30-04-2017

Deliverable DNA3.4: Recommendations on the grouping of entities and their deployment mechanisms in scalable policy negotiation

Deliverable DNA3.4

Contractual Date:	30-04-2017
Actual Date:	30-04-2017
Grant Agreement No.:	653965
Work Package:	NA3
Task Item:	NA3.4
Lead Partner:	STFC
Document Code:	DNA3.4
Authors:	David Kelsey (STFC), Licia Florio (GÉANT), David Groep (Nikhef), Mehdi Hached (Renater), Christos Kanellopoulos (GÉANT), Mikael Linden (CSC), Ian Neilson (STFC), Stefan Paetow (Jisc), Wolfgang Pempe (DFN), Mischa Salle (Nikhef), Hannah Short (CERN), Uros Stevanovic (KIT), Gerben Venekamp (SURFsara)

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

Abstract

Work Package NA3 (Policy and Best Practices Harmonisation) was charged with producing recommendations and defining best practices to implement a scalable and cost-effective policy framework for an integrated AAI eco-system. Sub-task NA3.4, in particular, deals with aspects of scalable policy and trust negotiation between the various entities involved in such an AAI. This deliverable report starts with a study of entity categories and the status of take-up by the eduGAIN inter-federation. The new "Snctfi" policy and trust framework has been developed for applying policies and best practices to an e-Infrastructure or Research Infrastructure using IdPs in the R&E federations via an SP-IdP proxy. After describing policy work for HNSciCloud and for the RCauth CA, we finally present our conclusions and future plans.

Table of Contents

Executive Summary

This deliverable presents the work of task NA3.4 during the two years of the AARC project. Work Package NA3 (Policy and Best Practices Harmonisation) was charged with producing recommendations and defining best practices to implement a scalable and cost-effective policy framework for an integrated authentication and authorisation infrastructure (AAI) eco-system. These recommendations and best practices have been developed with the objective of adoption by and integration into the existing AAI federations (through eduGAIN), e- infrastructures (such as EGI and PRACE) and the many Research Infrastructures, all following the recommendations of the JRA1 Blueprint Architecture [BPA]. Sub-task NA3.4, in particular, deals with aspects of scalable policy and trust negotiation between the various entities involved in such an AAI.

A general classification of entities in the identity and attribute ecosystem consists of the following three main entity types: identity providers (IdP), service providers (SP) and attribute authorities (AA). Because of the large number of such entities involved the bi-lateral negotiation of policies or contracts between SPs and IdPs/AAs does not scale. Other solutions must be used.

In addressing these challenges, AARC has used three complementary approaches: promotion of the use of 'entity grouping' and its associated specifications; the definition of a framework to organise service provider collectives and permit these to express their common qualities in a standard way; and the development of a concrete set of policies and practices for an SP-IdP proxy that enables researchers with an existing federated IdP to connect to existing e-Infrastructures without requiring modification of either.

The document starts by presenting the results of a study conducted in year 1 of the project. The study investigated the status of the adoption of entity categories within the eduGAIN inter-federation together with the experiences of three national federations. It goes on to consider what recommendations can be made to improve the relatively slow take-up. Section 2 also reminds the reader of the work done by AARC NA3 task NA3.2 [DNA3.2] on a new security incident response trust framework [SIRTFI].

The main activity of task NA3.4 during year 2 of the project has been the development of a new trust framework for e-Infrastructures and/or Research Infrastructures to be used when they are connected to the R&E Federation world via an SP-IdP proxy, or AAI Gateway. This framework, called the "Scalable Negotiator for a Community Trust framework in Federated Infrastructures" or "Snctfi", defines a mechanism for Infrastructures to express policy alignment across all of its constituent services and communities. Based on the Security for Collaborating Infrastructures framework [SCI], it takes specifically those elements related to infrastructure policy coherency and uses it to allow the proxy to assert (SAML) Entity Category attributes and assurance qualifiers, such as REFEDs Sirtfi and the GÉANT Data Production Code of Conduct [ENTCAT].

Other policy work has been performed by NA3.4 including advice requested from and provided to the HNSciCloud project [HNSCICLD] and policy work for the RCauth CA [RCAUTH].

The future challenges for Snctfi are twofold: work to address the 2018 EU General Data Protection Regulation on incorporating mechanisms that allow Infrastructures to collectively express compliance on behalf of all back-

end services and to do this for Infrastructures and collaborations that may not in themselves be legal entities; and secondly, to promote the wide adoption of the framework within both the generic e-Infrastructures and the Research Infrastructures. To this end, international groups such as the Interoperable Global Trust Federation [IGTF], the Wise Information Security for collaborating e-Infrastructures [WISE] community, and the Research Infrastructure requirements alignment group on federated identity management [FIM4R] appear to be the most appropriate forums. In the AARC2 project, these groups have been identified as partners for future development of the Snctfi framework.

1 Introduction

AARC Work Package NA3 (Policy and Best Practices Harmonisation) was charged with producing recommendations and defining best practices to implement a scalable and cost-effective policy framework for an integrated authentication and authorisation infrastructure (AAI) eco-system. These recommendations and best practices have been developed with the objective of adoption by and integration into the existing AAI federations (through eduGAIN), e-Infrastructures (such as EGI and PRACE) and the many Research Infrastructures, all following the recommendations of the JRA1 Blueprint Architecture [BPA]. Sub-task NA3.4, in particular, deals with aspects of scalable policy and trust negotiation between the various entities involved in such an AAI. The work done by task NA3.4 is presented in this deliverable document.

A general classification of entities in the identity and attribute ecosystem is simple and consists of the following three main entity types: identity providers (IdP), service providers (SP) and attribute authorities (AA). The bi-lateral negotiation of policies or contracts between SPs and IdPs/AAs does not scale. The AAI eco-system must be able to use policy templates, entity category/assurance qualifier specifications and/or accepted definitions of best practice to deploy a scalable and cost-effective policy and trust framework.

Due to the then current state of policy development at the start of the AARC project, defining effective operating models needed realistic policies to be developed, and for the policies to have a firm grounding in the community to which they apply. AARC is not, in itself, in a position to adopt policies, but instead mediates policy, proposes best practices and performs pilot-testing of the various technologies between a multitude of existing entities: Research Infrastructures, e-Infrastructures, IdPs, R&E Federations and SPs. Community engagement to gain such adoption has been essential. Much of the work in scalable policy negotiation has therefore involved outreach and participation in existing bodies, including REFEDS, IGTF, and FIM4R. Feedback received from these bodies has been included in the policies and frameworks developed.

We started with an investigation of entity categories within the eduGAIN inter-federation together with their deployment at that time (see Section 2). The experiences of three national federations were also sought and documented. We subsequently considered what recommendations could be made to help improve the slow take-up. Section 2 also reminds the reader of the work done by AARC NA3 task NA3.2 on the new Sirtfi security incident response trust framework, reported in detail elsewhere [DNA3.2]. Section 3 presents the main activity of task NA3.4 during the second year, i.e. the development of a new trust framework for e-Infrastructures and/or Research Infrastructures to be used when they are connected to the R&E Federation world via an SP-IdP proxy, or AAI Gateway. This policy and trust framework is called the “Scalable Negotiator for a Community Trust framework in Federated Infrastructures” or “Snctfi”. In section 4 we present other policy work performed in the task and finally our conclusions and future plans are presented in section 5.

2 Entities, groupings and scalable policies

This section presents a study conducted within task NA3.4 at the start of the AARC project, to investigate the status of the adoption of entity categories within the eduGAIN inter-federation.

The study focused on trans-border service providers (SP) that are republished in eduGAIN metadata. Studying the adoption of categories and policies of those SPs at that time helped identify the successful uses and the remaining challenges for upcoming SPs. It considered the extent to which the categories or policies had been adopted and used in practice and what could be the obstacles to further deployment. This study gave a view at that time of the remaining challenges for scalable policies that would make federated services more accessible for research communities.

2.1 Entity categories and attribute release challenges

Attribute release and trust between IdPs and SPs in identity federations are the key elements for a successful collaboration. Trust is settled through the federation's legal framework to which providers are bound. Attribute release on the other hand is a more complex problem, both technically and, for regulatory issues such as privacy and personal data protection, legally.

To overcome this problem, several attempts to categorise service providers have emerged in the recent years. Some categorisation involves data minimisation and purpose limitation on attribute release while others try to group service providers based on their business profile and common attributes requirements.

The most significant efforts in the policy area are the "GÉANT Data Protection Code of Conduct" [DPCOCO], either in its European version or in what was at that time a future international form, and the "REFEDS Research and Scholarship" category [R&S]. Other categories were already being considered or were in an advanced discussion (Academia, library...) [EC CAND]. Whilst the DP CoCo is more than just a SP category, on the metadata level, service providers and the agreeing identity providers are tagged as a SAML 2.0 Entity Category.

2.2 eduGAIN exploration tools

In order to find the current status, NA3 used the tools that eduGAIN offered at the time:

- Metadata exploration tool: https://met.refeds.org/
- eduGAIN entities search tool: https://technical.edugain.org/entities.php

This study was constrained to the publicly available information using these eduGAIN tools and reading the available aggregated metadata.

2.3 Statistics of entity category use in 2015

Using the tools described above, we could easily count the number of entities that had endorsed the R&S category or the GÉANT DP Code of Conduct. The figures are edifying, only 191 entities (without removing duplicates), out of a total of 2385 entities at that time (in 2015), were announcing the category/policy support. See Section 2.8 for some statistics in April 2017.

GÉANT (EU/EEA) Data Protection Code of Conduct version 1.0:

- 105 entities in total
  - 42 IdP
  - 63 SP

REFEDS Research & Scholarship Category:

- 86 entities in total
  - 39 IdP
  - 47 SP

Even though the GÉANT DP CoCo and the REFEDS R&S had been finalised for about a year, it was clear that their adoption was still very slow within the community.

The two documents are short and clear, and so we recommended that the national federation operators should help their deployment by:

- Advertising them toward the SPs;
- Translating them into the country's language if possible;
- Implementing a way of tagging the entities that support categories during the entity's recording (federation registry).

In practice, the IdPs are technically able to release the necessary attributes to the categorized SPs, in cases where they (the SPs) are legally and reliably processing the attributes.

The two categories described above do not cover all types of SPs (e.g. commercial, business models...) or the communities that run them (physics, libraries...). Other categories may be required but an effective way of presenting them and assistance for SPs to find the right category are both needed.

Some federations have already implemented national categories to help IdP managers to easily scale their attribute release mechanisms. Grouping SPs is the correct approach to avoid per-SP attribute filtering. Such categorisation requires clear definitions, and federation operators should be responsible for advertising and implementing the categories at the federation registry level. Indeed, categorizing SPs should also lead to common attribute requirements. These attribute subsets must be negotiated with SP managers to reduce the number of mandatory attributes and tag the possible other attributes as optional.

2.4 The Federation Éducation-Recherche experience (FR)

The French academic federation introduced SP categories in 2011. For each category, they distributed a set of attribute filters that automated the attribute release at the IdP level. They noticed that IdP managers are generally not keen to regularly update their IdP configurations. Indeed, there is always a risk of breaking the service. The vast majority started to simply use the filters containing all the categories even if it is not good to do this in terms of personal privacy.

Because the Shibboleth software was based on SAML2 specifications, the implementation allowed dynamic attribute release "negotiation", i.e. exploiting <md:requestedAttribute> elements [SAML2]. The French federation administrators will progressively abandon the automated filters distribution and rely more and more on that feature. But, even with this feature, IdP managers and data privacy officers want to have a way of distinguishing between SPs. The GÉANT EU/EEA DP CoCo is a way to scale attribute release. Federation operators should, in the first place, adopt this policy.

2.5 The DFN-AAI Experience (DE)

The Identity Federation operated by the German Research and Education Network (DFN) introduced Entity Categories (ECs), both for SPs and IdPs, in 2012 in order to support so-called "Virtual Sub-Federations". The setup is based on a whitelist maintained by a specific project or community and which is linked to the metadata registry. The project-specific EC is only available for entities listed on such a whitelist - a nightly check removes the EC automatically if an entity disappears from the respective whitelist. Using such an EC, (Shibboleth) SPs are able to select all project-related IdPs from the federation metadata and ignore the rest, while IdPs only have to set up one Attribute Filter Policy in order to release Attributes to a dynamic number of project-related SPs. This concept turned out to be quite popular; meanwhile (in 2015) three of these ECs are in use, and a fourth one has recently been requested.

The DP CoCo EC was introduced in July 2013, and R&S in 2015. While many SPs registered with the DFN-AAI committed specifically to the DP Code of Conduct, the acceptance by German IdPs is still not high. One reason for the reluctance of German IdP admins to support the DP CoCo and R&S ECs is the strictness and complexity of data protection laws and regulations in Germany cf. http://dariah-aai.daasi.de/attribute-release_and_legal-stuff_wp.pdf

2.6 The Greek Federation experience (GR)

The Greek Federation, operated by GRNET, has introduced Entity Groups in the published metadata, utilizing multiple EntitiesDescriptor elements. The groups match SPs with different trust levels in the federation (GRNET's own services SPs, Microsoft services for higher educational institutes, others) but are not formally defined. Since GRNET operates the majority of the IdPs of the Universities participating in the federation, respective attribute release policies have been deployed in the Identity Providers, utilizing AttributeRequesterInEntityGroup type rules for matching the SP and releasing the necessary attributes.

As this set up is neither optimal nor easily maintainable, the Greek federation is in the process of introducing a number of national Entity Categories (both for SPs and IdPs) in the federation. The main drivers for the change are:

- Simplicity in the federation metadata aggregation and publication
- Formal definition of Trust Levels
- Enhanced granularity
- The introduction of new Identity Providers (e.g. hospitals) in the federation that have stricter requirements for attribute release.

The work was ongoing and the preliminary plan was to introduce the Entity Categories in 2016. No decision had been made yet as to whether eduGAIN defined Entity Categories (GÉANT DP CoCo, R&S) will be used.

2.7 The issue with attribute values

Some community SPs, like libraries or research communities, need the provision (release) of some attributes (e.g. affiliation, entitlement, isMemberOf...) often having specific values tailored to those SPs. This scenario is the most complex to implement. Indeed, doing so forces the IdP manager to intervene and update the values in the IdP back-end (affiliations) or to dynamically generate a new value (entitlements) for a given new SP.

Defining categories for attribute values should also be considered to ease the deployment at IdP level. These categories would restrict or make use of common attributes and values for a given group of SPs or rely during the authorisation process on Attribute Authorities' (community IdPs) additional (custom) attributes.

2.8 Entity categories - conclusion

Defining common characteristics or standards that make a category or a policy a good and scalable "framework" is necessary. In the meantime, documentation and training should be improved within national federations and followed at the eduGAIN level. Community SPs and their requirements should be well defined and categorized before discussing with IdP managers.

Detailed analytics is now provided routinely by GÉANT and eduGAIN (through technical.edugain.org) allowing AARC to track developments in this area. Use of SAML-EC attributes has increased significantly over the past 2 years, and the adoption shows uptake of Sirtfi (137 entities with the SAML-EC "Sirtfi" tag within 1 year) and REFEDS R&S (646 entities out of ~4000) with an increasing overlap: 57 assert both R&S and Sirtfi as of April 2017, a combination that indicates the IdP meets the baseline requirements for Infrastructure access.

2.9 Security Incident Response Trust Framework (Sirtfi)

The concept of the Security Incident Response Trust Framework for Federated Identity [SIRTFI] was first proposed by FIM4R (Federated Identity Management for Research) in their 2012 paper [FIM4R]. It addressed

the Research Communities' need for a security incident response capability in the authentication infrastructure provided by identity federations and inter-federation. Work by the REFEDS Sirtfi Working Group [SIRTFI WG] has gained community-wide support for the framework and its use within federations. Federation entities are able to adopt Sirtfi by asserting each and every requirement of the framework and including appropriate extensions in federation metadata. There is ongoing work by the AARC project to develop a generic, scalable procedure for security incident response in identity inter-federation, based on Sirtfi [DNA3.2]

Sirtfi is an entry requirement for IdPs accessing multiple Service Providers owned by Research Communities, including CILogon, RCAuth and CERN. To enable users to contact their organisations directly and request that they become Sirtfi compliant, a helper tool was developed at http://sirtfi.cern.ch/. It is likely that this tool will spawn a project in GÉANT to provide users with a convenient channel of requesting that their home organisation adopt certain frameworks or release necessary attributes to grant access to service providers.

For a group of entities to support Sirtfi, it is necessary that the entity exposed to the identity federation assert Sirtfi requirements on behalf of the group. This includes providing a security contact able to satisfy security incident response responsibilities for the entire group.

3 Scalable Negotiator for a Community Trust Framework in Federated Infrastructures (Snctfi)

Studies in AARC [DJRA1.1] and experiences reported in [FIM4R] have shown that the majority of large-scale research communities have established their own SP-IdP proxies. Such a proxy is an excellent way of connecting a Research Infrastructure (RI) to the R&E federations and to eduGAIN. This SP-IdP proxy model (see figure 1), also known as an AAI Gateway, is well described in the AARC Blueprint Architecture [BPA] and identified as a recommendation for research collaborations engaging with R&E federations.

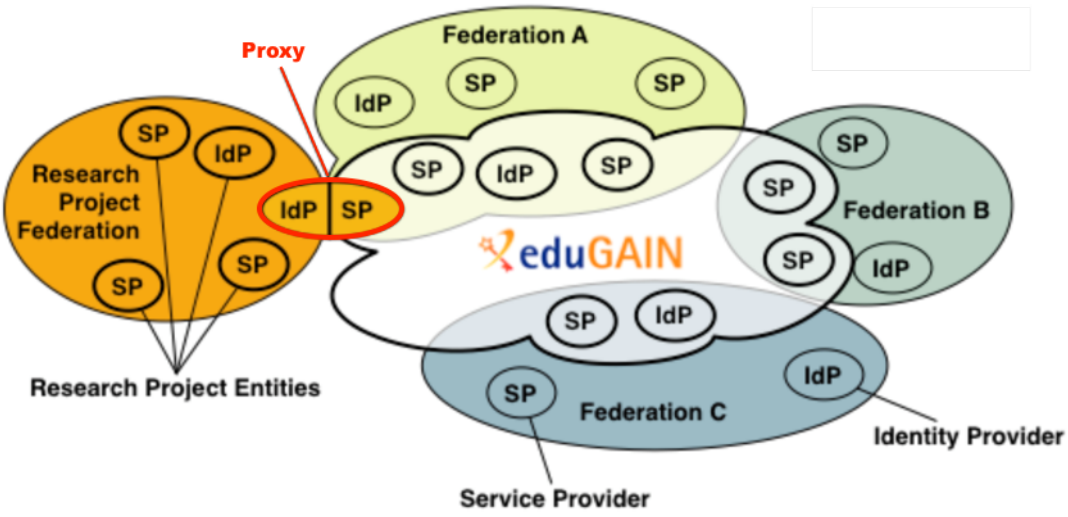

Figure 1: The SP-IdP Proxy Model. Source: GÉANT, GN3PLUS13-642-23

In the proxy model the RI can augment the authentication credential and attributes that come from the conventional R&E IdPs with qualifiers and other tags that are essential for authorisation within that RI: a persistent non-reassigned identifier, an assurance level, community membership roles and groups, and optionally a reputation qualifier. In this way, the RIs shield themselves from the heterogeneity of the global R&E federation space, and they make their operation more efficient since they can now 'hide' all their services behind just a single proxy. Once the boundary SP-IdP proxy joins a federation, and also joins eduGAIN, the RI services see just one IdP, perhaps augmented with an IdP-of-last-resort for their own community, and their users can all securely authenticate at their home institutes and can gather any community attributes they require from the proxy.

Such a proxy does however pose considerable policy challenges: first of all, the services 'internal' to the community see only the single IdP, and they will have to trust it. Relying parties and data centres that service multiple research communities will have to trust several of these proxies, as in EUDAT or in EGI. Basically, each proxy creates a silo, and ways are needed to merge these silos for multi-tenant resource providers. Conversely, the SP-IdP proxy hides all of the research services behind a single SP, so the home organisations and R&E federations will see a single entity, even if the service behind that is provided by many hundreds of different administrative domains and service providers. Without direct knowledge of all these entities, how can the R&E federations be sure what, by accepting the SP-IdP proxy, they allowed into their federation space?

The problem task NA3.4 had to solve was to develop a scalable policy and trust framework that allows the determination of the 'quality' of such SP-IdP proxies. For example, an EGI SP-IdP proxy for all its HTC compute and mass storage services, would be able to express to the R&E federation space that is has an internally-consistent policy set, that it can make assured collective statements about all its 'upstream' services and resource providers, and that it can and will abide by the requirements of the GÉANT Data Protection Code of Conduct, by REFEDS Research & Scholarship, and by REFEDS Sirtfi [ENTCAT]. The new framework should help the management of Research Infrastructures and e-Infrastructures to publish assertions for the proxy on

behalf of the Infrastructure as a whole. An additional problem relates to how one compares the EGI proxy with other infrastructures, e.g. the one from ELIXIR.

We decided that the best approach to solving our problem was to evolve the Security for Collaborating Infrastructures trust framework [SCI] which was also the foundation of the Sirtfi work (see above). The SCI document addresses more than just operational security and incident response by also including traceability, participant responsibilities, community membership documentation requirements, and legal issues such as data protection. SCI is a framework, so it does not in itself set standards, but does establish a set of recommendations by which you can see whether a peer infrastructure has a policy or process in place addressing a topic of concern. SCI was, however, primarily written to compare Distributed IT Infrastructures, and does not deal in detail with the IdP and identity management. We concluded that we had to build on the SCI foundations and to develop a new trust and policy framework that would work for the SP-IdP proxy federation model. This would need to help both on the R&E federation side ("do I want this RI and proxy in my federation?"), and on the SP side where they would find it easier to join more federations, would find IdPs that were more willing to release attributes to them, and for the back-end e-Infrastructure services ("yes, I want assertions from this proxy, since I know it meets my resource provider trust requirements").

AARC NA3.4 staff working in scalable policy negotiation have developed a policy document. We have presented the work at every opportunity and taken advice and input from many different stakeholders, including FIM4R, IGTF, REFEDS, and the Vienna TIIME meeting [TIIME].

We are proposing “Scalable Negotiator for a Community Trust framework in Federated Infrastructures (Snctfi)" as a policy and trust framework allowing determination of the 'quality' of SP-IdP proxies and the RI/EI services they support. The framework places requirements on compliant RIs/EIs for an internally consistent policy set covering critical areas of best practice such as the protection of personal data and security incident handling capabilities. Compliant RIs/EIs are encouraged to assert relevant entity categories and assurance attributes to indicate to federations and their IdPs that they can be trusted to act appropriately. The assertion of the SP-IdP proxy membership of the entity categories Research and Scholarship, GÉANT Data protection code of conduct and Sirtfi allows the Infrastructure to request attributes from eduGAIN IdPs without each and every one of the Infrastructure constituents having to join a Federation/eduGAIN and in doing so, being required to assert its own compliance. Furthermore, by addressing the structure of the security policy set that binds services supported by the SP-IdP proxy, Snctfi facilitates comparison between RIs/EIs.

The full text of the Snctfi document is presented in Appendix A.

4 Other policy work in NA3.4

In this section, we present additional policy work undertaken by task NA3.4. Firstly, this relates to guidance provided to HNSciCloud and secondly, work on proxying federated identity through RCauth.

4.1 Policy challenges for HNSciCloud

HNSciCloud is a pre-commercial procurement project [HNSCICLD] intended to enable the incorporation of commercial cloud service providers into publicly funded e-Infrastructures. The AARC Project was approached by HNSciCloud for guidance on Federated Identity Management. The use of SAML for authentication was specified as a functional requirement in the tender specification, with integration in eduGAIN and ELIXIR AAI being later identified as necessary to interoperate with the procurers' existing Identity Providers.

Appendix B contains a checklist of questions for a new infrastructure or platform, such as that being defined in HNSciCloud, to support and guide the policy and architectural decisions. This list was also circulated within the REFEDS community for their use.

4.2 Proxying federated identity towards the Infrastructures through RCauth

The AARC Blueprint Architecture identifies both the IdP-SP proxy as well as token translation as technical mechanisms to support federated authentication by the Infrastructures. With Snctfi providing one 'direction' of the bridge, there is a need for a complementary policy framework to convey trust from the federated identity ecosystem towards the Infrastructures. But whereas the majority of the currently-existing Infrastructures are already coherent and likely to meet the Snctfi requirements, the reverse is fraught with all the complexities and heterogeneity that has been identified in the work on assurance profiles (baseline and differentiated assurance), incident response (Sirtfi), and (lack of) attribute release (adoption of REFEDS R&S).

Fortunately, the Infrastructures are conceptually simpler, as they share a common and narrower scope (research collaboration, collective service provisioning). This coherency, achieved over the past decade through initiatives such as the Interoperable Global Trust Federation (IGTF) for authentication, and FIM4R for joint requirements on the broader AA domain, provides a solid target to which a policy bridge can be built. Providing the policy for this bridging element in a scalable way, dealing with policy, technical, and governance heterogeneity on the federation side, has been the aim of the harmonization work for what has been labelled the "RCauth.eu" service [RCAUTH].

The RCauth.eu service policy (part of the implementation and sustainability of the CILogon-like TTS for Europe pilot as described by AARC SA1 and in [DNA3.3]) can leverage the existing policy-requirements coherency of the Infrastructures. In constructing the policy for the RCauth.eu proxy elements, the following principles of scalability and trust establishment were followed:

- The service must be accessible to any qualified person in Europe that has a federated identity, or can be provided with an appropriate federated identity by an Infrastructure. *The service should be obtainable for any organization, regardless of the (European) country in which they are based, and regardless of business model considerations by the country, NREN, or (national) bodies where the home organization (or person) is based. If at all possible, it should be accessible to researchers world-wide. It should be possible to join the service as an identity provider independent of whether there is a*

national federation, whether a national federation has joined or not joined an inter-federation scheme (eduGAIN) – within reasonable limits of available effort on behalf of the RCauth.eu service. In particular, Research and e-Infrastructures must be able to join as entities providers of last resort.

- The service must meet both the baseline assurance requirements (see AARC MNA3.1 for details), as well as meet international trust levels for identifier-only assurance (the IGTF DOGWOOD assurance profile [DOGWOOD] that ensure unique non-re-assignable identifiers to users, and be formally accredited. *The formal accreditation of the service means by the relevant IGTF chapter (in this case the EUGridPMA) through a rigorous peer-review process of the policy and implementation practices ensure that the necessary trust anchors (for PKIX) or service endpoints (for OIDC) will be trusted implicitly and automatically by all relying party members of the IGTF. This has a significant advantage in scalability, in that no further negotiations are needed with the e-Infrastructures in Europe (EGI, EUDAT, PRACE), the Americas (XSEDE, Open Science Grid, ComputeCanada, and in the RedCLARA region), and the Asia Pacific (HPCI, PRAGMA).*

- Wherever possible, it leverages eduGAIN and the relevant Entity Categories and assurance profiles. Yet the policy allows for exceptions and explicit configuration where this aim would conflict with the primary aim of ubiquitous availability. *The REFEDS R&S category provides for both a (reasonable representation of) the name of the person, as well as for a persistent, non-reassigned identifier. The identifier is not omnidirectional, but as RCauth.eu can be made into one logical service end-point, a targeted identifier on the federation side results in an omnidirectional identifier on the service side. Adding the Sirtfi incident response, interventional, and traceability requirements on the federation side ensures long-term uniqueness and auditing, and is essential to provide the trust I credential revocation and management towards the 'infrastructure side' of RCauth.eu. In absence of the federation (or IdP) asserting R&S + Sirtfi (e.g. because of technical and local federation policy reasons), a materially equivalent commitment by the IdP may be used in lieu of an entity category tag – thus RCauth.eu can break the 'adoption barrier' and serve its primary aim of ubiquitous service – as long as this commitment is substantiated, scoped, and open to external assessment (e.g. by the RCauth.eu policy management authority).*

- The service must be resilient to incidental compliance failures on the federation (IdP) side, and protect trust in accordance with the Infrastructure requirements, but in a scalable way (i.e. there must be no need for assessment of specific individual transactions). *In any large distributed system incidents (be they technical misconfigurations, policy misunderstandings, or integrity compromises) will occur. And although Sirtfi version 1 addresses the response to such incidents, it does not define a model for identifying incidents, nor does it in itself provide any immediate, scalable way to address incidents. Given the potential number of transactions in the system, detection and mitigation of risks should be largely automatic, and err rather on the side of caution than be too lax. A slightly higher false-positive rate can be accepted as long as the mitigation is automatic and leads to only limited inconvenience to the user (e.g. the user can re-start an operation herself without the need for administrative intervention if thereby the incident is mitigated).*

- The RCauth.eu service operation should not inadvertently take on liability or risks that are beyond its domain of authority or influence. *The service may 'front' the entire global federation world as well as many other entities, yet were it to assume full liability for anything it bridges from that domain no entity would be willing to operate the proxy service. Liability must be deferred back 'upstream' to the*

connected IdPs for any incidents originating there without compromising the trust given to the service itself by the Infrastructures. Since the Infrastructures (being direct stakeholders) are inherently motivated to collaborate, a peer-review, transparent and assessable policy model will suffice to maintain such trust – the more so since most stakeholders are directly represented as relying parties in the IGTF peer review process and assessments.

To demonstrate that a scalable policy model based upon these principles would work in practice, a full policy for RCauth.eu was developed for its proxy model to PKIX credentials. In the PKIX model, the RCauth.eu service becomes a Certification Authority (CA) towards the Infrastructures that act as relying parties (RPs). The identity providers (both federated and directly connected) are Registration Authorities (RAs) and the researchers and people are the Subscribers (and act always on behalf of their professional organisation).

Policies for CAs follow a rigorous standard format, the "Certification Policy" and the "Certification Practice Statement" (CP/CPS) as described in RFC 3647, which ensures that all relevant elements of trust are addressed in the policy. For RCauth.eu, the full policy text (combined CP/CPS) in RFC3647 format amounts to 51 pages, and is for technical trust reasons (the ability to recover from incidents) backed by an air-gapped ('off-line') higher-level CA for which another such policy was developed. Also, the RCauth.eu CP/CPS leverages and includes by reference explicitly the elements that make it into a scalable service: the eduGAIN Declaration (v2.0 at the time of accreditation), the eduGAIN Attribute Profile and its successor specifications, and the REFEDS Research and Scholarship (R&S) and Sirtfi entity category specifications. The combined policy hierarchy (higher-level CA and RCauth.eu Issuing CA) were jointly submitted for assessment and review by the EUGridPMA in February 2016 and went through the peer review process in Spring 2016. The RCauth.eu Pilot ICA G1 and its higher-level CA were formally accredited to the IGTF on May 9th, 2016.

Following a year of pilot operations based on these policy principles, RCauth.eu now services the following identity providers and user groups:

- the eduGAIN entities asserting the REFEDS R&S plus Sirtfi entity categories: 119
- directly connected community (Infrastructure) identity providers: 2 (EGI CheckIn, ELIXIR)
- other directly connected IdPs (materially equivalent IdP without entity category support):1 (KIT, Germany)

Amongst the eduGAIN-originating entities is also the global IGTF eduGAIN proxy bridge, which provides for higher-assurance identity provider of last resort for those in the global research community that already avail over existing research and e-Infrastructure credentials. It also includes the XSEDE identity provider, serving the distributed research community in the USA through a domain-oriented vetting model.

On the 'infrastructure' side of RCauth.eu, the connection model through centralized credential management services (Master Portals, as described by the SA1 pilot) introduces the concept of a central hub for a research community, thereby off-loading the (potentially unbounded) registration of all individual research services off the single RCauth.eu instance.

The Master Portals connect once to the RCauth.eu instance, but then themselves can connect any service within their research infrastructure, domain organization, or e-Infrastructure. Basic security requirements are laid out in the Guidelines on Trusted Credential Stores and the Guidelines on Private Key Protection for End-users, which put conditions on the acceptability of such portals for the Rcauth.eu CA instance. This layered

approach allows security-critical elements to be concentrated in a few, well controlled, locations, whilst simultaneously permitting deployment of many community-run services with limited exposure of user credentials. It provides a hierarchical deployment model for RCauth.eu by the relying parties that scales linearly with the number of research domains instead of the number of research services.

The adoption of RCauth as discussed in the sustainability model study and evidenced by the incorporation of the RCauth.eu model as a service by the Infrastructures, reinforces the conclusion that this policy model, using both the grouping of entities as well as for allowing incidental integration of non-eduGAIN entities in an open way, is effective in providing for the research and collaboration needs.

5 Conclusions and future plans

Creating the common federation 'ecosystem' for research collaboration requires that everyone is effectively connected to it, including not only those organisations whose primary purpose is collaborative research, but also institutions (both identity providers and those providing services) where the number of individuals collaborating is small relative to the size of the organisation. Similarly, there are too many organisations worldwide (eduGAIN alone already has about 4000 entities; there are even more national entities) to make bilateral discussion and agreement a feasible model to negotiate authentication and release of attributes.

In addressing these challenges, AARC has used three complementary approaches: promotion of the use of 'entity grouping' and its associated specifications; the definition of a framework to organise service provider collectives and permit these to express their common qualities in a standard way; and the development of a concrete set of policies and practices for an SP-IdP proxy that enables researchers with an existing federated IdP to connect to existing e-Infrastructures without requiring modification of either.

The entity grouping study leverages the eduGAIN "Entity Category" (SAML-EC) mechanism: both controlled and self-asserted attributes can be associated with IdPs and services in eduGAIN to make them 'recognisable' and to permit filtering. In other policy work, AARC used these entity categories for expressing operational security capabilities ("Sirtfi") or leveraged existing definitions such as the REFEDS "Research and Scholarship" controlled group or the (European) "GÉANT Data Protection Code of Conduct v1". Initial surveying indicated that this grouping concept is applicable to the scalable policy negotiation challenge. The wider applicability of the SAML-EC mechanism already ensured its promotion through the GÉANT project and REFEDS and AARC has been able to leverage this complementary work effectively both by proposing new SAML-ECs (e.g. in the area of Assurance Profiles as discussed in DNA3.1) and by using SAML-ECs in SP-IdP proxies (as described here in RCauth, and as used in the SA1 Pilot on the IGTF Certificate Proxy). Detailed analytics is now provided routinely by GÉANT and eduGAIN (through technical.edugain.org) allowing AARC to track developments in this area. Use of SAML-EC attributes has increased significantly over the past 2 years, and the adoption shows uptake of Sirtfi (137 entities with the SAML-EC "Sirtfi" tag within 1 year) and REFEDS R&S (646 entities out of ~4000) with an increasing overlap: 57 assert both R&S and Sirtfi as of April 2017, a combination that indicates the IdP meets the baseline requirements for Infrastructure access. GÉANT and eduGAIN are now routinely tracking the statistics of the deployment of entity categories, e.g. showing that in April 2017, 57 entities assert the combination meeting the baseline requirements for infrastructures access (i.e. R&S and Sirtfi together).

The entity category attribute mechanism will be pursued also in the future for expressing qualifiers, although depending on technology evolution and practices their specific format (SAML meta-data attributes) may well change. The adoption rate is increasing, and although specific per-entity 'bypasses' will always be needed to support specific collaborators or use cases, the incorporation of the mechanism in GÉANT production is an assurance of its long-term future.

The ability to connect Infrastructures (in the R&E federation world typically seen as 'collectives of service providers working together') has been structured through the introduction of SP-IdP proxies (AAI gateways) in the AARC Blueprint Architecture. This, on the one hand simplifies the technical connection process in the R&E federation space, since it exposes a single entity (the proxy) towards the federations as a service provider. Yet, unless a coherent policy statement can be made about the 'qualities' within this collective, the negotiation challenge remains, with some IdPs still requiring bilateral contracts. This is particularly true in the case of attribute release, where some of the R&S attributes like the real name of the person and email address (although every user will happily send email anyway) are seen by IdPs as subject to data protection regulations, and their GDPR compliance requires specific assurances regarding onward processing of the attributes by the proxy. The "Scalable Negotiator for a Community Trust framework in Federated Infrastructures" or "Snctfi" framework introduced here defines a mechanism for Infrastructures to express policy alignment within the Infrastructure. Based on the SCI framework, it takes specifically those elements related to infrastructure policy coherency and uses them to allow the proxy to assert (SAML) Entity Category attributes from the REFEDS controlled vocabulary (specifically R&S) and also lays the basis for asserting alignment with the GÉANT Data Protection Code of Conduct v1 (DPCoCo) and Sirtfi. The framework leverages complementary work in operational security (it includes Sirtfi by reference) and accounting data protection (as described for example in AARC deliverable DNA3.5).

Snctfi, however, is conceptually different from other specifications, such as the REFEDS R&S and the GÉANT DP CoCo, in that it provides a framework for the Infrastructures (not the federations) to ensure internal consistency. As such, the details of Snctfi are not necessarily disclosed to the R&E federations, making both the consultation and the document management process different. Although Snctfi has been discussed with the REFEDS community, its primary stakeholders are the generic e-Infrastructures (e.g. EGI, PRACE, and EUDAT in Europe, and XSEDE, Open Science Grid, HPCI in the rest of the world), and the Research Infrastructures (e.g. the ESFRI clusters in Europe, the WLCG collaboration, etc.). These Infrastructures are explicitly global in scope, and policy alignment mechanisms are needed that work also outside of the EU and its list of countries with equivalent protection for personal data. The future challenges for Snctfi are thus twofold: work – within the context of the 'new' General Data Protection Regulation – on incorporating a mechanism that allows Infrastructures to collectively express alignment with this process, in a way that allows the proxy to do that on behalf of all back-end services and also to do this for Infrastructures and collaborations that may not in themselves be a legal entity. Secondly, to be successful in its aims the framework needs wide adoption within both the generic e-Infrastructures and the Research Infrastructures. To this end, international groups such as the Interoperable Global Trust Federation (IGTF), the WISE Information Security for collaborating e-Infrastructures community, and the Research Infrastructure requirements alignment group on federated identity management (FIM4R) appear to be the most appropriate forums. In the AARC2 project, these groups have been identified as the partners for future development of the Snctfi framework.

To investigate the 'end-to-end' scalability of the proxy model for policy harmonisation, the AARC SA1 technical pilot on a CILogon-like TTS for Europe was complemented by a complete suite of policies and operational practices ("RCauth.eu") that allowed the pilot to be trusted in the production environments of e-Infrastructures

globally through the IGTF accreditation process. The accreditation process includes a thorough external review by peer infrastructures and identity service providers, addressing technical, policy, security, and business model issues. The developed policy set intentionally leveraged all existing federation mechanisms, including the REFEDS R&S entity category, the Sirtfi security trust framework, and guidance on credential management and trusted credential stores developed in collaboration with the Infrastructures in the IGTF. The RCauth IGTF accreditation in May 2016, closely following the technical implementation of the AARC SA1 pilot, demonstrates the feasibility of the scalable policy methodology through SAML-ECs and Sirtfi. The 'reverse' proxy bridge, another SA1 Pilot deploying an "IGTF Certificate Proxy" identity provider in eduGAIN, demonstrates that the accreditation processes used within the Infrastructures are aligned with the eduGAIN entity category model by its ability to assert both REFEDS R&S as well as Sirtfi.

Both RCauth.eu and the IGTF Certificate Proxy have a long-term sustainability plan through their adoption by the major e-Infrastructures in Europe and by national research and educational infrastructures – of which the details are given in the AARC deliverable DNA3.3 on sustainability and in the specific "Sustainability models for the AARC CILogon-like TTS Pilot and RCauth.eu".

The necessary evolution of Snctfi and the close involvement of the Infrastructures and R&E federations in the alignment of the entity categories and policy frameworks will be taken up in the successor AARC2 project through its Competence Centre.

Appendix A The “Snctfi” document

The Scalable Negotiator for a Community Trust Framework in Federated Infrastructures (Snctfi) paper identifies operational and policy requirements to help establish trust between an Infrastructure and identity providers either in an R&E Federation or in another Infrastructure, in each case joined via a Service Provider to Identity Provider proxy. It is intended for use by the personnel responsible for the management, operation and security of an Infrastructure and those wishing to assess its trustworthiness, and reproduced herein as one of the results of the AARC task on scalable policy negotiation mechanisms.

Scalable Negotiator for a Community Trust Framework in Federated Infrastructures (Snctfi)

Licia Florio (GÉANT), David Groep (Nikhef), Christos Kanellopoulos (GÉANT), David Kelsey (STFC), Mikael Linden (CSC), Ian Neilson (STFC), Stefan Paetow (Jisc), Wolfgang Pempe (DFN), Vincent Ribaillier (IDRIS-CNRS), Mischa Salle (Nikhef), Hannah Short (CERN), Uros Stevanovic (KIT) and Gerben Venekamp (SURFsara)

AARC - Version 1.0 - 26 Apr 2017

e-mail: david.kelsey@stfc.ac.uk

Abstract: This paper identifies operational and policy requirements to help establish trust between an Infrastructure and identity providers either in an R&E Federation or in another Infrastructure, in each case joined via a Service Provider to Identity Provider proxy.

Audience: This document is intended for use by the personnel responsible for the management, operation and security of an Infrastructure and those wishing to assess its trustworthiness.

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

1. Background

Research Infrastructures (RI) and e-Infrastructures (EI) increasingly make use of national and global "Research and Education" (R&E) identity federations to facilitate their users' access to RI/EI services. When requesting access to RI/EI services, users are directed to authenticate at their home organisation Identity Provider (IdP) using their home organisation credentials. The RI/EI may enrich the resulting authentication credential with community information to allow authorisation decisions to be made on the combined assertions. For example, information about community roles, added to the token, may be mapped to rights and privileges.

Studies in the AARC project [1] have shown that research communities often connect to a R&E federation using a Service Provider to Identity Provider proxy (SP-IdP proxy). In this model a single component, the SP-IdP proxy, also known as an AAI-Gateway, negotiates between the services in the RI/EI and the IdPs in the federation as shown in Figure 1. By positioning all RI/EI services behind a single proxy IdP, the RI/EI is shielded from the heterogeneity of the global R&E federations and itself need only be registered once, for all its services, as a single SP in the R&E federations. More details are presented in the AARC Blueprint Architecture [2] which identifies this model as a recommendation for research

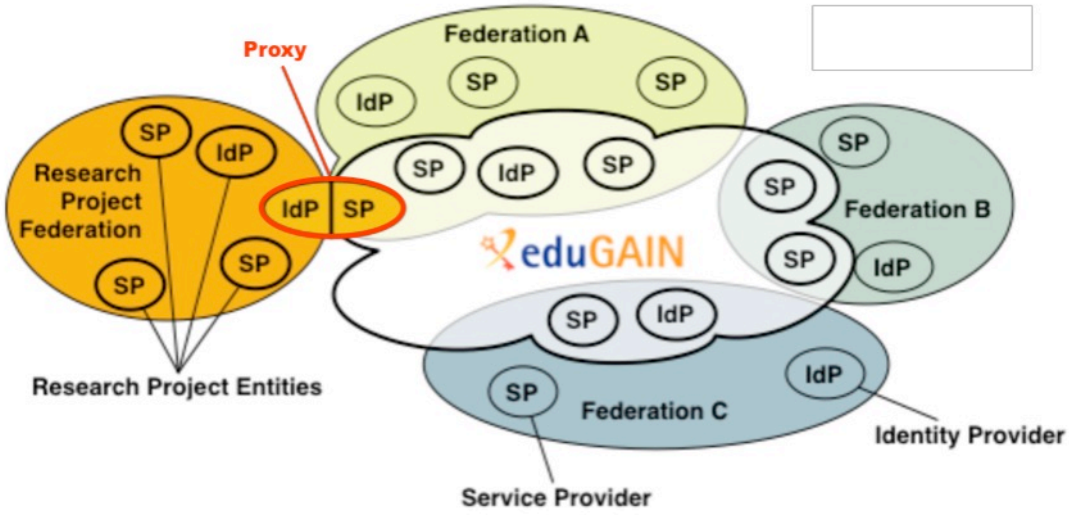

Figure 1: The SP-IdP Proxy Model. Source: GÉANT, GN3PLUS13-642-23

collaborations engaging with R&E federations.

The use of the proxy model, however, poses policy challenges in establishing a sufficient level of trust between the RI/EI SPs and the federation IdPs, as the IdPs must be assured that the identity information they release will be treated appropriately by the RI/EI and its SPs.

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

Depending on the number of communities supported, some RI/EI services may only see a single proxy IdP which they have to trust, but more generic service providers or EIs, supporting multiple research communities or RIs, may have to trust many proxy IdPs or many RIs. R&E federations can feel confident in releasing attributes to SP-IdP proxies that assert entity attributes related to, for example, REFEDS Research and Scholarship, data protection and/or security incident response [3]. The framework introduced in this paper enables the management of Research Infrastructures and e-Infrastructures to publish such assertions for the proxy on behalf of the Infrastructure as a whole.

2. Introduction to the Snctfi Trust Framework

This document addresses the problem of establishing the transitive trust described above. Building on the "Security for Collaboration among Infrastructures (SCI)" framework [4], it proposes the "Scalable Negotiator for a Community Trust framework in Federated Infrastructures" (Snctfi) as a policy and trust framework allowing determination of the 'quality' of SP-IdP proxies and the RI/EI services they support. This framework places requirements on compliant RIs/EIs for an internally consistent policy set covering critical areas of best practice such as the protection of personal data and security incident handling capabilities. Compliant RIs/EIs are encouraged to assert relevant entity categories and assurance attributes [3] to assure federations and their IdPs that they can be trusted to act appropriately. The assertion by the SP-IdP proxy of additional qualifiers or tags, for example REFEDS Research and Scholarship, GÉANT data protection code of conduct and REFEDS Sirtfi, encourages the release of attributes from eduGAIN IdPs to the Infrastructure. The benefit of this approach is that each of the Infrastructure constituents no longer has to join an R&E federation and eduGAIN in order to assert its own compliance. Furthermore, by addressing the structure of the security policy set that binds services supported by the SP-IdP proxy, Snctfi facilitates comparison between RIs/EIs.

3. Scope

This document applies to the set of SPs, group- and VO-management systems acting as Attribute Authorities, and the SP-IdP proxy, together comprising an e-Infrastructure or Research Infrastructure (hereafter called the "*Infrastructure*"). The individual SPs, Attribute Authorities and SP-IdP proxies are hereafter called the "*Constituents*" of the "*Infrastructure*".

4. Normative Requirements

We present normative requirements in this document in three areas: Operational Security, User Responsibilities and the Protection and Processing of Personal Data.

An *Infrastructure* must address these requirements if asserting conformance with the Snctfi Trust Framework.

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

4.1 Operational Security [OS]

The aims of Operational Security in an *Infrastructure* include:

- Preventing security incidents, wherever possible, via the timely handling of and patching of software vulnerabilities;
- Minimising the impact of those security incidents that do occur by implementing appropriate logging, monitoring and incident handling capabilities sufficient to understand the causes and the controls necessary to contain the impact and prevent recurrence.

The *Infrastructure* must:

[OS1] define a set of common security requirements including stipulations on: authentication, authorisation, access control, physical and network security, security vulnerability handling and security incident handling, together with compliance mechanisms ensuring appropriate implementations.

[OS2] ensure that its *Constituents* abide by the stipulations of the *Infrastructure* security requirements by means of, for example, binding contracts, MoUs, SLAs, OLAs, policies, or a suitable combination of these.

[OS3] ensure that its *Constituents* meet all relevant requirements specified in REFEDS Sirtfi version 1.0 [5].

[OS4] define appropriate policies and procedures necessary to coordinate the implementation of [OS2] and [OS3] commensurate with the scale of the *Infrastructure*.

4.2 User Responsibilities [UR, RU, RC]

To establish trust between the *Infrastructure* and the R&E federations, and between *Infrastructures*, the *Infrastructure* relies on appropriate behaviour by its users and user communities.

[UR1] The *Infrastructure* must ensure that its users and user communities are aware that they have the responsibilities documented in this sub-section.

4.2.1 Individual Users [RU]

Each SP or the *Infrastructure* must provide:

[RU1] an Acceptable Use Policy (AUP). The AUP must at least address the following areas: defined acceptable use, non-acceptable use, user registration, protection and use of credentials, data protection and privacy.

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

[RU2] a process to ensure that all users are aware of, and accept the requirement to abide by, the AUP.

[RU3] communication to their users of any changes to the AUP and/or additional restrictions or requirements on acceptable use that arise out of new collaborative partnerships (if any).

4.2.2 Collections of Users [RC]

A Collection of users is a group of individuals, organised with a common purpose, jointly granted access to the *Infrastructure*. It may serve as an entity which acts as the interface between the individual users and the *Infrastructure*. In general, the members of the Collection will not need to separately negotiate access with Service Providers or *Infrastructures*.

Examples of Collections of users include, but are not limited to: User groups, Virtual Organisations, Research Communities, Research Infrastructures, Virtual Research Communities, Projects, Communities authorised to use particular portals or gateways, and geographically organised communities.

Each *Infrastructure* must have:

[RC1] policies and procedures regulating the behaviour of the management of the Collection of users in relation to individual user registration and membership management (registration, renewal, suspensions, removal, and banning). At a minimum, these must address the accuracy of individual user contact information both for initial collection and periodic renewal and related Data Protection issues (see later).

[RC2] a process to ensure that all Collections of users using the *Infrastructure* are aware of, and accept the need to abide by, applicable *Infrastructure* policy requirements.

The *Infrastructure* policies must require that Collections of users must:

[RC3] be aware that inappropriate actions by individual members of the Collection may adversely affect the ability of other members to use the *Infrastructure*.

[RC4] ensure there is a way of identifying the individual responsible for an action.

[RC5] record membership management actions as these may be needed in security incident response.

[RC6] define their common aims and purposes, i.e. the research or scholarship goals of the group. They should make this available to the *Infrastructure* and/or Service Providers to allow them to make decisions on resource allocation.

[RC7] inform the *Infrastructure* of any significant changes to common aim and purposes (see above).

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

4.3 Protection and processing of Personal Data [DP]

Infrastructure Constituents and, where necessary Collections of users, must have policies and procedures addressing the protection of the privacy of individual users, i.e. members of the Collections, with regard to the processing of their personal data (also known as Personally Identifiable Information or PII) collected as a result of their access to services provided by the *Infrastructure*.

The *Infrastructure* must:

[DP1] have a Data Protection Policy binding those *Constituents* and Collections of Users who process personal data to an appropriate policy framework, e.g. the GÉANT Data Protection Code of Conduct [6] or, for example, as recommended by AARC [7].

[DP2] ensure that all *Constituents* must provide, in a visible and accessible way, a Privacy Policy covering their processing of personal data for purposes that are necessary for the safe and reliable operation of their service, compliant with the *Infrastructure* policy (or policy framework). The availability of a Privacy Policy template for the *Constituents* to follow, provided by the *Infrastructure*, would help the easier production of such a policy.

5. References

[1] https://aarc-project.eu/wp-content/uploads/2015/10/AARC-DJRA1.1.pdf

[2] https://aarc-project.eu/roadmap/blueprint-architecture

[3] https://refeds.org/specifications; https://wiki.refeds.org/display/CODE/Data+Protection+Code+of+Conduct+Home

[4] https://pos.sissa.it/archive/conferences/179/011/ISGC%202013_011.pdf

[5] https://refeds.org/sirtfi

[6] https://www.geant.org/uri/Pages/dataprotection-code-of-conduct.aspx

[7] https://aarc-project.eu/wp-content/uploads/2016/12/AARC-DNA3.5_Recommendations-for-Processing-Personal-Data_2016_11_07_v4_DG.pdf

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965 (AARC).

Appendix B Considerations for designing an Infrastructure

The following questionnaire prompts infrastructure management to articulate requirements for policy and architecture when designing a federated identity management solution.

This effort began as input to the HNSciCloud Project [http://www.hnscicloud.eu] and reflects input from members of the REFEDS community including: Niels van Dijk (SURFnet), Mikael Linden (CSC), Mark Jones (UTHealth), Michal Prochazka (CESNET), Lukas Haemmerle (SWITCH), Hannah Short (CERN), Romain Wartel (CERN), Wolfgang Pempe (DFN).

User Groups

- Do all R&E users belong to an eduGAIN enabled home organisation?
- Are there any users outside the R&E Community? E.g. citizen scientists.

User Data

- Which attributes will be required from home organisations?
- Which additional attributes, if any, will be required from the users directly?
- Which additional attributes, if any, will be required by the infrastructure/community? (e.g. group memberships or entitlements, cf. also section “Authorisation” below)
- What user data will be stored centrally within the infrastructure?
- For how long will user data be stored?
- How will a user be able to request a change to or deletion of their user data?
- How will an individual be identified?
- How will multiple accounts belonging to an individual, either consecutively or in parallel, be associated?
- How will variation in user data from multiple home organisations be resolved?
- How will assurance of the accuracy of user attributes be established, for the authentication token overall and on a per-attribute basis?
- Will each identity provider populate a global, non-reassignable unique ID?
- Will all user data remain within the organisational domain of the collaborative organisation? (or are external services (PAAS/IAAS) used?)

Participant Policies

- Will a common policy set be defined for all participating organisations; e.g. security response, data protection?
- Is there an acceptable use policy for end users?
- Which takes precedence if a participant policy for this infrastructure contradicts a local policy of a home organisation?
- How will participating organisations assert compliance with relevant policies?
- How will end users assert compliance with relevant policies?

Security

- How will the risk profile of the services accessed via the infrastructure, including external services, be established?
- Is the security capability of each participating organisation sufficient to cover the risk profile of the services?
- Is the security capability of each participating organisation known and documented?
- Is adequate proactive and reactive software defence employed at each participant organisation? How is this defined?
- How will traceability of user actions be established?
- Will shared accounts be granted access to the infrastructure?
- Can we ensure that participants proactively contribute to incident response?
- Are participating organisations Sirtfi compliant? https://refeds.org/sirtfi

Trust

- Do we have sufficient assurance that the identity belongs to the asserted person?
- How will participating organisations' membership be established?
- Are face-to-face passport checks or other forms of identity proof required?
- Can existing identity proofing at participating organisations be reused?
- Does identity vetting need to be controlled centrally, by the infrastructure?
- Do all participating organisations agree to abide by confidentiality protocols during incident response and general communication?
- Is there an established network of trust groups to provide coverage of all the participants?
- Are there individuals identified at participating organisations to assist with trust and security?

Authorisation

- Will an authorisation system be used to define user roles within the infrastructure? Would such an authorisation system be externalised?
- Will there be any automatic mapping between attributes and user roles?
- Will a blacklist of authorisation be implemented?
- How will membership be expressed?
- Are multiple authorisation roles needed?
- Is there a relation between authentication assurance and authorisation roles? How is that expressed?

References

[BPA]	https://aarc-project.eu/blueprint-architecture/
[DJRA1.1]	https://aarc-project.eu/wp-content/uploads/2015/10/AARC-DJRA1.1.pdf
[DNA3.2]	https://aarc-project.eu/wp-content/uploads/2017/02/DNA3.2-Security-Incident-Response-Procedure-v1.0.pdf
[DNA3.3]	https://aarc-project.eu/wp-content/uploads/2017/04/DNA3-3-final.pdf
[DOGWOOD]	https://www.igtf.net/ap/loa/
[DPCOCO]	https://wiki.refeds.org/display/CODE/Data+Protection+Code+of+Conduct+Home
[EC CAND]	https://wiki.refeds.org/display/ENT/Academic-Academia; https://wiki.refeds.org/display/ENT/Library-Affiliation
[ENTCAT]	https://wiki.refeds.org/display/ENT/Entity-Categories+Home
[FIM4R]	D. Broeder et al., Federated Identity Management for Research Collaborations; CERN-OPEN- 2012-006; http://cds.cern.ch/record/1442597/files/CERN-OPEN-2012-006.pdf
[HNSCICLD]	http://www.hnscicloud.eu
[IGTF]	https://www.igtf.net/
[R&S]	https://refeds.org/category/research-and-scholarship
[RCAUTH]	http://rcauth.eu/
[SAML2]	http://docs.oasis-open.org/security/saml/v2.0/saml-metadata-2.0-os.pdf
[SCI]	https://pos.sissa.it/archive/conferences/179/011/ISGC%202013_011.pdf
[SIRTFI]	https://refeds.org/sirtfi
[SRTFI WG]	https://wiki.refeds.org/display/GROUPS/SIRTFI
[TIIME]	https://tiimeworkshop.eu/tiime2017/index.html
[WISE]	https://wise-community.org/