Guidelines for evaluating the combined assurance of linked identities

Publication Date	4/24/2018
Authors:	Davide Vaghetti (ed.);Mikael Linden;Nicolas Liampotis;David Hübner;Jens Jensen
Document Code:	AARC-G031
DOI:	*deferred*
Grant Agreement No.:	730941
Work Package:	JRA1
Task Item:	TJRA1.3
Lead Partner:	Consortium GARR

The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 730941 (AARC2).

Abstract

In order to assign an assurance profile to their users, Infrastructures shall evaluate the assurance of the linked identities used for registration and authentication. These guidelines provide a method to combine assurance information and to compensate for the lack of it.

Table of Contents

1. Introduction

The Research Infrastructures (including e-infrastructures and cyber-infrastructures - from now on jointly referred to as “Infrastructures”) that follow the AARC Blueprint Architecture [AARC-BPA] set up their own AAI to grant access to their services. The AAI is typically based on a central IdP-SP proxy that acts as a gateway for the Infrastructure services and resources. Infrastructures AAIs rely on external identity providers and employ identity linking services in order to authenticate their users.

The Infrastructures also define one or more assurance profiles [AARC-G021], or a combination of assurance components, tailored to a specific risk assessment. To assign an assurance profile (or a set of assurance components values) to their users, the Infrastructure shall evaluate the assurance of the linked identity, or identities, used to register to the Infrastructure’s AAI.

2. Definitions

2.1. Infrastructure identity

In the context of research collaborations, the user is typically assigned an identity by the Infrastructure. This “Infrastructure identity” consists of a personal, unique, non-reassignable, non-targeted identifier, and additional attributes containing profile information about the user, as well as group membership and role information. The Infrastructure identity can be associated with a set of credentials issued by the Infrastructure itself, but the identity bootstrap is generally accomplished through an external identity provider.

2.2. External identities

In this document, it is assumed that the user links external identities [AARC-G008] in the proxy [AARC-BPA]. The user has access to more than one external identity provider (external to and independent of the infrastructure), be they home organisation, social media, community managed virtual organizations, etc. [AARC-MJRA1.2]. Each identity provider (hereinafter, IdP) provides different personal identity attributes (name, email), affiliation (organisational affiliation, community membership) and assurance information which the proxy combines together to create the infrastructure identity.

2.3. Identity linking

The Infrastructure leverages external identities to authenticate the users and to bootstrap the Infrastructure Identity. External identities are thus linked to the Infrastructure Identity [AARC-

G008]. The minimal technical requirement to link an external identity is the availability of a persistent, non-reassignable, and unique external identifier.

2.4. Effective identity

When multiple external identities are linked to the Infrastructure Identity, the user has multiple authentication options as well. In this context, we will refer to the identity that will be used to authenticate to the Infrastructure as the effective identity.

3. Combined assurance evaluation

In the context of this document, we will use the definition of Assurance as expressed in the REFEDS Assurance Framework [RAF]. The RAF is still a draft, but the REFEDS community extensively discussed the key concepts defined by the framework and used in this document. When the final version of the RAF will be published this document will be updated accordingly.

Along the lines of other recent assurance guidelines [NIST.SP.800.63.3] and proposed standards [VOT], the RAF does not use the concept of level(s) of assurance, rather it splits assurance into separate components. The RAF considers the following three components:

- Identity uniqueness.
- Identity proofing and credential issuance, renewal and replacement.
- Attribute quality and freshness.

The assurance values are represented using the eduPersonAssurance attribute [eduPerson] in case of SAML 2.0, or using the eduPersonAssurance claim as defined by the REFEDS OIDCre working group [OIDCre] in case of OIDC. Please refer to the RAF current draft for more detailed information.

A requirement for the assurance evaluation is that assurance components related to the same individual, but coming from different IdPs, are defined along the lines of the RAF, or can be translated into those definitions. When no assurance information is directly provided by the IdP during the authentication, the Infrastructure SHOULD NOT make any assumption on the assurance of the external identities, but it can rely on other evidences and compensatory control to ascertain the relevant assurance features of the incoming identity, as it will be shown in the following sections on a component by component base.

The components SHOULD eventually be collapsed to compose assurance profiles, each consisting of a set of values for one or more of these components. Please consult AARC-G021: Guideline on the exchange of specific assurance information between Infrastructures [AARC-G021] as a reference for the available assurance profiles.

3.1. Identifier uniqueness (ID)

The RAF ID component describes “how a CSP expresses that an identifier represents a single natural person and if that person remains the same over time” [RAF].

When an external identity provider asserts the ID component value unique, no further evaluation is to be made by the Infrastructure, and the value SHOULD be treated verbatim. The evaluation SHOULD be performed at the time of the identity linking.

3.1.1. Combined evaluation

When combining ID component values that belong to two or more linked identities, the value for the Infrastructure identity SHOULD be calculated with an AND operation where a value `unique` is equal to TRUE and a value N/A (no available value) is equal to FALSE. As in:

```
ID_value = ID_value_1 AND ... ID_value_n
```

Possible combinations and values with two linked identities:

Linked Identity 1 ID value	Linked Identity 2 ID value	Infrastructure Identity ID value
`unique`	`N/A`	`N/A`
`N/A`	`unique`	`N/A`
`unique`	`unique`	`unique`

Table 3.1. ID component combinations with two linked identities

As it is apparent, the outcome of the combined evaluation will make it impossible to assert the value `unique` for the Infrastructure Identity when one of the linked identities lacks it. This is wanted to prevent the use of shared and reassignable accounts associated with properly ID unique value accounts.

3.1.2. ID component compensatory controls

When an external identity provider does not assert the ID component value unique the Infrastructure SHOULD perform compensatory controls as defined by Expression of REFEDS RAF assurance components for identities derived from social media accounts [AARC-G041]. Failure to do so will expose the Infrastructure to unreasonable risks (for example non-traceability of users, shared accounts use and the like).

In the table below, you will find the compensatory controls needed to raise the assurance of the ID component to the value `unique`.

External identity provider	Compensatory controls (short name)
`Any IdP (included social media IdPs)`	`R&S_EC || (im_a_person && contacts)`[1]

[1] See section 4, “Compensatory controls” for technical details about each compensatory controls.

Table 3.1. ID component compensatory controls to assert the value unique

3.2. Identity proofing and credential issuance, renewal and replacement (IAP)

The RAF IAP component describes the quality of the identity proofing, credential issuance, renewal and replacement processes.

The possible values are:

- `https://refeds.org/assurance/IAP/low`
- `https://refeds.org/assurance/IAP/medium`
- `https://refeds.org/assurance/IAP/high`

The IAP component value MUST be asserted incrementally, that is: when asserting a value `medium`, the value `low` MUST be asserted too; when asserting a value `high`, the values `medium` and `low` MUST be asserted too [RAF].

When an external identity provider asserts the IAP component value, no further evaluation is needed.

3.2.1. Combined evaluation

When combining IAP component values that belong to two or more linked identities, the value for the Infrastructure identity will be equivalent to the value of the effective identity.

3.2.2. IAP component compensatory controls

When an external identity provider does not assert any IAP component values and a value is required by the service providers beyond the proxy, the Infrastructure SHOULD perform compensatory controls as defined by *Expression of REFEDS RAF assurance components for identities derived from social media accounts* [AARC-G041].

In the table below, you will find the compensatory controls needed to raise the assurance of the IAP component from no value to `low`.

External identity provider	Compensatory controls (short name)
`Any IdP (included social media IdPs)`	`conf_email`[2]

Table 3.2. IAP component compensatory controls to assert the value low

The table above set the requirements to assert the value low for the IAP components without the need to manage policies on a per IdP and/or Identity Federation base. However, it may well be the case that an eduGAIN IdP would qualify for higher IAP values.

[2] See section 4, “Compensatory controls” for technical details about each compensatory controls.

When the Infrastructure needs to assert values above `low`, it SHOULD consider evaluating both the Identity Federation policy and the assurance information published in the metadata of the incoming IdP. All the policies of the Identity Federations that belong to eduGAIN are published on the eduGAIN Technical site [eduGAIN-TECH].

3.3. Attribute quality and freshness (ATP)

The ATP component describes the quality and the freshness of the attributes the IdP delivers to the SP (in this case the SP side of the IdP/SP proxy of the Infrastructure). Current values are limited to represent the freshness of the affiliation attributes defined in eduPerson: eduPersonAffiliation, eduPersonScopedAffiliation and eduPersonPrimaryAffiliation [eduPerson].
The permitted values are:

- `https://refeds.org/assurance/ATP/ePA-1m`
- `https://refeds.org/assurance/ATP/ePA-1d`

The values reflect the latency in updating the affiliation status of a user in case of departure or role change. The values are hierarchical, that is when asserting `ePA-1d` also `ePA-1m` MUST be asserted.

Try to compensate for missing ATP component values is both difficult and pointless, since, according to the *Guideline on the exchange of specific assurance information between Infrastructures* [AARC-G021], the ATP value for the Infrastructure identity "SHALL reflect the affiliation of the identity with the Infrastructure" in compliance with guideline *Exchange of affiliation information* [AARC-G025].
Nonetheless, when the ATP component is expressed by an external identity provider, the Infrastructure MAY rely on the ATP component value to compute its own [AARC-G041].

4. Compensatory controls

The list of compensatory controls proposed here is not meant to be exhaustive. The Infrastructure will define additional control if deemed necessary and according to its own policy and risk assessment.

4.1. I'm a person

The user registering to the Infrastructure will be required to confirm that she is a single natural person and that she will not share the account with other people. Those requirements MAY also be included in the Infrastructure AUP.

Rationale	Be sure that the user is a single natural person, and have a simple way to ban users that share their account for policy/AUP violation.
RAF requirement	The “I’m a person” statement is meant to meet one of the four requirements for asserting the value unique of the ID component: the “User account belongs to a single natural person” [RAF].
Enforcement	The “I’m a person” statement itself cannot prevent bad actors and misbehaviour, but it gives a solid ground for banning or suspending malevolent or careless users. Failure to confirm the statement will prevent the user to access the Infrastructure.
Short name	`im_a_person`

4.2. Contacts

When a user register to the Infrastructure, their (external) identity providers will be required to release contacts information as email or mobile phone number. The “Confirmation mail” compensatory control can substitute “Contacts”, but not vice versa.

Rationale	Have a mean to contact the user.
RAF requirement	The “Contacts” control is meant to meet one of the four requirements for asserting the value unique of the ID component: the “CSP can contact the person to whom the account is issued” [RAF].
Enforcement	The failure to release contact information by the external IdP can have two different outcomes: the user cannot access the Infrastructure or she will be asked to insert the missing information.
Short name	`contacts`

4.3. Research and Scholarship entity category

eduGAIN IdPs asserting the support for the REFEDS Research and Scholarship entity category [REFEDS-R&S] commit to release a set of attributes following specific rules on the quality of the identifier.

Rationale	Reuse the entity category rules about the identifier.
RAF requirement	Support for REFEDS R&S meet all the requirements of the value unique of the ID component.

Enforcement	Failure to detect support for the entity category in the IdP metadata should activate the other compensatory controls.
Short name	`R&S_EC`

4.4. Confirmation email

When a user wants to register to a service, it is common practice to send an email to the provided address with a confirmation link. Once received, the user will follow the link to complete the registration process. This process guarantee that the email is both valid and in control of the user. The Infrastructure will embrace the same process for the users' registration.

Rationale	Obtain a verified email address for each user registering to the Infrastructure.
RAF requirement	The confirmation email the basic requirement for the value low of the IAP component.
Enforcement	Failure to provide a valid email address, or to follow the link sent via the confirmation email, will prevent the user to access the Infrastructure.
Short name	`conf_email`

5. Authentication assurance

The RAF does not cover the assurance quality of the authentication process. However, the REFEDS Assurance Working Group [REFEDS-AWG] has defined two authentication assurance profiles that MAY be paired with the RAF assurance components:

- REFEDS Multi-Factor Authentication (MFA) Profile (https://refeds.org/profile/mfa)
- REFEDS Single-Factor Authentication (SFA) Profile (https://refeds.org/profile/sfa)

Please refer to the working group website [REFEDS-AWG] for further details on the profiles.

Whether or not the Infrastructure will evaluate the authentication assurance expressed by external identity provider, the authentication assurance value cannot be combined. The value will always be equal to the one expressed for the effective identity.

References

AARC-BPA	AARC Blueprint Architecture. https://aarc-project.eu/architecture/
AARC-MJRA1.2	Design for Deploying Solutions for “Guest Identities”. https://aarc-project.eu/wp-content/uploads/2016/06/MJRA1.2-Design-for-Deploying-Solutions-for-Guest-Identities.pdf
AARC-G008	Implementing SAML authentication proxies for social media identity providers. https://aarc-project.eu/guidelines/aarc-g008/
AARC-G021	Guideline on the exchange of specific assurance information between Infrastructures. https://aarc-project.eu/guidelines/aarc-g021/
AARC-G041	Expression of REFEDS RAF assurance components for identities derived from social media accounts. https://aarc-project.eu/guidelines/aarc-g041/
RAF	(Draft) https://docs.google.com/document/d/15v65wJvRwTSQKViep_gGuEvxLl3UJbaOX5o9eLtsyBI
NIST.SP.800.63.3	NIST Special Publication 800-63-3, *Digital Identity Guidelines*, June 2017, https://doi.org/10.6028/NIST.SP.800-63-3
VOT	(Draft) https://www.ietf.org/id/draft-richer-vectors-of-trust-09.txt
eduGAIN	https://edugain.org/
eduGAIN-TECH	https://technical.edugain.org
eduPerson	Internet2/MACE. eduPerson Object Class Specification (201602). http://software.internet2.edu/eduperson/internet2-mace-dir-eduperson-201602.html
REFEDS-OIDCre	https://wiki.refeds.org/display/GROUPS/Mapping+SAML+attributes+to+OIDC+Claims
REFEDS-AWG	https://wiki.refeds.org/display/GROUPS/Assurance+Working+Group