2022-06-08 R&S 2.0 Notes

Attendees

- Andrew Morgan
- Pål Axelsson
- Heather Flanagan
- Christos Kanellopoulos
- Björn Mattsson
- Jií Pavlík
- Martin Stanislav
- Nicole Harris
- Alan Buxey
- Jens Jensen - STFC UKRI

Pre-reading

- Anonymous Access Draft
- Pseudonymous Access Draft
- Personalized Access Draft
- Federated Authorization Best Practices

WG Consensus

- The Anonymous Access, Pseudonymous Access, and Personalized Access Entity Categories shall be harmonized based on the decisions made around Personalized Access.
- Authorization guidance shall be split out into a separate, descriptive paper and not be part of any of the entity categories.
- The names should be "Access Entity Category" not "Authorization Entity Category" - 10 January 2022
- We will not include assurance requirements to the Anonymous Access Entity Category - 10 January 2022
- We will take out wording in Anonymous that Section 4 that requires proof while leaving in wording that requires documentation for Registration Requirements - 24 January 2022
- We will remove the technical requirements for SAML2 and metadata refresh - 7 April 2022
- Pseudonymous is done (modulo any changes identified as we work through personalized) - 20 April 2022
- Federations should allow SPs to request multiple ECs - 4 May 2022

Agenda

- Continue discussion re: requiring a graceful failover from Personalized to Pseudonymous, and consider what that means for the data minimization statement
    - Homework - group to consider what and where text should go in Personalized, and should consider how to clarify Pseudonymous to help reduce confusion and ambiguity.
- Review draft slides for REFEDS 44
- Review initial draft for authorization - Federated Authorization Best Practices
    - "I think it's important that a service that requires only the former but can do the latter be able to assert both. We should take care to author the changes to both of them to ensure that's sensible. It shouldn't worded so strictly that you have to pick only one."

Notes

- If the SP should only request Personalized, and if the IdP only supports Pseudonymous, the IdP will respond with Pseudonymous even if the request is for Personalized. The SP will present a different level of service if it can't get Personalized, thus resolving the concern that the SP cannot operate with less than Personalized. It can, it just won't operate to the same level without additional action directly on the part of the user.
- This won't work as well to fall back all the way to Anonymous, but there is a narrow use case where it might work depending on what the SP can handle.
- If the IdP sets Personalized, it has to send all the attributes (like R&S).
- Is it ok for an SP to both requests Personalized and Pseudonymous, so they can use Pseudonymous as a fall back? No. If an SP only has Personalized in their metadata, when they send an authN request to an IdP that supports Pseudonymous, the IdP would understand to release Pseudonymous.
    - This is similar to authN context class - you release the best available.
    - We do not want an SP to assert multiple categories. We discussed that GDPR would require the IDP to only release the lesser attribute set to an SP that asserts multiple categories.
- ECs cannot cross-reference; we cannot build in dependencies.
- From a technology perspective, this is simple. From a rule of data minimization, this is very problematic. You cannot send more than the least you can send, so you would only ever send Pseudonymous.
    - Pseudo-config in Shibboleth IDP:
        - If entityCategory == Personalized or entityCategory == Pseudonymous { release Pseudonymous attribute set }
    - If an IdP supports Personalized, I will release all the attributes. If I have Pseudonymous, I only release those. If I don't want to support Personalized, then I should release less, which means I still need to configure Personalized in my IdP. Either I configure two ECs and follow the rules, or only configure one and Personalized will never work.
    - The SP does require everything, but is trying to allow for a fallback flow.

- Given the intersection here of legality, policy, and technology, are we really going outside the role of the EC?
- By adding the implementation text to Pseudonymous, one thought is that we make this more useful. When we discuss GDPR, we're approaching this from a taboo perspective and will cause us to stall entirely. Either we see the whole picture of the federation ecosystem or we only work with a minimal slice.
- We're trying to solve a problem of not releasing attributes. Is there a different way of solving the fallback issue? Is this out of scope for the EC? The EC should subscribe cleanly to what it can do, not what to do if the EC isn't supported. The SP can't assert both because then they wouldn't know what they would get. Some SPs will filter if the ECs don't line up between the SP and IdP, and some won't.
    - If the SP isn't filtering on what ECs the IdPs support, they are going to have a higher rate of error presented to the users.
- Next Steps
    - Heather to summarize the discussion and options at REFEDS 44