01-04-2017

Deliverable DNA3.1: Differentiated LoA recommendations for policy and practices of identity and attribute providers, applicable to research use cases

Deliverable DNA3.1

Contractual Date:	31-03-2017
Actual Date:	01-04-2017 (rev1.1)
Grant Agreement No.:	653965
Work Package:	WP3 Policy and Best Practice Harmonisation
Task Item:	TNA3.1
Lead Partner:	CSC
Document Code:	policy harmonisation differentiated assurance
Authors:	Mikael Linden (CSC), Niels van Dijk (SURFnet), Licia Florio (GÉANT), David Groep (Nikhef), Christos Kanellopoulos (GÉANT), Ian Neilson (STFC), Hannah Short (CERN), Uros Stevanovic (KIT), and Gerben Venekamp (SURFsara); *with contributions to the REFEDS Assurance Profile from wider membership of the REFEDS Assurance Working Group, including Pål Alexsson (SUNET), Jim Basney (NCSA), Tom Barton (UChicago), Nicole Harris (GÉANT), Nick Roy (Internet2),*

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

Abstract

Differentiated Assurance Recommendations for identity attributes originating at federated home organisations (and attribute authorities for AAI gateways) were developed. The recommendations are based on research and e-Infrastructure requirements and assessed feasibility of implementation by the federations and home organisations. Using a component-based approach to *assurance profiles*, and with broad input from the global community through the REFEDS Assurance Framework working group established at the initiative of AARC to attain global consensus, a limited set of profiles was developed that addresses three recognised risk profiles of the Infrastructures.

Table of Contents

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. Yet when leveraging federated identity management for access to services, the Research and e-Infrastructures have an expressed and a concrete need to evaluate identity assurance information for the attributes and authenticator presented to them by the user's home organisations via the R&E federations. The balance between service provider requirements and feasibility at the identity provider (IdP) side is complex, as each of them has different interests. The service provider, off-loading the authentication but having to protect its services based on its own risk assessment, will prefer to off-load much of the complexity of identity vetting to the IdPs, and maintain sufficient confidence in the assertions presented in a simple way. The IdPs on their side need specific guidance as to what is needed in terms of assurance, expressed in a form that allows them to make confident assertions about specific components of the assurance assertion: identifier uniqueness, identity proofing, authentication, and attribute 'freshness'. To attain this balance, a combination of structured interviews, open community consultation, and stakeholder engagement through existing groups (REFEDS and FIM4R) was used.

The resulting Assurance Framework includes a baseline assurance profile (presented earlier in AARC milestone document MNA3.1 [MNA3.1]), a set of explicit assurance components that can be implemented and assessed by the federated identity providers, and a grouping of these assurance components into non-hierarchical profiles, allowing differentiated assurance for Research and generic e-Infrastructures based on their own service risk assessments. In this deliverable we present the requirements collection process, the community engagement methodology through REFEDS and FIM4R, and the principles behind the Assurance Framework. The version of the Assurance Framework as recommended by AARC is included in the Appendix. It will subsequently be sent out for further consultation to the entire REFEDS community.

1 Introduction

“*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)”:* this definition, introduced in the Recommendations on Minimal Assurance Level Relevant for Low-risk Research Use Cases [MNA3.1], remains a useful tool when approaching ‘differentiated assurance’ since also in the Blueprint Architecture (BPA) published by the AARC project [BPA] the home organisations remain an anchor point for both ‘authenticator’ trust (the password, token, or other credentials) as well as for many of the attributes used to identify users within their communities.

The BPA also introduces other actors in the collaborative workflow: attribute authorities operated by or on behalf of Research and e-Infrastructures, token translation services, and AAI gateways. Depending on the usage scenario, these can either employ existing user information to organise them into groups and assign roles and privileges, or they can augment existing information with user assurance elements. Both of these are seen in real-life deployments: e.g. many community membership services in the EGI infrastructure add group membership to their users, but rely entirely on the initial authentication authority for the binding of credential to user, and for the name of the users (and then use such information to manage the community). Communities such as ELIXIR add also trust information to user records in the AAI gateway (‘bona-fide researcher’), and for the communities in WLCG the identity information is entirely held in the community-managed attribute authority, and it relies for the home organisations only for the binding of credential to the user.

Yet in all these cases, the role of credential service provider (CSP) and relying party (RP) can be distinguished, with the RPs putting specific requirements (and implicit expectations) upon the CSPs when authenticating entities. Attribute authorities can then take both roles simultaneously, yet the logical concepts remain equally valid. So although this document provides recommendations for CSPs and RPs, it is similarly applicable for the attribute authorities and AAI gateways that customarily front the Research and e-Infrastructures.

In describing assurance recommendations, it is impossible to ignore the large body of existing work in this area. Without unduly reiterating existing work, we take note of the Entity Authentication Assurance Framework of ISO/IEC 29115:2013 [ISO29115], the Kantara Identity Assurance Framework [KIAF], and the (more USA-centric) NIST SP800-63 guidance [SP80063]. All of ISO/IEC, Kantara, and NIST SP800-63 (revs. 1 and 2) provide for four levels of assurance, in which a combination of controls and assurance elements are combined. The eIDAS regulation and its implementation directives [EIDAS] introduce an assurance framework based on three levels (“low”, “substantial”, and “high”), and similarly combines various controls and elements when defining the (minimum) requirements for each level. All these frameworks basically define a monolithic structure by which to express levels of assurance.

Especially SP800-63 in its early version (revisions 1 and 2) significantly influenced the assurance framework that was introduced early 2013 in the InCommon research and education federation in the US [INCAP] This InCommon framework provided for two identity assurance profiles, Bronze and Silver, that were closely aligned with the SP800-63 levels 1 and 2, including the need for external audits and certification. Outside of the InCommon federation, this model was not adopted, and within InCommon there is – as of early 2017 – no single identity provider certified to the Silver level. Although the framework names Bronze and Silver as profiles, they align closely with the levels from SP800-63 and are monolithic in nature.

Meanwhile, other work in the IETF around Vectors of Trust (VoT [VOT]) and the consultation process of NIST SP800-63 revision 3 [SP80063V3] have promoted a component-based approach. Instead of amalgamating assurance elements, they are decomposed into independently-assessable 'vectors'. In the IETF VoT strawmen for example, these are identity proofing, credential strength, assertion presentation, and operational management.

Decomposition aids in identifying the source of a basic imbalance in the (R&E) federated identity management system: the service providers have ambitions expectations in what level of identity trust can be provided by the home organisations ("everything", but at zero cost to the SP), and what the identity providers are willing to offer ("good enough for internal systems", at no additional cost, and explicitly described in detail so as not to incur additional liability).

The Baseline Assurance requirement identified in the Recommendations on Minimal Assurance Level [MNA3.1] listed six elements, which – to the service providers in the Infrastructure – were clear and unambiguous. Yet in the consultation phase with federations and identity providers, it became clear that statement such as "there must be certain widely approved good practices for the password quality, such as length, complexity, and change cycle", or "follow a widely established structure" were not sufficiently explicit to allow implementation on a wide scale. The Infrastructures (for their 'limited risk' use cases) were deliberately open to variations in an implementation process, but to obtain common understanding more detailed guidance was needed. The component-based approach allows this level of detail to be expressed – both in the requirements as well as by the federated identity provider in their issued assertions.

Yet the level of complexity introduced by the component approach ("how should the components be combined to assure sufficient trust levels for this application?") significantly complicates the job of the service provider. Access and authorization frameworks, even if they are able to process large numbers of attributes for each decision they need to make), require a risk assessment and a policy that is an expression of that risk assessment. Complicating this already challenging task by assigning assurance vectors in a multi-dimensional space is not helpful in addressing the issue at hand: granting trustworthy access to potentially very valuable resources.

To support the service providers, Research and e-Infrastructures, and AAI gateways in assessing the assurance quality, this Differentiated Assurance Recommendation introduces groupings of assurance components, based on the analysis of both the 'low-risk' use cases introduced in MNA3.1, as well as more complex access cases for Infrastructure use and enrolment in dynamically-structured communities (where the assurance requirements on the initial identity assertion are higher), as well as for use cases (mainly from the biological and medical sciences domain) where there is the need to ensure truly verified identities.

However, contrary to the monolithic approach used for the assurance levels by Kantara, ISO, NIST and eIDAS, the component framework recommended here allows for more dynamic groupings. It is for this reason that we explicitly refer to these groupings as assurance profiles (APs), and the profiles are – by design – not on a linear scale. The APs can be combined in authorization decisions alongside complementary assurance from other sources (combined assurance or combined adequacy, as is being used in e.g. the EGI infrastructure), and with other attributes. To emphasise the non-linear nature of the APs, it was decided to assign to each of them an non-orderable name, akin to the naming scheme of the IGTF for their assurance profiles. In this case, the names of popular coffee flavours have been choses after a community consultation.

2 Description of the assurance requirements and recommendations process

Federated identity assurance is most effective when definitions and profiles are common among all participants in the identity ecosystem: so not only the Infrastructures (and their closely associated identity providers of last resort), but also by at least the majority of the home organisations across the world. This means that bodies such as REFEDS, the Research and Education FEDerations group that articulates the mutual needs for collaboration and harmonisation of R&E federation operators, play an important role. This role is two-fold: it is a vital input in determining which assurance requirements are feasible (i.e. implementable and managerially supported by federations and their members), but they also provide the primary means of communicating with the federations and – through them – the user's home organisations.

Similarly, the 'generic' e-Infrastructures (such as EGI and PRACE) and the collaboration group of research infrastructures (FIM4R, the Federated Identity Management for Research group that articulates the joint requirements of many of the research infrastructures [FIM4R]) helps in arriving at a grouping of assurance components into a limited and yet widely supported set of assurance profiles.

This section describes the process to develop the differentiated LoA assurance recommendations.

2.1 Requirements Gathering among Research Communities

To understand the needs of the research communities, structured interviews were carried out with representatives of six research infrastructures and two e-infrastructures. The structured interview method was chosen - over a form based query - because it provides the interviewer with an opportunity to acquire a deeper understanding of the infrastructure's underlying needs and use scenarios.

The research infrastructures interviewed were selected both from the AARC project participants as well as from the wider FIM4R community, and include CLARIN (language research); DARIAH (arts and humanities); ELIXIR (life science); LIGO (physics); photon/neutron facilities (multi-disciplinary); and WLCG (physics).
Representing the generic e-Infrastructures (which serve a more diverse set of applications and – also through national contributions – a part of the 'long tail' of research) were EGI (the federated e-Infrastructure set up to provide advanced computing services for research and innovation), and PRACE (the Partnership for Advanced Computing in Europe).

2.2 Drawing the Recommendations from the Requirements

Based on the requirements gathered, a Recommendation on minimal assurance level relevant for low-risk research use cases was developed and published as the milestone MNA3.1 in November, 2015.

The Recommendation on minimal assurance level can be summarised with the following six requirements:

1. The accounts in the Home Organisations must each belong to a known individual person
2. Persistent user identifiers (i.e., no re-assignment of user identifiers)
3. Documented identity vetting procedures (not necessarily face-to-face)
4. Password authentication (with some good practices)
5. Departing user's eduPersonAffiliation must change promptly
6. Self-assessment (supported with specific guidelines)

The Recommendation was further exposed to a public consultation within the REFEDS and FIM4R communities in December 2015 and January 2016. During the public consultation, in total 25 comments on the recommendation was received.
A key finding was that the recommendation alone isn't sufficient as a framework that the Home Organisations managing the Identity Provider servers could deploy. The recommendations need to be extended to specific requirements that can be mounted on the identity federations utilising the federated identity protocols. This led to the project joining forces with the REFEDS to develop a REFEDS assurance profile that is described below.

The comments also suggested that more research communities need to be involved in the work to establish assurance level recommendations which are widely accepted. In order to enable researchers to collaborate internationally, a global instead of a European approach was seen as preferred.

Other comments during the public consultation expected more detailed requirements, such as on the unique identifiers, requirements for password authentication and freshness of the eduPersonAffiliation attribute.

2.3 Spin-off: Self-assessment tool

One of the recommendations was that, instead of external or internal audits, a self-assessment should be sufficient as the method to evaluate the assurance level an Identity Provider can provide. On the other hand, the research community interviews emphasised that the self-assessment should be supported with specific guidelines, such as checklists, that the Identity Provider administrators need to go through to back up their positive self-assessment.

In the appendix of the milestone document MNA3.1, the Level of Assurance task proposed designing and implementing a self-assessment tool that the federation operators could use to invite the Identity Provider administrators to make a self-assessment. The tool would then present structured questions to lead the administrator through the self-assessment and, once complete, submit the results back to the federation operator. If an assurance profile expected peer review of the assessments, the tool could help to carry out them, too.

It turned out that the Level of Assurance task shared the requirement for a self-assessment tool with the Incident response task (NA3's task 2). Together the tasks developed a software requirements specification for the tool and the requirements specification was then handed to the GÉANT project (GN4-2 JRA3 Task 1) which is studying the alternatives to implement the tool and offer it as a service to the federations.

2.4 Working with REFEDS assurance working group

After the public consultation of the Recommendation on a minimal assurance level, it was found that further assurance level development was required

1. in an *open* forum. A closed project like AARC is not optimal for developing specifications who needed to be adopted widely.

2. in an *international* forum. Research is international and a specification developed and adopted in Europe only does not remove obstacles from international collaboration.

3. in a forum which has *close ties with the identity federations*. The actual roll-out of any profile requires actions from the Home Organisations managing the Identity Provider servers, and the federation operators are the actors who have connections to them in their own constitutions.

In June 2016, the AARC project proposed to REFEDS that it establishes an open and international working group for assurance that continues the work the AARC project had started. The proposal was approved by REFEDS and the AARC experts continued the work in the newly established REFEDS assurance working group. More experts joined the work from the United States, including the Home Organisations and research communities. The working group had biweekly calls and in February 2017, it delivered the first draft of a

REFEDS assurance framework (Appendix) which will be exposed to a community consultation following the REFEDS procedures.

Since also the use of multi-factor authentication (MFA) is an assurance element, but the mechanics of MFA are technologically different from the other assurance elements (identifier uniqueness, identity proofing, and attribute freshness), within the context of REFEDS and the GÉANT project a dedicated profile on expressing MFA is being developed. In order to align with that work, the consultation processes of the Assurance Framework and the MFA profile have been coordinated, and will finish together so that the Assurance Framework can leverage the then-existing standard representations to express the MFA authentication (“authentication context”). The work on MFA is deferred to the REFEDS GÉANT work, and further described there. The expected results (single factor, so a “good entropy”, or at least two-factor, when expressing “https://refeds.org/profile/mfa” in the authentication context) are used within this Assurance Framework, but their detailed syntax left until the MFA consultation has completed.

2.5 Future Steps

After the consultation, it is expected that a pilot will take place on the REFEDS assurance framework. Although the baseline assurance requirements have already been used in selected AARC pilots (in particular in the CILogon-like service for Europe, where these are used in conjunction with the REFEDS Research and Scholarship specification and Sirtfi), the specific assurance profiles (groupings of assurance components) proposed in the REFEDS Assurance Framework should be evaluated in a real-life production environment. This requires that the consultation process has been completed and the profiles are stable. Thus, piloting the result will happen only once that global consultation process has converged, and within the context of the AARC2 project. After the pilot, the uptake and the necessary outreach and training effort is expected to be carried out together with the GÉANT project (GN4-2 JRA3).

3 Description of the Assurance Profile

Because of the need for global contributions and wide adoption, the differentiated assurance profiles have been developed as a separate document, the REFEDS Assurance Framework. The details thereof are presented here in Appendix A, subject to the ongoing global consensus process which is expected to finish by the end of the second quarter of 2017. In this section, we describe the key properties of the profile and the rationale for the choices made therein.

3.1 Distributing responsibilities

The REFEDS assurance framework identifies two actors:

- Credential Service Providers (CSP) who manage the users' identities and attributes and authenticate them when they access the Service Providers. In the context of research and education, CSPs are typically researcher's Home Organisations who manage an Identity Provider server, but also research and e-infrastructures can assume the role of a CSP in their constituencies.

- Relying Parties (RP) who are the Service Providers relying on the identity assertions made by the CSPs

The REFEDS assurance framework focuses mostly on CSPs who both authenticate the users and manage their attributes. Some sections may be also useful for Attribute Providers who do not authenticate users but just serve the RPs with extra attributes on authenticated users.

3.2 Differentiated assurance level recommendations

The REFEDS assurance profile adopted an approach where assurance was split to four components:

1. Identifier uniqueness
2. Identity proofing and credential issuance, renewal and replacement
3. Authentication
4. Attribute quality and freshness

The assurance profile defines one or more values for each components. A Credential Service Provider (CSP, such as a Home Organisation or research or e-infrastructure running an Identity Provider server and the associated identity management system) can manage and issue one or more values from one or more components to each authenticated user and pass them to the Service Provider.

This gives the Service Provider the flexibility to require those values they need. The specific components can be used individually based on a per-service risk assessment. For example:

- some Service Providers are supposed to ensure the access is closed for a researcher who departs from their home organisation. They are supposed to observe the freshness of the eduPersonAffiliation attribute

- some Service Providers emphasise that the researcher's identity must be verified carefully when they register an account in their Home Organisation. Those Service Providers are supposed to pay attention to the Identity proofing component

- some Service Providers are known to completely ignore the Identity Proofing carried out by the CSP. Instead, the researchers can use whatever identity they want to register to the Service Provider and their identity is verified out-of-band by the research community. Those Service Providers can ignore the Identity proofing component.

However, it was also concluded from the interviews and the consultation process that many Service Providers seek for ‘simplicity’ and have no need to each by themselves make decisions on the detailed components if they share a common risk profile (or risk perception, which may not be the same but results in the same behaviour towards identity assurance). To serve those Service Providers (or service provider groups such as the Infrastructures), assurance profiles were assembled using the components. So as to make clear that no hierarchy is implied, and opaque naming scheme was chosen. Such a naming scheme aids service providers in that it is clear that ‘higher is not always better’, and that each profile should match a specific service risk assessment. Opaque naming schemes without implied hierarchy are non-trivial (even “A”, “B”, “C” implies a hierarchy, and even has an ordering that may be culturally determined). The Interoperable Global Trust Federation IGTF has pioneered such schemes based on tree-names (birch, dogwood), which is used in many e-Infrastructures. For this global endeavour, a community survey on suitable naming was conducted in REFEDS, with the result that the assurance profiles are named after coffee drinks.

For instance, “Cappuccino” profile has the following component properties:

- user identifiers are unique,
- their identity proofing is done for instance by shipping their credentials to their registered address of record,
- their authentication is done with quality passwords, and
- their eduPersonAffiliation attribute is supposed reflect their departure in one month.

A Service Provider can observe which CSPs can provide these factors to (at least some of) their users and require their use for user authentication.

At the moment, two profiles are defined. The “Cappuccino” profile aligns with the risk assessment done by the Infrastructures for access to compute and research data that is not otherwise sensitive personal data. The “Espresso” profile, requiring also multi-factor authentication and verified identity, is driven by the risk assessment for the biological and medical sciences where also personal research data is processed and access needs to be verifiably real-person controlled.

4 Conclusions

When leveraging federated identity management for access to services, the Research and e-Infrastructures have an expressed and concrete need to evaluate identity assurance information for the authentication presented to them by the user’s home organisations via the R&E federations. The balance between service provider requirements and feasibility at the identity provider (IdP) side is complex, as each of them has different interests. The service provider, off-loading the authentication but having to protect its services based on its own risk assessment, will prefer to off-load much of the complexity of identity vetting to the IdPs, and maintain sufficient confidence in the assertions presented in a simple way. The IdPs on their side need specific guidance as to what is needed in terms of assurance, expressed in a form that allows them to make confident assertions about specific components of the assurance assertion: identifier uniqueness, identity proofing, authentication, and attribute 'freshness'. To attain this balance, a combination of structured interviews, open community consultation, and stakeholder engagement through existing groups (REFEDS and FIM4R) was used.

The structured interviews with Infrastructures and research communities resulted in a baseline assurance profile consisting of six elements. Broad consensus was reached on the need for these elements, but pilots in the R&E federation community identified to need to express the requirements on IdP in an explicit, unambiguous manner. Although the infrastructure service providers deliberately left open some of the implementation details (based on their implicit understanding of 'reasonable'), such implicit mechanisms do not work beyond a single community. To attain broader understanding and consensus, a REFEDS working group was formed with participants from both R&E federations, selected IdPs and the Infrastructures to arrive at an Assurance Profile specification that addressed both the need for explicit guidance (in terms of assurance elements) as well as the need of the service providers for combinations (profiles) that match common risk assessments. REFEDS provided the open forum in which this discussion could be held as a global level, and in the framework of a defined consultation process.

The (consultation) version of the REFEDS Assurance Framework, presented in appendix A, embodies the consensus of the AARC and REFEDS Assurance WG participants, and the non-hierarchical profiles ("Cappuccino" and "Espresso") align with identified use cases in the existing infrastructures. In leverages the work in the GÉANT project lead REFEDS multi-factor authentication (MFA) working group on authenticator assurance, and both profiles (Assurance Framework and MFA Profile) are expected to be used in conjunction by the IdPs and service providers.

The baseline assurance profile has been used in the AARC CILogon-like token translation service for Europe pilot ("RCauth.eu"), and the Assurance Framework developed here will be piloted by the Infrastructures in their production infrastructure alongside the assurance harmonisation that is ongoing between the convergent infrastructures in the European Open Science Cloud.

Appendix A

REFEDS Assurance Profile (Consultation version)

The following text is copied verbatim from the REFEDS Assurance Framework draft that was developed through the AARC contribution to the REFEDS working group. The elements highlighted in yellow are deliberate placeholders that identify where input from the REFEDS MFA Profile - under simultaneous development in the GÉANT Project - is to be merged.

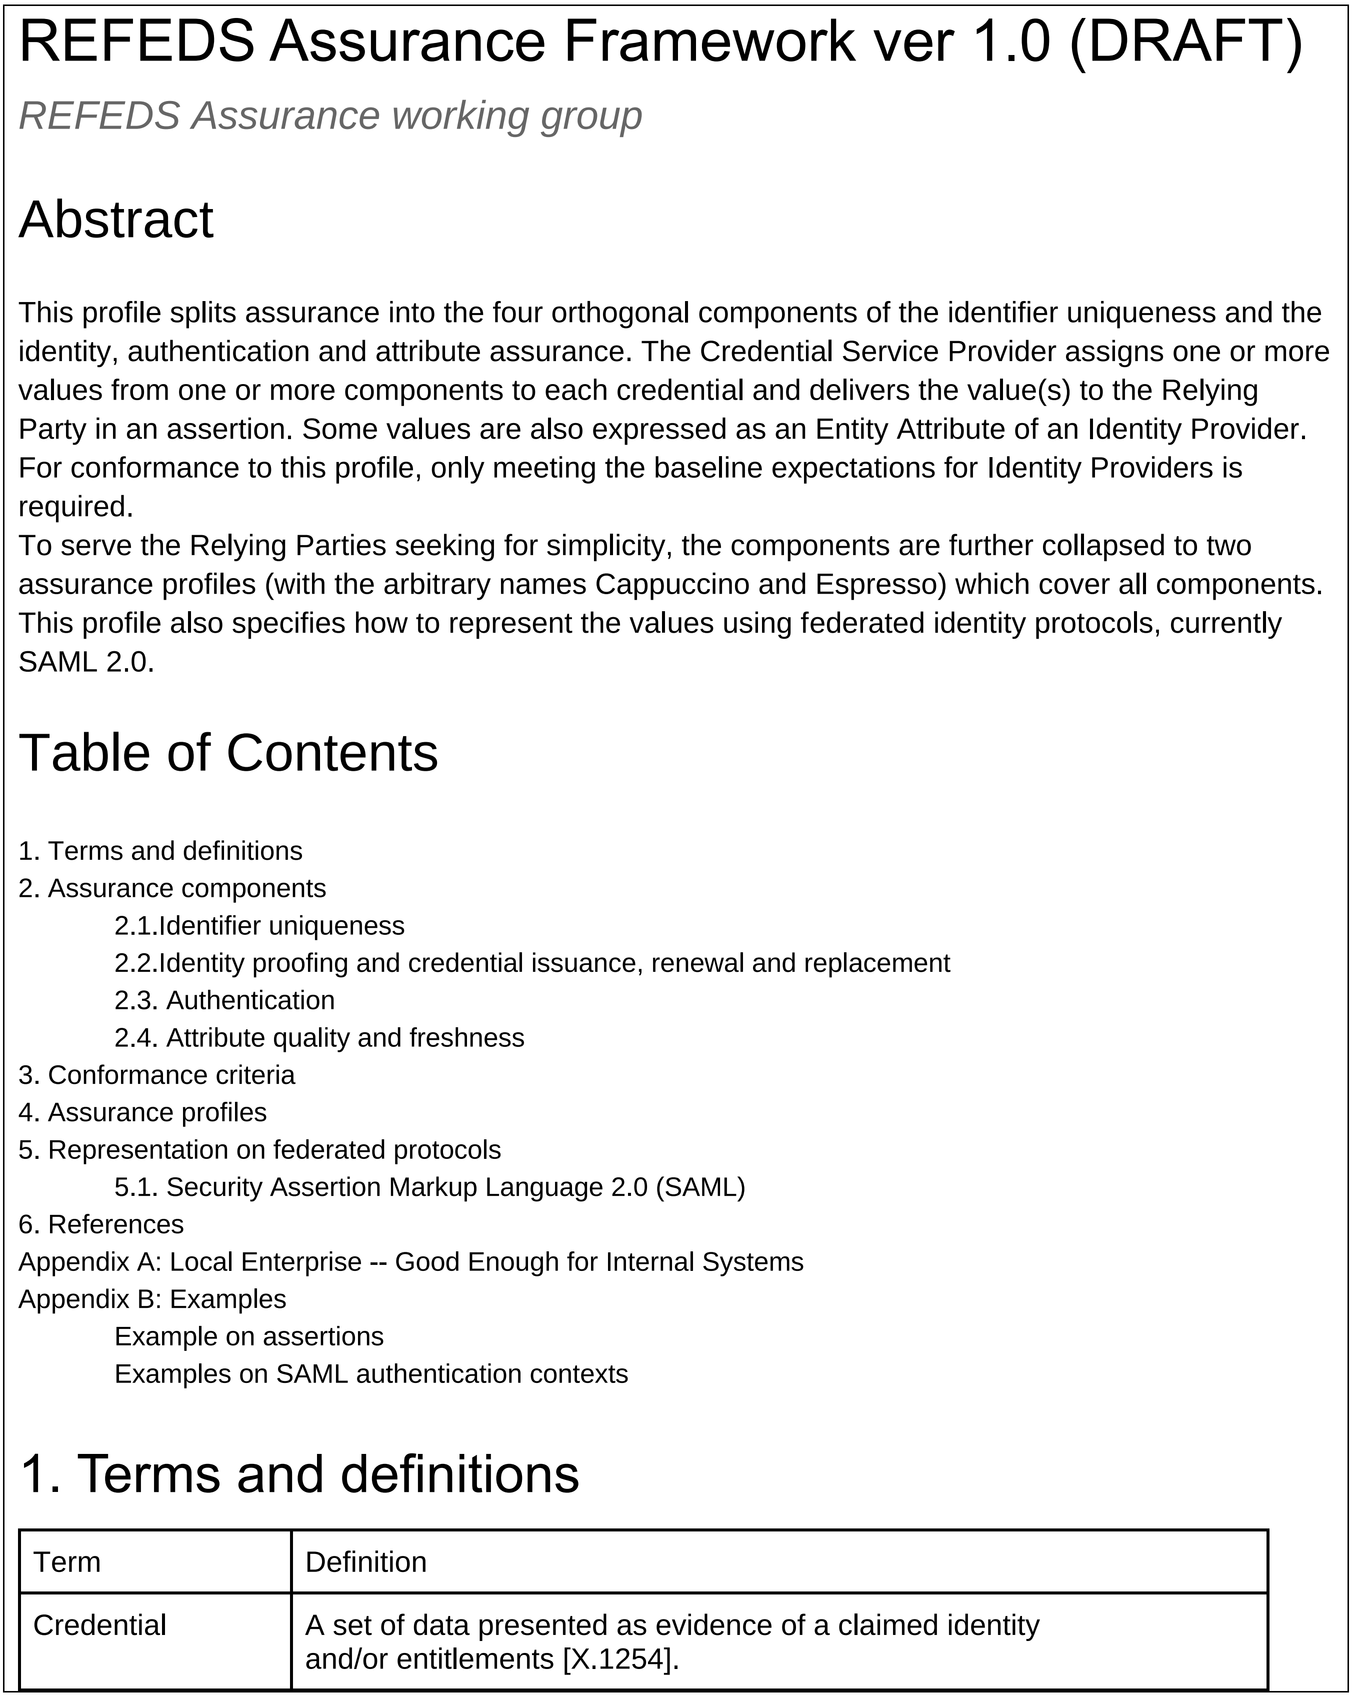

REFEDS Assurance Framework ver 1.0 (DRAFT)

REFEDS Assurance working group

Abstract

This profile splits assurance into the four orthogonal components of the identifier uniqueness and the identity, authentication and attribute assurance. The Credential Service Provider assigns one or more values from one or more components to each credential and delivers the value(s) to the Relying Party in an assertion. Some values are also expressed as an Entity Attribute of an Identity Provider. For conformance to this profile, only meeting the baseline expectations for Identity Providers is required.

To serve the Relying Parties seeking for simplicity, the components are further collapsed to two assurance profiles (with the arbitrary names Cappuccino and Espresso) which cover all components. This profile also specifies how to represent the values using federated identity protocols, currently SAML 2.0.

Table of Contents

1. Terms and definitions

Term	Definition
Credential	A set of data presented as evidence of a claimed identity and/or entitlements [X.1254].

Credential Service Provider (CSP)	A trusted actor that issues and/or manages credentials [X.1254]. In the context of this specification, CSP refers to the Identity Provider and the associated Identity Management system that manages the user identities, attributes and authentication observed by the Relying Parties.
No re-assignment (of an identifier)	No re-assignment means that while a user can be assigned a new identifier value (such as, an eduPersonUniqueID attribute value [eduPerson]), the old value MUST NOT be recycled to another user. However, the identifier value can be assigned back to the same user (for instance, if a departed person later returns back to the organisation).
Relying Party (RP)	Actor that relies on an identity assertion or claim [X.1254].

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

To assert the values defined in this profile to the RPs the CSPs will use URIs which has the following prefix:
`$PREFIX$=https://refeds.org/assurance`

2. Assurance components

This section introduces four assurance components which each represent a different aspect of assurance. The components are orthogonal i.e. a CSP can assert one or more values from different components independently. The value pertains to the user represented in the assertion and different users or the same user in different authenticated sessions can qualify to different values.

2.1.Identifier uniqueness

This component describes how a CSP expresses that an identifier represents a single natural person and if that person remains the same over time.

Value	Description
`$PREFIX$/ID/unique`	- User account belongs to a single natural person - The person and the credential they are assigned is traceable i.e. the CSP knows who they are and can contact them - The user identifier will not be re-assigned - The user identifier is one of these: eduPersonUniqueID, SAML2 persistent ID or eduPersonTargetedID[1]

[1] eduPersonTargetedID is a legacy attribute. The use of the SAML 2.0 persistent nameID is encouraged, instead.

Within the REFEDS community there is a long legacy of using eduPersonPrincipalName (ePPN, [eduPerson]) attribute as a human-readable user identifier despite its undefined re-assignment practice. The table below defines two alternative values the CSP can use to indicate its ePPN re-assignment practice to the RPs that prefer to use ePPN.

The values are mutually exclusive. A CSP MAY assert one of them but MUST NOT assert several.

Value	Description
`$PREFIX$/ID/no-eppn-reassign`	eduPersonPrincipalName values will not be re-assigned.
`$PREFIX$/ID/eppn-reassign-1y`	eduPersonPrincipalName values may be re-assigned after a hiatus period of 1 year or longer.

The intention is that

- if the Home organisation asserts `unique` and `no-eppn-reassign`, then also the ePPN attribute value shares the same uniqueness properties as eduPersonUniqueID (ePUID, [eduPerson]), SAML2 persistent ID and eduPersonTargetedID (ePTID, [eduPerson]).
- If the Home organisation asserts `unique` only, an ePPN value released by it is not assumed to fulfill the uniqueness property
- A user may have more than one ePPN at one time or over time, but non re-assignment means that the same ePPN value shall never refer to two different users

The expected Relying Party behaviour for observing ePPN re-assignment

- If the Home organisation asserts `no-eppn-reassign`, the Relying party knows that when it observes a given ePPN value it will always belong to the same individual
- If the Home organisation asserts `eppn-reassign-1y`, the Relying party knows that if an ePPN holder doesn't show up for one year, the ePPN holder may have been changed. A safe practice for the Relying Party is to close a user account or remove the ePPN value associated to it if the user hasn't logged in for one year.
- If the Home Organisation asserts neither `no-eppn-reassign` nor `eppn-reassign-1y`, the Relying Party cannot rely on ePPN as a unique user identifier but should use it only in combination with another identifier that is unique (such as ePTID, SAML2 persistent nameID or ePUID).

2.2.Identity proofing and credential issuance, renewal and replacement

This section describes the requirements for

- Identity Proofing, which is the process by which the CSP captures and verifies sufficient information to identify a user to a specified or understood level of assurance [X.1254].
- Credential issuance, which is the process of providing or otherwise associating a user with a

particular credential, or the means to produce a credential [X.1254].
- Renewal, which is the process whereby the life of an existing credential is extended [X.1254].
- Replacement, which is the process whereby a user is issued a new credential, or a means to produce a credential, to replace a previously issued credential that has been revoked [X.1254].

These values are incremental i.e. constitute an ordered set of levels with increasing requirements. The CSP asserting a value MUST also assert all preceding values (i.e. a CSP asserting `assumed` must also assert `local-enterprise` and a CSP asserting `verified` must also assert `assumed` and `local-enterprise` for a given user).

Value	**Description**
`$PREFIX$/IAP/local-enterprise`	The identity proofing and credential issuance, renewal and replacement are done in a way that is less than `assumed` but qualifies (or would qualify) the user to access the Home Organisation’s internal administrative systems (see appendix A).
`$PREFIX$/IAP/assumed`	Identity proofing and credential issuance, renewal, and replacement qualify to any of - sections 5.2.2-5.2.2.9, section 5.2.2.12 and section 5.2.3 of Kantara assurance level 2 [Kantara SAC] - IGTF level BIRCH [IGTF] - section 2.1.2, section 2.2.2 and section 2.2.4 of eIDAS assurance level low [eIDAS LoA]
`$PREFIX$/IAP/verified`	Identity proofing and credential issuance, renewal, and replacement qualifies to any of - section 5.3.2-5.3.2.9, section 5.3.2.12 and 5.3.3 of Kantara assurance level 3 [Kantara SAC] - section 2.1.2, section 2.2.2 and section 2.2.4 of eIDAS assurance level substantial [eIDAS LoA]

2.3. Authentication

This section describes the requirements for the user authentication. These values are incremental.

Value	**Description**
	Placeholder for a reference to REFEDS authentication context definition for good-entropy
	Placeholder for a reference to REFEDS authentication context definition for Multi-factor authentication

2.4. Attribute quality and freshness

This section describes the requirements for the quality and freshness of the attributes (other than the unique identifier) the CSP delivers to the RP.

The requirements are limited to the eduPersonAffiliation and eduPersonScopedAffiliation attributes defined in [eduPerson]. The freshness of eduPersonAffiliation and eduPersonScopedAffiliation are further limited to the following attribute values: faculty, student and member[2]. Other values and attributes are out of scope.

The freshness of eduPersonAffiliation and eduPersonScopedAffiliation intends to serve the RPs who want to couple their users' access rights with their continuing institutional role.

Value	Description
`$PREFIX$/ATP/ePA-1m`	eduPersonAffiliation and eduPersonScopedAffiliation attributes (if populated) reflect user's departure within 30 days time

"A departure" takes place when the organisation decides that the user doesn't have a continuing basis for the affiliation value (i.e., can no longer speak for the organisation in that role). The practices here may vary; for instance

- In some organisations a researcher ceases to be a faculty member the day their employment or other contract ends, in some organisations there is a defined grace period
- In some universities a student ceases to be a student the day they graduate, in some organisations the student status remains effective until the end of the semester

This value is intended to indicate only that there is a maximum latency of one month for the CSP's identity management system to reflect the user's affiliation change in their attributes.

Notice also that this section does not require that the departing user's account must be closed; only that the affiliation attribute value as observed by the RPs is updated.

3. Conformance criteria

For a CSP to conform to this profile it is REQUIRED to conform to the following baseline expectations for Identity Providers:

1. The Identity Provider is operated with organizational-level authority
2. The Identity Provider is trusted enough to be used to access the organization's own systems
3. Generally-accepted security practices are applied to the Identity Provider
4. Federation metadata is accurate, complete, and includes site technical, admin, and security contacts, MDUI information

A CSP indicates its conformance to this profile by asserting `$PREFIX$`.

[2] Values faculty, student and member appear to be used consistently across federations [ePSA Comparison].

4. Assurance profiles

To serve the RPs seeking for simplicity, this section collapses the components presented in section 2 into two assurance profiles Cappuccino and Espresso.

The CSPs who populate the assurance assertions presented in the section 2 MUST populate also all assurance profiles to which they qualify.

A CSP that asserts the assurance profile Espresso MUST assert also the assurance profile Cappuccino.

The table below defines the following assurance profiles:

- Assurance profile Cappuccino for low-risk research use cases (`$PREFIX$/AP/cappuccino)`
- Assurance profile Espresso for use cases requiring verified identity and two factor authentication (`$PREFIX$/AP/espresso)`

Value	Cappuccino	Espresso
`$PREFIX$/ID/unique`	X	X
`$PREFIX$/ID/no-eppn-reassign`		
`$PREFIX$/ID/eppn-reassign-1yr`		
`$PREFIX$/IAP/local-enterprise`	X	X
`$PREFIX$/IAP/assumed`	X	X
`$PREFIX$/IAP/verified`		X
`$PREFIX$/AAP/good-entropy`	X	
`$PREFIX$/AAP/multi-factor`		X
`$PREFIX$/ATP/ePA-1m`	X	X

For instance, if a user qualifies to all values required according to the column “Espresso” (including their multi-factor authentication was performed during the session) the CSP MUST assert also both Espresso and Cappuccino for this user. However, if multi-factor authentication was omitted and authentication qualifying only to `good-entropy` was carried out during the session, the CSP MUST assert Cappuccino and MUST NOT assert Espresso.

5. Representation on federated protocols

This section specifies how the values presented in the previous section shall be represented using federated identity protocols.

5.1. Security Assertion Markup Language 2.0 (SAML)

The table below presents how this assurance profile is represented using the SAML framework. Following presentations are used:

- **eduPersonAssurance** attribute, as defined in [eduPerson].
- **AuthenticationContextClassRef**, as defined in section 2.7.2.2. of [SAML Core].
- **SAML2 metadata entity attributes**, using the EntityAttribute name "urn:oasis:names:tc:SAML:attribute:assurance-certification" [TO BE DONE]

Value	eduPersonAssurance	Authentication ContextClassRef	SAML2 Metadata entity attribute
`$PREFIX$`			X
`$PREFIX$/ID/unique`	X		
`$PREFIX$/ID/no-eppn-reassign`	X		
`$PREFIX$/ID/eppn-reassign-1y`	X		
`$PREFIX$/IAP/local-enterprise`	X		
`$PREFIX$/IAP/assumed`	X		
`$PREFIX$/IAP/verified`	X		
`$PREFIX$/AAP/good-entropy`		X	
`$PREFIX$/AAP/multi-factor`		X	
`$PREFIX$/ATP/ePA-1m`	X		
`$PREFIX$/AP/cappuccino`	X		X
`$PREFIX$/AP/espresso`	X		X

The CSPs are expected to populate the `$PREFIX/AP/cappuccino` and `$PREFIX/AP/espresso` metadata entity attributes if they are capable of fulfilling those profiles at least for a subset of their users. The Relying Parties can make use of that information to manage their list of CSPs who can provide assurance that meets their requirements.

The CSP MUST present the values a particular authenticated user qualifies to in an assertion which the Relying Parties are advised to observe.

6. References

eduPerson	Internet2/MACE. eduPerson Object Class Specification (201602). http://software.internet2.edu/eduperson/internet2-mace-dir-eduperson-201602.html
eIDAS LoA	European Commission. Commission Implementing Regulation (EU) 2015/1502 of 8 September 2015 on setting out minimum technical specifications and procedures for assurance levels for electronic identification means. http://eur-lex.europa.eu/legal-content/EN/TXT/?uri=OJ:JOL_2015_235_R_0002
ePSA Comparison	Cormack, A., Linden, M. REFEDs ePSA usage comparison, version 0.13. https://blog.refeds.org/wp-content/uploads/2015/05/ePSAcomparison_0_13.pdf
IGTF	Groep, D (editor). IGTF Levels of Authentication Assurance, version 1.0. https://www.igtf.net/ap/authn-assurance/
Kantara SAC	Kantara Initiative. Kantara Identity Assurance Framework. Kantara IAF-1400 Service Assessment Criteria v5.0. https://kantarainitiative.org/confluence/display/LC/Identity+Assurance+Framework
RFC2119	Bradner, S. Key words for use in RFCs to Indicate Requirement Levels. RFC2119. https://www.ietf.org/rfc/rfc2119.txt
SAML Core	Cantor, S., Kemp, K., Philpott, R., Maler, E (editors). Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) V2.0. OASIS Standard. http://docs.oasis-open.org/security/saml/v2.0/saml-core-2.0-os.pdf
X.1254	International Telecommunication Union. Series X. Data Networks, Open System Communication and Security. Cyberspace security – Identity management. Entity authentication assurance framework. Standard X.1254.https://www.itu.int/rec/T-REC-X.1254

References

[MNA3.1] https://aarc-project.eu/wp-content/uploads/2015/11/MNA31-Minimum-LoA-level.pdf

[BPA] AARC Blueprint Architecture, https://aarc-project.eu/blueprint-architecture/

[ISO29115] ISO 29115:2013 *Information technology -- Security techniques -- Entity authentication assurance framework*; https://www.iso.org/standard/45138.html

[KIAF] Kantara Initiative, *Kantara Identity Assurance Framework*; https://kantarainitiative.org/confluence/display/LC/Identity+Assurance+Framework

[SP80063] NIST SP800-63 *Electronic Authentication Guideline*; http://dx.doi.org/10.6028/NIST.SP.800-63-2; http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-63-2.pdf

[SP80063V3] NIST Draft Publication Digital Identity Guidelines; https://pages.nist.gov/800-63-3/sp800-63-3.html

[EIDAS] Digital Single Market, Digital Economy & Society: *Trust Services and eID*; https://ec.europa.eu/digital-single-market/en/trust-services-and-eid

[INCAP] *InCommon Assurance Program*; https://www.incommon.org/assurance/

[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

[VOT] IETF Vectors of Trust BoF; https://www.ietf.org/mailman/listinfo/vot