 to 80 according to how each Member State' implements their model) for a total of over 130 IdPs at European level.

At the end of STORK 2.0, the outcomes of the STORK 1.0 and STORK 2.0 projects were passed on to the Connecting Europe Facility [CEF], which is responsible for defining the building blocks for eID, eSIGNATURE, eDELIVERY, eINVOICING and eTRANSLATION within eIDAS. STORK 1.0 and 2.0 are both based on SAML2.0, the same protocol used in the academic federations in eduGAIN. As SAML 2.0 leaves implementers many options, (e.g. on how to pass attributes, what bindings to use, how to use PKI, what should be signed and what should be encrypted), the actual implementations in eduGAIN and STORK are not interoperable out of the box.

The Interoperability subtask of the Trust and Identity Harmonization task, looked at the differences between the SAML2Int standard and the STORK2.0 specification and piloted possible scenarios for the interoperability between the implementations. As part of its work on interoperability within eduGAIN, the subtask made use of the Federation Lab – a test suite for SAML interoperability developed as part of GN3 JRA3. In addition, the Trust and Identity Harmonisation Team also focused on the current behaviour within federations as relating to authentication statistics. The results of this work were included in a white paper describing the current status of Federation Statistics [Statistics Paper].

2.5.1 Development

During GN4-1, development work started by JRA and SA during GN3plus was completed on the pilot use cases with STORK 2.0. Two high-level use cases were identified and developed:

- Use case 1: A user with a STORK eID accesses a service provided by an SP in one of the eduGAIN Federations.
- Use case 2: A user with an account on an IdP in one of the eduGAIN Federations accesses a service provided by STORK in one of the STORK-enabled member states.

To enable interoperation, a new proxy element, eduPEPS, was added. eduPEPS can be situated between a STORK federation and an eduGAIN federation and acts as a protocol translation bridge, while at the same time providing a trusted entity for both the STORK and the eduGAIN federations.

The eduPEPS service is comprised of two components, the eduPEPS Proxy and the eduPEPS Attribute Translation Library. The eduPEPS Proxy provides the necessary interfaces to enable it to act as a proxy between entities in the eduGAIN and STORK Federations. In academic federations, the eduPEPS service can participate as a virtual SP or IdP, with a built-in Discovery Service, depending on whether users are authenticated on an academic IdP or using their governmental eID. In a STORK infrastructure, eduPEPS can participate as a PEPS or IdP, or as an Attribute Authority or SP, depending again on the actual user flows and the chosen deployment scenario. The eduPEPS Translation Library is an integral part of the eduPEPS service and provides configurable attribute mappings. Primarily it is used to map between the eduPerson schema and the STORK attribute schema for Academia, but it can easily be configured to map between any kinds of attribute schemas.

The Trust and Identity Harmonisation task has also been working with FedLab, which was developed as part of previous JRA work during GN3, to pilot its approaches to testing entities for interoperability and compatibility. The aim of this work was to define recommendations on future service models for FedLab as well as to improve interoperability within eduGAIN.

One of the most often requested features for FedLab that is not found in the current toolset and does not fit well with the scope of that toolset, is the ongoing need for federation statistics and high-level monitoring of the number of authentications carried out by federations. Obtaining such data is problematic for most federations as this information is typically only captured at the institutional level.

The subtask carried out an initial review of the current status of statistics gathering within federations [Fed Statistics], discussing with federation operators their current and future needs and potential developments in this respect. The results were included in a white paper on Issues and Solutions for SAML Identity Federation Statistics [Statistics Paper] which identifies some pragmatic first steps to demonstrate the value of statistics and thereby encourage greater engagement with the topic.

2.5.2 Pilots and Production Deployments

The pilot eduPEPS service was deployed and tested in a testbed environment comprised of:

- An SP using the SAML2Int profile operated by GRNET in Greece.
- An SP using the STORK SP interface operated by the University of Murcia in Spain.
- An IdP connected to the GRNET Academic Federation.
- A test STORK PEPS installation operated by the University of Murcia in Spain.
- The pre-production STORK PEPS installation operated by the Ministry of Interior and Administrative Reconstruction in Greece.

Two deployment scenarios for the eduPEPS service have been identified. In the first scenario, the eduPEPS service is deployed as a bridging element between the academic federations in eduGAIN and the STORK infrastructures in the European member states. In the second scenario, the eduPEPS service is deployed within the member states and acts as a gateway between the STORK infrastructure and the academic federation in each member state.

Both these scenarios have pros and cons, but they allow the interoperation between the academic federations in eduGAIN and the STORK infrastructures without having to introduce multiple software stacks for federated access at the participating SPs.

During development of the eduPEPS service, the STORK PEPS pilot implementation presented a moving target. Until the last month of the STORK 2.0 project, new functionality was being added to the STORK 2.0 PEPS, which in many cases was not backwards compatible with the previous versions of STORK. Furthermore, some of the most advanced scenarios that were described in STORK 2.0, such as cross-border attribute aggregation, were never tried in real-life implementation.

Under these circumstances, the proof-of-concept implementation of the eduPEPS service had to rely on assumptions about the STORK 2.0 functionality that could not be fully tested. Furthermore, when the STORK 2.0 project ended in September 2015, it handed over its outcomes to CEF, which is responsible for the definition of the building blocks for eID, eSIGNATURE, eDELIVERY, eINVOICING and eTRANSLATION within eIDAS. The eID building blocks for eIDAS are expected to follow the STORK architecture to a certain extent, especially for cross-border authentication interoperation.

Fedlab [FedLab], a highly valued tool for federation operators, is not a clean fit for the GÉANT's Product Management Processes' production-oriented methods. This is due to the unique scope and impact of the toolkit, and the need for it to be constantly innovated. GÉANT partners with REFEDs and the Vietsch Foundation [Vietsch] to help ensure the tool's continued availability and evolution.

Coherent statistical information remains a significant issue for federations and federation funders, and at the time of writing is not an area that is mature enough for an interoperable approach. Individual national approaches are sparse and do not standardise the type of information requested, and different architectures pose radically different challenges for data collection and aggregation.

2.5.3 Recommendations for Further Work

The interoperability subtask item has drawn up the following recommendations:

- The SAML F-TICKS specification should be completed to make it deployable.
- A central service to aggregate statistics from federations, in the same model as monitor.eduroam.org, should be developed, starting with hub-and-spoke/centralised federations as a pragmatic baseline.
- Further work should be done to develop RAPTOR based on the Edugate approach used in Ireland. This depends heavily on increasing uptake of the RAPTOR tool by organisations.
- More effort should be directed at enabling IdPs to commit time to developments such as support for aggregated statistics. GÉANT should look at producing an accessible set of recommendations for IdPs' local development plans for use by senior management at campuses, highlighting the areas where time and effort should be invested over the next one or two years.
- REFEDS / eduGAIN should look at tracking changes in entity information over time in MET and eduGAIN tools.
- Work with the STORK 2.0 project has shown that the technical interoperation between the academic federations in eduGAIN and the eGOV eIDs services is achievable. As eIDAS is being implemented by the member states, close collaboration with the eIDAS Task Force and CEF towards achieving policy and technical interoperation at an infrastructural level should continue.

3 Conclusions and Recommendations

During its year of operation, the Trust and Identity Harmonisation task has had an impact on federation harmonisation that is measurable in part in terms of adoption figures, but also by the development of new guidelines and tools and demonstrations that are able to support future interoperability work. There are clearly many challenges still to be faced by federations and eduGAIN as they grow, but the task has laid strong groundwork that can be built upon in the areas of future services, service development and research activities.

The recommendations set out in the main sections of this report have been fed into the development of future GÉANT projects, AARC, REFEDS and other AAI initiatives. In addition to the support from these projects and initiatives, ongoing funding and support for federation operators, identity providers and service providers will clearly be needed if they are to deliver a cost-effective, secure and efficient AAI for e-Infrastructures in the future.

References

[AARC]	https://aarc-project.eu
[AAI Study]	https://www.terena.org/publications/files/2012-AAA-Study-report-final.pdf
[Assurance Paper]	http://www.geant.org/Resources/Documents/Service%20Aspects%20of%20Assurance.pdf
[Attribute Training]	https://wiki.edugain.org/AttributeReleaseTraining2015
[CEF]	https://ec.europa.eu/digital-single-market/en/connecting-europe-facility
[Code of Conduct]	https://wiki.refeds.org/display/CODE/.
[eduGAIN]	https://www.edugain.org
[eduGAIN Policy Review]	https://wiki.geant.org/download/attachments/45844969/SA5T1-edugain-policy-changes.pdf.
[eduGAIN Town Hall]	https://wiki.edugain.org/EduGAIN_Town_Hall-20151201.
[FedLab]	http://fed-lab.org
[Fed Statistics]	https://wiki.geant.org/display/gn41sa5/Current+Federation+Statistics.
[FIM4R]	http://cds.cern.ch/record/1442597/files/CERN-OPEN-2012-006.pdf?version=2
[InCommon]	http://www.incommon.org/federation/categories/index.html
[MNA3.1]	https://aarc-project.eu/wp-content/uploads/2015/11/MNA31-Minimum-LoA-level.pdf
[MRPS]	https://docs.google.com/document/d/1CrvhSPfWRc_afILKHPURGbNuasKMyzLfnuFARN3CFp4/
[NIST SP 800-63]	http://csrc.nist.gov/publications/PubsSPs.html#800-63-Rev1
[REFEDS]	https://refeds.org
[REFEDS HfD]	https://refeds.org/category/hide-from-discovery/
[REFEDS R&S]	https://refeds.org/category/research-and-scholarship/
[REFEDS Resolution]	https://refeds.org/resolutions.
[Statistics Paper]	https://wiki.geant.org/download/attachments/45844980/Issues%20and%20Solutions%20for%20SAML%20Identity%20Federation%20Statistics.pdf
[STORK]	https://ec.europa.eu/digital-single-market/en/content/stork-take-your-e-identity-you-everywhere-eu
[STORK2]	https://www.eid-stork2.eu/
[TNC15]	https://tnc15.terena.org/
[Vietsch]	http://www.vietsch-foundation.org
[VoTWG]	https://www.ietf.org/mailman/listinfo/vot

Glossary

AAI	Authentication and Authorisation Infrastructure
ABFAB	Application Bridging for Authentication Beyond the Web working group
CEF	Connecting Europe Facility
ECP	Enhanced Client or Proxy
eduGAIN	Global Authentication Infrastructure
eduroam	A global service that provides secure roaming connectivity.
eID	Electronic Identification
eTS	Electronic Trust Services
IdP	Identity Provider: the federation entity that issues assertions about identity on behalf of end users who use them to access the services of service providers (SPs).
FIM	Federated Identity Management
IETF	Internet Engineering Task Force
IGTF	Interoperable Global Trust Federation
LoA	Level of Assurance
MDS	Metadata Distribution Service (The Metadata Aggregator used within eduGAIN)
NREN	National Research and Education Network
OLA	Operational Level Agreement
RP	Relying Party: a federation entity that evaluates an assertion from an Identity Provider and uses the information from the assertion for controlling access to protected services. Often a synonym for an AAI-enabled application.
SAML	Security Assertion Markup Language
SP	Service Provider
SSO	Single Sign-On (can also be described as WebSSO or Unified SSO for hybrid approaches)
STORK	Secure idenTity acrOss boRders linked
WP29	Article 29 Working Party
VoT	Vectors of Trust