THE DOCUMENT IS NOW CAST ON THE PROJECT MILESTONE TEMPLATE. THIS DOCUMENT IS OBSOLETE AND NOT UPDATED ANY MORE.

AARC project

# Recommendations on minimal assurance level relevant for low-risk research use cases

Milestone MNA3.1 (due in November 2015)

## Scope

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.

## Method

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

## 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 query based on a web form) 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 (physics), and WLCG (physics). Following e-infrastructures were interviewed: EGI and PRACE.

## 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 WLGC) 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 use. Example: LIGO and WLCG.

The interview questions and results are summarised in another document.

## Related work

In parallel, the GEANT GN4-1 project (Service Activity 5, Task 1.4) 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”.

# Minimal assurance profile

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

## Requirements for the minimal assurance profile

It is not expect 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/logins (see “Implementation note” below).

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

No shared accounts must exist and the Home Organisation must be able to trace an 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 reassign of user identifiers)

The Home Organisation must provide a persistent identifier for a user. The identifier must have the property that it is never reassigned, i.e., circulated 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 should be certain widely approved good practices for the password quality, such as length, complexity, and change cycle.

5. **Departing user'seduPersonAffiliation changes promptly**

When a user departs from his/her Home Organisation, his/her eduPersonAffiliation value (and derivatives, such as eduPersonScopedAffiliation) should 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 (see REFEDS ePSA paper). 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 should 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.

# Implementation note

How the requirements presented above are implemented to 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.
- When a user is authenticated, the Home Organisation releases an indication to the Service Provider if the authenticated user complies to 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 authentication context assertion.

# Appendix 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 IdP's SAML metadata and sending a log-in link to that email address
- The tool presents structured self-assessment questions to the IdP/IdM admin
    - Quantitive ("do accounts belong to an individual")
    - 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 assurance.edugain.org/<idp-entity-id>
- The tool would evaluate whether 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 tag for the IdP, possibly directly to eduGAIN metadata (MDS)
    - an Entity Category attribute needed to be defined for that
- The tool would asks the Identity Provider admin to do a new self-assessment every year
    - Missing self-assessment would make the Entity Category attribute disappear from eduGAIN metadata
- The tool could assist in the 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.