REFEDS assurance framework public consultation results June 2017

Consultation document: https://wiki.refeds.org/download/attachments/22544423/REFEDS-Assurance-Frameworkv1.0.pdf

Type:

- Content (C): The comment proposes changes to the profile
- Presentation (P): The comment proposes clarifications or reformatting of the text

Number	Type	Line / Reference	Proposed Change or Query	Proposer	Action / Decision by the working group
25	P	Section 1	My comment on the document itself is that it is difficult to understand what it is for at a high level. It would be helpful to have a purpose statement at the beginning.	David Dykstra, Fermilab	Added one more paragraph to the abstract.
12	P	Section 1, Line 44	Proofreading s/has/have/	Andrew Cormack	Adopted.
16	P	Section 2, Line 51	"authenticated sessions can qualify to different values." This possibility is not expanded on in the SAML representation section, only the entity level assertions are included.	Hannah Short	This comment became obsolete when the working group decided to drop the assurance of the authenticated session from the framework.
2	C	Section 2.1	There is active discussion around the major shortcomings of SAML persistent IDs because of their case-sensitivity requirements, so I worry this is going to be obsolete on day one. Perhaps the reference to specific identifiers can be left to implementation guidance and not baked in.	Scott Cantor	Amended the RAF document to make ePUID a preference. Made the text and footnote more abstract for the directed identifiers. Need to clarify

					them in accompanying non-normative material.
3	C	Section 2.1	In light of recent concerns raised about the relationship between EPPN and mail attributes, it may be wise to be explicit about any assumptions that one is intended to make (or not) about that. For example, if it's meant to be assumed that a given EPPN value if found in a mail attribute refers to the same person, that isn't mandated by eduPerson, and should be called out here.	Scott Cantor	Added to section 2.1: "Finally, the reader is reminded that they should not assume any uniqueness property that goes beyond the specification of the attribute. For instance, a Relying Party should not assume that the holder of an ePPN value is the receiver of an email message sent using the ePPN value as the receiver address."
23	C	Section 2.1, Line 55	Requirements on ID traceability need to be further specified, so that institutions can set up their policies properly: e.g., what happens if an institution goes out of business - with and without business continuity?	Petr Holub, BBMRI-ERIC	The proposal was not adopted. The RAF initial profiles are responsive to prior research infrastructure survey.
14	P	Section 2.1, Line 57-ish	Maybe worth stating explicitly that none of the following applies to the identifiers actually listed under "unique"? "In addition to the three identifiers mentioned in the definition of "unique", within the REFEDS community..." and "... two additional values that a CSP declaring uniqueness can use to indicate the extent to which this applies to ePPN"	Andrew Cormack	Adopted.

17	C	Section 2.1, Line 65 - 78	Should this be linked somehow to the Research and Scholarship entity category? R&S makes almost the same requirements on ePPN. Likewise, the fact that no ePPN reassignment properties are included in the cappuccino and espresso profiles imply that ePPN should not be used (to my mind). I can imagine this being confusing when it comes to adoption by entities already using R&S.	Hannah Short	There is no conflict between R&S and RAF; both approaches can be deployed in parallel (see table: https://docs.google.com/document/d/1MKZuKlDBnSggM4gPYY_CpuCQ6_LaVPfqE-Bqe9NVeGA/edit). Clarified that in a footnote.
13	P	Section 2.1, Line 88	Or the SP could invoke some strong out-of-band process to verify whether the user **is** still the same? I'm thinking of courses that involve a student having a year in industry/another country, or academics taking sabbaticals. Either of those could result in a year of non-use of a service. Or, indeed, there might be some services (e.g. enrollment) that are naturally only used once a year!	Andrew Cormack	Added an extra sentence to section 2.1.
1	C	Section 2.2	What about the other IGTF APs? Presumably CEDAR is also here if CEDAR>=BIRCH. What about ASPEN? (suggestion for Appendix)	Jens Jensen	Added CEDAR to the list.
24	C	Section 2.2	There might be even higher level of ID vetting/proofing LoA: when you validate the presented IDs with the government, and then even higher when you capture biometric information from the user as a part of the registration process. The latter one is not very likely for life sciences, but the earlier might be needed to avoid problems with stolen/counterfeit IDs	Petr Holub, BBMRI-ERIC	The proposal was not adopted. RAF initial profiles are responsive to prior research CI survey. New profiles can be added to this extensible framework when it becomes valuable to do so.
18	P	Section 2.2, Lines 110 - 114	The wording here is potentially misleading. Does "assert" mean comply with, or broadcast compliance with (e.g. in SAML metadata)? If they are incremental, is it really necessary to broadcast both?	Hannah Short	Added “assert and comply with”.

22	C	Section 2.3	Authentication LoA: level for X.509 certificates - they are not necessarily multi-factor, depending on how they are implemented practically (token stored certs vs. browser-stored certs) **Mischa Salle:** Hi Petr, token stored certs and browser-stored certs still both have an extra password (pin for token vs. keystore or browser master password). Only hostcerts typically don't. Are you thinking about a specific scenario?	Petr Holub, BBMRI-ERIC	MFA is defined in the REFEDS MFA profile. It is being considered to extend REFEDS good-entropy-passwords to good-single-factor
20	C	Section 2.4	Other attributes LoA: project affiliation from institutions - for projects that are given to the institution 1. when people are kicked out of the projects and still retain their institutional affiliation, we have no way to figure this out on the infrastructure level then their good will to report it back to us) 2. for projects given directly to the people (like ERC grants) this is not needed	Petr Holub, BBMRI-ERIC	The proposal was not adopted. RAF initial profiles are responsive to prior research CI survey. Freshness of other attributes than ePAffiliation can be addressed in a separate profile.
26	C	Section 2.4	Making ePAffiliation freshness a requirement makes no sense for the IdPs that have adopted a policy that they never release ePAffiliation	Pål Axelsson	Changed the wording to reflect the idea that the ePAffiliation needs to be 1m fresh only if it is released.
10	C	Section 2.4	Attribute freshness is also determined by the operational security of the IdP, which can have an impact on the LoA of the asserted attributes. We were wondering whether this should be added as a 5th point determining the overall LoA (i.e. as a 2.5). Section 3 point 4 mentions that "Generally-accepted security practices are applied to the Identity Provider." but this is too vague and still does not take operational security into account in the determination of the overall	Mischa Sallé & the rest of AARC-JRA1	The proposal was not adopted. RAF and SIRTFI are not in conflict and an RP can require them independently. Section 3 is borrowed from

			LoA. We feel that this could be made explicit by e.g. requiring SIRTFI for this baseline LoA.		InCommon's baseline expectations. Making higher requirements than InCommon would break the RAF's compatibility with them.
19	C	Section 3, Line 164	Given the lack of other "Generally-accepted security practices" in this domain, could we tie the framework to Sirtfi? +1 to Mischa's comment above	Hannah Short	See above
4	P	Section 3, Line 166	s/securitybcontacts/security contacts/	Scott Cantor	Adopted
9	C	Section 3.2	"The Identity Provider is trusted enough to be used to access the organization's own systems." This cannot never be fulfilled by some organisations in Italy because they don't use SAML internally, so their IdP is used only to access external systems. At the same time all the identities managed by the IdP are trusted at the same level of the identities used to access internal systems.	Lalla Mantovani	Reformulated: "...is trusted enough that it is (or it could be) used to access... " "Is trusted enough to be used" does not mean that they actually must be used for accessing internal systems.
5	C	Section 4	I'm really surprised that there's still no profile defined for "strong authentication, no verification" given the very consistent feedback I think IdPs have always gotten that that's the dominant requirement of a lot of projects. I think, as with InCommon Silver before, the verification requirements will prove impractical for most IdPs to meet whereas MFA is becoming very common, at least in the US.	Scott Cantor	That's mostly covered by the REFEDS MFA alone and a new coffee drink adds negligible additional value.
21	C	Section 4	Another profile is common: verified identity (f2f) + passwords	Petr	The proposal was not

				Holub, BBMRI-ERIC	adopted. That can be covered by an RP observing Cappuccino + verified
11	P	Section 4, etc	The profile name "espresso" duplicates the name NISO ESPRESSO report, which is directly mentioned in the REFEDS discovery guide. This runs the risk of confusing service providers who might be referred to both the ESPRESSO report for discovery best practices and the espresso profile for assurance, and will be looking for both of these on URLs within refeds.org namespace. I know profile names were voted on (and indeed, I stupidly voted for espresso). However, I believe avoiding this confusion is a compelling reason to choose another coffee beverage.	Guy Halse	The working group believes the risk of confusion is low because NISO ESPRESSO is a report and the Espresso profile is a specification. If confusion arises the Espresso profile can be called REFEDS Espresso for disambiguation.
15	C	Section 4, Lines 177 & 186	It bothers me that line 177 implies that Espresso is a superset of Cappuccino, but the table in 186 says they aren't, The implication is that anything qualifying as "mfa" also satisfies "good-entropy", but I'd much prefer that (if it's true, which depends on what replaces the placeholders s2.3) to be handled explicitly by making good-entropy a requirement of Espresso as well (in the same way as you've done for the ranked IAP values)	Andrew Cormack	Rejected the proposal but clarified the table. A SAML authentication response can have only one Authentication context value. If SAML2 Authentication context is used for asserting MFA and good-entropy, the Espresso profile cannot require them both to be present at the same time.
6	C	Section 5.1	The overall assurance profiles really need to be handled via AuthnContextClassRef if you intend for RPs to be able to request their use.	Scott Cantor	Requesting/asserting coffee drinks via AuthnContextClassRef was rejected to avoid IdP configuration complexity.

					Only authentication components are requested using AuthnContextClassRef. The RP cannot request the others, including cappuccino and espresso
7	C	Section 5.1	The text around the use of the EntityAttribute(s) needs to clarify whether there's an intended "authorization" semantic such that CSPs are not meant to assert the profile(s) if they don't carry the attribute. That tends to be inferred by people and that's one reason the InCommon MFA profile elected not to suggest that approach. If the profile is self-asserted by the CSP or at least if this document doesn't purport to suggest otherwise, I would reconsider that mechanism.	Scott Cantor	To foster adoption, pulled back the proposal to use SAML2 metadata Entity Attributes in the framework.
8	C	Section 5.1	The inclusion of SAML metadata entity attributes in the framework makes me nervous that the entire framework will stall due to unresolved questions around handling of metadata entry attributes by federation operators. However, as far as I can tell, the metadata entity attributes are only a hint for IdP discovery, and I understand the recommended practice is to list all eduGAIN IdPs in our discovery interface then gracefully handle errors rather than preemptively restricting the list of IdPs to only those with certain entity attributes.	Jim Basney	See above.