Attribute Release Process

The default privacy stance of most deployments, and indeed most organizational (as opposed to consumer) IdPs, is one of minimal, if any, information released to SPs for which special arrangements have not been made. This creates a serious barrier to adoption for SPs that need only basic/common identity attributes, but are not designed around heavy-weight contractual vehicles.

It can be argued (and has been argued by some) that these are not appropriate uses of organizational IdPs, and that contracts are necessary to ensure appropriate data use and regulatory compliance. Such IdPs will simply not be usable for the vast majority of "long tail" services expected to emerge.

For those that do choose to support these use cases, there are a number of options:

1. a model based on user consent
2. a relaxed stance to privacy control for some subset of users
3. service categories

No matter what path the IdP chooses, a documented policy and process must exist around the release of attributes such that SPs can engage with the IdP in as efficient a manner as possible, and in particular so that privacy-related failures can be reported to users in a useful fashion. The user should not bear the burden of understanding the technical details involved.

The `<PrivacyStatementURL>` element, an important IdP user interface element in metadata, is a useful place to capture, if not directly, then indirectly on a broader privacy policy page, a link to this information.

The "administrative" contact in metadata is an appropriate way to designate the point of contact for users in referring privacy issues to their IdP for resolution (and for the appropriate staff to ask questions of the SP).

For any these alternatives to work well, SPs need to articulate their Requested Attributes to IdP partners, preferably in metadata. Thus there is overlap between the kinds of SPs for which such definition is possible and useful, and those to which the release of attributes must be streamlined.

Recommended Practice

- IdPs make common identity attributes (identifiers, displayName, mail) available to educationally-useful/non-commercial SPs for significant user populations, either subject to opt-in user consent, or with an opt-out process.
- IdPs document and publish their policies and procedures for the release of attributes. The `<PrivacyStatementURL>` element should link directly or indirectly to this information.
- An "administrative" contact in metadata, is documented for each IdP and SP identifying a point of contact for attribute release issues.

Opt-In

The *goal* is to facilitate the release of basic identity information without the involvement of a manual workflow at the IdP. The most restrictive mechanism for this is to require users to opt-in to the release of data, usually at the time the service is used.

This is supported in software through built-in or add-on modules for "consent management." Such modules are necessarily a significant part of the effort involved in deploying an IdP, but the flexibility is extensive in terms of recognizing changes to specific attributes and values, tracking consent across client devices, etc.

A less invasive addition might be a simple page that collects a generic "opt-in" to the release of an abstracted set of attributes (e.g., "your name and email address") to specific SPs or a generic class of them, tracked via cookies. This may be suitable for some IdPs at significantly less cost.

A major advantage to a consent-based model is that universal coverage (even for FERPA-suppressed students) is potentially possible.

Opt-Out

If consent is deemed too onerous or undesirable, an alternative is to identify classes of users for whom default release of basic information is acceptable to SPs that meet a commonly-accepted set of criteria (such criteria definitely a work in progress at this stage of discussion).

If the classes of users are broad enough to be useful (e.g., faculty and staff, perhaps even non-suppressed students), then the utility of the IdP may be high enough to realize significant benefits.

Of course, an opt-out mechanism of some kind is an essential component of such an approach.

Service Categories

Another approach is to define a service category that IdPs can choose to support or not. A specific example is the Research & Scholarship Category of services, which currently is gaining traction in the international R&E community.

To operationalize a service category, entity attributes are used in IdP attribute release policy configurations in lieu of Entity IDs. The advantage of doing so can not be overstated.