IDESG Authenticate Person Use Case Annotated to Highlight Attribute and Identifier Issues

Status: Proposed This Use Case is under development by members of the use cases ad hoc group.
Title: Authenticate Person
Use Case Description: A Claimant browses to a website which requires authentication. The web site provides the Claimant the ability to authenticate their identity using an Identity Service Provider of the Claimant's own choice through the use of privacy enabling and standards based protocols.
Use Case Category: Authentication Related
Contributor: Standards Committee Use Cases Ad Hoc Group

Use Case Details

Actors:

- Claimant – wants to obtain access to a web resource
- Identity Service Provider – performs primary authentication of the Claimant's credentials
- Relying Party – wants to have some level of assurance about the identity of the Claimant
- User Agent – enables communication between the Claimant, Identity Service Provider, and Relying Party

Goals / User Stories: The Claimant is able to gain authenticated access to the Relying Party web site without having to provide the Relying Party with a primary credential. The Claimant is able to perform primary authentication with an Identity Service Provider of their choice and the manner in which the Identity Service Provider is identified is an intuitive process. An example of an intuitive process (but not a requirement) would be to identify the Identity Service Provider via the Claimant's email address. If the Claimant has previously established a trusted relationship with the Identity Service Provider then a session management design should enable the authentication to take place without requiring an additional prompt for the primary credential.

    - *NOTE: This annotation describes how to manage attributes and preserve privacy in the interaction that the user has, via their ISP, with the relying party RP. That RP may subsequently put up a screen that asks the user for other attributes, potentially degrading privacy. These other attributes are now being self-asserted,*

but many RP's have other mechanisms for determining the validity of the attributes.

Assumptions: It is assumed that the Claimant has already been identity proofed to some LOA and has already received credentials binding their identity to one or more tokens.

Requirements:

- The Claimant must be capable of selecting an Identity Service Provider of their choice (provided that the Identity Service Provider meets the LOA requirements of the Relying Party)
    - *NOTE: This is known as the "discovery problem". In general if a claimant first browses to a federated Relying Party (RP), the RP has no way to know which Identity Service Provider (IdP) the claimant may wish to use for this online interaction. Common approaches are either providing a type-ahead drop-down list of choices based on the RP's Identity Federation metadata or having the RP "guess" the IdP endpoint by asking the claimant for an email address and then looking at the portion of the email address to the right of the "@" sign. Clearly the latter approach is fragile. Other approaches under development include the work being done under the heading "Account Chooser" (see, e.g., http://accountchooser.net/).*
- The Identity Service Provider must present the Claimant with privacy protection choices that minimally include the ability to not disclose their true identity (e.g. use a pseudonym)
    - *NOTE: For this to work in practice, both the IdP and the RP must support it and be configured accordingly. In some federations and with some software implementations, it is a straightforward matter of configuration to specify which identifiers will be released by an IdP to a given RP or class of RPs. In general it will not work to have the IdP side alone provide the claimant with this choice if the RP is not prepared to handle both identity-revealing identifiers and pseudonymous identifiers and to recognize which is being presented during a given user interaction. If this is not achieved by configuration it must be arranged ahead of time through out of band means.*
    - *NOTE: There is a complicated set of issues around different types of identifiers, their properties and their fitness for purpose in various scenarios. See, for example, the Identifier document at https://spaces.internet2.edu/download/attachments/38670741/identity,+identifiers+and+attributes.pdf?version=1&modificationDate=1378931647484*
    - *Advice for RPs*
        - *DO communicate with the IdPs of interest to understand the characteristics of the identifiers they are prepared to present to you.*
        - *DON'T assume that an IdP-provided identifier obtained at one point in time will always refer to a valid entity (it may be impermanent)*
        - *DON'T assume that a particular value for an IdP-provided identifier will always point to the same entity (it may be reassignable)*

        - *DO be prepared to handle the fact that a given person's identifier with one RP may be different from their identifier with a different RP (it may be "directed" and non-correlatable)*
        - *DO be prepared to accept identifiers that may fall at different places on the continuum from identity revealing (such as an email address) to privacy protecting (such as a random bit string or UUID)*
        - *DON'T base your operational model on the notion that an email address can double as a unique, permanent user identifier*
- The Identity Service Provider must present the Claimant with an option to not track the Relying Party
    - Advice for IdPs
        - *DO communicate with the RPs of interest to understand their operational assumptions about the characteristics of the identifiers you send them*
        - *DO expect that many RPs will treat any identifier value they receive as permanently valid and non-reassignable*
        - *DO expect that many RPs will assume that a scoped identifier that looks like an email address will be a deliverable email address*
        - *DO be prepared for an RP to treat any user email address you provide them as a permanent unique account identifier for that user*
        - *DO maintain identifiers of multiple types for each person or other entity so that you can support a variety of claimant and RP requirements around such issues as privacy, human convenience, and non-correlatability*
        - *DO expect that regardless of the privacy protections you set up as an IdP, some RPs will immediately ask newly authenticated users to provide all kinds of self-asserted attributes and identifiers on a "profile" form.*

Process Flow:

1. The URL is loaded and the site requires an authenticated identity in order to proceed.
2. The Claimant is able to intuitively indicate to the Relying Party their preferred Identity Service Provider
3. The Relying Party directs the Claimant to their Identity Service Provider, likely via their web browser User Agent (UA)
4. The Identity Service Provider authenticates the Claimant. The Identity Service Provider may accomplish this authentication either via performing primary authentication of the Claimant, or via the Claimant's possession of a bearer token showing that authentication has already taken place, e.g. via the presentation of a session or persistent cookie. Note that some Relying Parties might have differing requirements that dictate whether or not cookies (or other session tokens) may be used, and if so what lifetime is acceptable before the Relying Party requires the Identity Service Provider to perform Primary Authentication of the Claimant. Such requirements should be indicated by the Relying Party in its request to the Identity Service Provider when asking for an Identity Token for the Claimant.
5. Upon successful authentication of the Claimant, the Identity Service Provider generates an Identity Token which contains the claimed identity of the Claimant and possibly other personally identifiable information. The Claimant must be able to indicate what

personally identifiable information is included in the Identity Token, including the usage of real name vs. pseudonym or other personally identifiable information such as email address, street address, or birthday. The token is signed by the Identity Service Provider and possibly encrypted with a key known to the Relying Party. The token is sent back to the Relying Party via the Claimant's web browser UA.

6. The Relying Party validates the signature of the Identity Service Provider and extracts the claimed identity and possibly other personally identifiable information. The Relying Party may optionally query a third party attribute provider for additional attributes bound to the claimed identity or may map the claimed identity to local attributes.
7. The Relying Party makes authorization decisions based on the claimed identity, attributes of the identity, or both and returns resource (as applicable) to the Claimant's web browser.

Success Scenario: The Relying Party returns the requested resource to the Claimant's web browser

Error Conditions:

- Relying Party cannot validate assertion
- Identity Service Provider cannot authenticate the Claimant
- The Relying Party rejects the LOA of the Identity Token
- The Relying Party is unable to authorize the Claimant, even after validating the claimed identity
- The Claimant is not authorized to access the requested application, resource or service

Relationships

Extended by: Authenticate Using Pseudonymous Identity Use Case

References and Citations

- NIST_SP_800-63-1