

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

Version 1.0-2017

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.

Authors

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 Ribailier (IDRIS-CNRS), Mischa Salle (Nikhef), Hannah Short (CERN), Uros Stevanovic (KIT) and Gerben Venekamp (SURFsara)

e-mail: snctfi@igtf.net

Table of Contents

1	Background.....	3
2	Introduction to the Snctfi Trust Framework.....	4
3	Scope.....	4
4	Normative Requirements.....	5
4.1	Operational Security [OS].....	5
4.2	User Responsibilities [EU, RU, RC].....	6
4.2.1	Individual Users [RU].....	6
4.2.2	Collections of Users [RC].....	6
4.3	Protection and processing of Personal Data [DP].....	7
5	References.....	7

Identifications

This document: **urn:oid:1.2.840.113612.5.2.8.1**
DOI: **10.5281/zenodo.1173109**

License

This document is licensed under [CC BY-NC-SA 4.0](https://creativecommons.org/licenses/by-nc-sa/4.0/)

This work, "Scalable Negotiator for a Community Trust Framework in Federated Infrastructures", is a derivative of "[A Trust Framework for Security Collaboration among Infrastructures](#)" by D. Kelsey, K. Chadwick, I. Gaines, D. Groep, U. Kaila, C. Kanellopoulos, J. Marsteller, R. Niederberger, V. Ribaillier, R. Wartel, W. Weisz and J. Wolfrat, used under [CC BY-NC-SA 4.0](#).

Acknowledgements

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

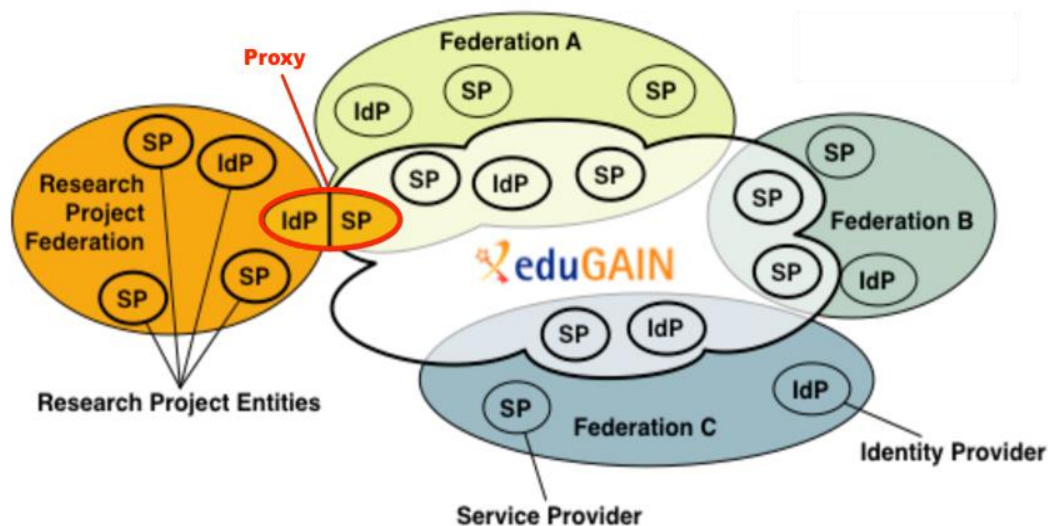


Figure 1: The SP-IdP Proxy Model. Source: GÉANT, GN3PLUS13-642-23

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 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. 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.

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 [EU, 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.

[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).

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/sirffi>
- [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