



30-Nov-2015

Milestone MNA3.1: Recommendations on Minimal Assurance Level Relevant for Low-risk Research Use Cases

Milestone MNA3.1

Contractual Date: 30-Nov-2015
Actual Date: 30-Nov-2015
Grant Agreement No.: 653965
Work Package: NA3 « Policy and Best Practices Harmonisation »
Task Item: Task 1 Development of best practices for Levels of Assurance
Lead Partner: CSC – IT Center for Science Ltd
Document Code: MNA3.1
Authors: Mikael Linden, David Groep, Daniela Pöhn, Tangui Coulouarn, Wolfgang Pempe, Hannah Short

© GÉANT on behalf of the AARC project.

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

Abstract

In a federated AAI, the user's Home Organisation issuing and managing user's credentials determines the assurance level available for the user identity. For the risk management of the research services relying on the federated AAI, it is important to determine the assurance level available for the authenticated users. This document defines requirements for a minimal assurance level which is still relevant for low-risk research use cases.



Table of Contents

Executive Summary	1
1 Introduction	2
1.1 Scope	2
1.2 Method	2
1.2.1 Infrastructures interviewed	2
1.2.2 General findings	3
1.3 Related work	3
2 Minimal assurance profile	3
2.1 Requirements for the minimal assurance profile	4
2.2 Implementation note	5
3 Conclusions	6
Appendix A A proposed tool to support Home Organisations' self-assessments	7
References	9
Glossary	10



Executive Summary

In identity management, assurance level means a level of confidence in the binding between an entity (such as a user) and the presented identity information. In federated identity management, the assurance level is determined by the user's Home Organisation who issues and manages the user identities (accounts) and carries out the user authentication. In the context of research and education identity federations, the Home Organisation is typically the university or research institution to which the user is affiliated (e.g. the employer of a researcher).

Currently, in research and education, there is no well-established assurance level framework. The assurance levels available depend on the policies and practices of the user's Home Organisation and the identity federation to which it belongs. Research infrastructures have seen the need for different Levels of Assurance (LoA) with provenance [FIM4R]. This document defines a "LoA floor" – a minimal assurance level which is still relevant for low-risk research use cases.



1 Introduction

1.1 Scope

In identity management, assurance level means a level of confidence in the binding between an entity and the presented identity information [ITU-T X.1252].

This document defines requirements for a minimal assurance level which is still relevant for low-risk research use cases.

The document focuses on a federated AAI (Authentication and Authorisation Infrastructure) use scenario where the end user's Home Organisation (e.g., the university that is the employer of a researcher) issues him/her an identity and authenticates him/her for the services provided by a research infrastructure or e-infrastructure.

1.2 Method

This recommendation is based on interviewing e-infrastructures and research infrastructures from different disciplines. This section presents the way the interview was done, the infrastructures that were interviewed and provides some general findings from the interviews.

1.2.1 Infrastructures interviewed

To develop the minimal assurance level, a structured interview was carried out to six research infrastructures and two e-infrastructures. The structured interview was chosen (instead of a form based query) because it provided the interviewer with an opportunity to acquire a deeper understanding of the infrastructure's underlying needs and use scenarios.

The infrastructures interviewed were identified from the AARC project participants and from the FIM4R community.

The following research infrastructures were interviewed for this document:

- CLARIN (language research)
- DARIAH (arts and humanities)
- ELIXIR (life science)
- LIGO (physics)
- photon/neutron facilities (multi-disciplinary)

- WLCG (physics)

Following e-infrastructures were interviewed:

- EGI
- PRACE

1.2.2 General findings

In general, the results provided by research infrastructures were more specific than the results from the e-infrastructures. The e-infrastructures themselves are service platforms for different research infrastructures and communities and need to adapt to their LoA needs.

In general, some of the infrastructures and facilities rely mostly on their own AAI that is used for management of users' roles, attributes and authorisations. Those infrastructures use the federated AAI primarily for authenticating the user. Because they have compensating controls in place, their LoA needs are limited. Examples of such infrastructures are LIGO and photon/neutron facilities. Also the X.509 technology based infrastructures that make use of the VOMS system (EGI and WLCG) share some of these properties.

In general, the research infrastructures who need a higher LoA need it for either of the following reasons:

1. They have sensitive research data. The research infrastructure has an obligation to protect the confidentiality of the research data they host. Example: ELIXIR (sensitive human data)
2. They have expensive research instruments. Some research infrastructures have highly expensive research instruments and need to protect their integrity and availability against the risk of downtime and unauthorised access. Example: LIGO and WLCG.

The interview questions and results are summarised in a separate document [SurveyRes].

1.3 Related work

In parallel, the GEANT GN4-1 project has interviewed identity federations and Home Organisations who operate an Identity Provider server regarding their current LoA offering and their possibilities to improve it. The results are published in a white paper "Service aspects of assurance" [GN4-1SA5.1.4].

2 Minimal assurance profile

Based on the interviews, this section introduces the proposed minimal assurance profile. Finally, some remarks on how to implement the profile on the technical infrastructure are proposed.

2.1 Requirements for the minimal assurance profile

It is not expected that a Home Organisation must comply with these requirements for all of its user accounts. Instead, the Home Organisation must be able to tag the compliant accounts and logins (see section 2.2 “Implementation note”).

1. The accounts in the Home Organisations must each belong to a known individual person

Accounts must not be shared and the Home Organisation must be able to trace each account back to its holder.

This requirement follows from the need for a reliable audit trail. For instance, in the Service Provider side, the user may need to register to services and commit to their licence terms, and if the Service Provider does not know who exactly was the individual user who logged in, those service terms become unenforceable.

2. Persistent user identifiers (i.e., no re-assignment of user identifiers)

The Home Organisation must provide a persistent identifier for a user. The identifier must have the property that it is never re-assigned, i.e. recycled to another person. This is due to Service Providers assigning sensitive files to the user identifier and the files can be there even for a long time.

Currently, in the federated AAI, the most widely used identifier `eduPersonPrincipalName` is lacking this property.

3. Documented identity vetting procedures (not necessarily face-to-face)

The Home Organisations must have a documented identity vetting process for its user accounts. The documentation must be available in English and follow a widely established structure.

4. Password authentication (with some good practices)

For the low-risk research, authentication with passwords is sufficient. However, there must be certain widely approved good practices for the password quality, such as length, complexity, and change cycle.

5. Departing user’s `eduPersonAffiliation` must change promptly

When a user departs from his/her Home Organisation, his/her `eduPersonAffiliation` attribute value (and derivatives, such as `eduPersonScopedAffiliation`) must reflect the change promptly, within one month of the departure at most. This must cover at least the `eduPersonAffiliation`=“faculty”, “student”, and “member” values, which are found to be consistently interpreted in different federations [REFEDS ePSA]). This profile does not introduce requirements for other `eduPersonAffiliation` values (such as, “affiliate”, “alum” or “library-walk-in”).

6. Self-assessment (supported with specific guidelines)

A regular self-assessment is deemed as a sufficient way for having a cost-effective audit of the Home Organisation’s Identity management practices. However, there must be a self-assessment framework that is

complete and specific enough. See the proposal in Appendix A regarding a tool that helps the Home Organisations in doing the self-assessments.

2.2 Implementation note

How the requirements presented above are implemented in the technical infrastructure (e.g., the SAML-based federated AAI, such as, eduGAIN) is out of scope for this document. However, the minimum LoA level could be achieved, for instance, with the following two-layered approach:

- The Home Organisation indicates that it conforms to this recommendation, at least for some subgroup of its user identities. This enables the Service Providers to blacklist those Home Organisations who do not conform to the recommendation. For SAML Identity providers, it could be, for instance, an Entity Category attribute in the Identity Provider's SAML metadata.
- When a user is authenticated, the Home Organisation releases an indication to the Service Provider if the authenticated user complies with this recommendation. It is then up to the Service Provider to decide if it denies access for the user. For SAML Identity Providers, the indication could be for instance a SAML attribute (e.g., eduPersonAssurance) or an authentication context assertion.

3 Conclusions

Based on interviews made on eight research and e-infrastructures, this document presented six requirements for a minimal assurance level which is sufficient for low-risk research use cases. The requirements covered properties of the user identities (accounts must belong to known individuals) and identifiers (identifiers must be persistent), initial proof of identity (identity vetting procedures must be documented), authentication (passwords with sufficient quality), freshness of user data (eduPersonAffiliation attribute values must reflect the departure of a user) and audits (a self-assessment based on a complete and specific framework).

This document proposes also designing, developing and rolling out a tool to assist self-assessment of a Home Organisation's Identity management practices. The tool could be useful also in other context.

Appendix A **A proposed tool to support Home Organisations' self-assessments**

One of the pain points for deploying a LoA framework is its actual uptake in the Home Organisations. Without paying extra attention to support the Home Organisations in deploying the LoA, there is a risk that the LoA framework does not get adopted. Because the identity federation operators are in a direct relationship with their federations' Home Organisations, the operators are likely to have a central role in the communication and outreach. Providing a central tool to support Home Organisations would be a cost-effective way to ease the operators' burden.

It is therefore suggested to design, deploy, and roll out a tool to support the Home Organisation in doing the LoA self-audits. It is further suggested that the tool is operated centrally for the whole federated AAI (eduGAIN) community.

The tool is proposed to have the following main functionalities:

- The tool is an eduGAIN Service Provider to which any eduGAIN Identity Provider admin can log in
 - Identification/authorisation of the Identity Provider admins could be done by picking the contact information from the Identity Provider's SAML metadata and sending a log-in link to that email address
- The tool presents structured self-assessment questions to the Identity Provider/Identity management admin
 - Quantitative ("do accounts belong to an individual? [Yes/No]")
 - Qualitative ("explain how you ensure accounts belong to an individual").
 - The language would be English (only)
- The tool publishes the results for anyone to read
 - In a well-known URL, for instance <http://assurance.edugain.org/<idp-entity-id>>
- The tool would evaluate if the LoA minimum is fulfilled
 - And show the Home Organisation's IdM admin where the pain points are and why the LoA minimum is not met
- If an Identity Provider qualifies the LoA minimum, the tool would output an Entity Category attribute for the Identity Provider, possibly directly to eduGAIN metadata (MDS)
 - An Entity Category attribute needed to be defined for that
- The tool would ask the Identity Provider admin to do a new self-assessment every year
 - A missing self-assessment would make the Entity Category attribute disappear from eduGAIN metadata
- The tool could assist in a LoA peer review
 - If peer review becomes a requirement, for instance, for a higher LoA level

There is interest in the tool in other community efforts as well, for instance, in the Sirtfi working group where the tool could assist the Home Organisations to self-assess if they qualify to the Sirtfi framework. Therefore, it is proposed that

- the tool and the content (i.e. the self-assessment question) are kept separate so that the tool can be used in different contexts
- there is a flexible option to nominate organisations to complete peer-reviews
- the project to design, develop, and deploy the tool are detached from AARC NA3.1 as a separate effort common to AARC NA3.1, AARC NA3.2, and GN4SA5.1.4.

References

- [FIM4R]** Federated Identity Management for Research Collaborations.
<https://cds.cern.ch/record/1442597/files/CERN-OPEN-2012-006.pdf>
- [GN4-1SA5.1.4]** <https://wiki.geant.org/display/gn41sa5/1.4+Service+Aspects+of+Assurance>
- [ITU-T X.1252]** Recommendation ITU-T X.1252. Cyberspace security – Identity management.
Baseline identity management terms and definitions.
- [REFEDS ePSA]** REFEDS ePSA usage comparison v0.13.
https://blog.refeds.org/wp-content/uploads/2015/05/ePSAcomparison_0_13.pdf
- [SurveyRes]** Level of Assurance survey for SP communities - Summary of interviews
<https://wiki.geant.org/display/AARC/LoA+Summary+of+Interviews>



Glossary

AAI	Authentication and Authorisation Infrastructure
IdP	Identity Provider
LoA	Level of Assurance
SAML	Security Assertion Mark-up Language
SP	Service Provider