

How Better Attribute Management Helps Federation

This is a position statement from the InCommon Technical Advisory Committee.

Abstract: New methods of managing attributes promise to make federation easier to use and to operate. The key elements are: publishing of attribute requirements, support for user consent, and common attribute policies. Software and services that provide these features are becoming available, but will require InCommon participants to align their policy and technology deployments to actually realize the potential benefits across the federation.

Attributes: the State of the Art

This note describes some new approaches to managing user attributes in InCommon. These approaches promise to make federation easier and more scalable for end users, system owners, and policy makers. Our intent with this note is to introduce the InCommon community to these methods and their potential, and to encourage participation in making this happen.

InCommon makes it easier for its participants to federate securely with many partners by being a trusted registry of site data (metadata). Today this data includes site names and locations, chosen methods for transactions, public keys, and contact information. Successful user login requires more than this, however. The correct user attributes must also be provided so that services (SPs) can make informed access decisions. This is true even for simple SPs that only require the equivalent of a userid. Today, configuring an IdP to release the desired user attributes to an SP is almost always a manual process; far too often it is also mysterious, error-prone, labor-intensive, and user-unfriendly. In some cases the process starts with a bad user experience when a user tries to access an SP and the needed attributes aren't provided. The IdP admin is expected to match requested attributes with the ones that are actually released by the IdP, all the while dealing with a trouble ticket and an impatient user.

Once an IdP and SP decide to work together, typically the IdP and SP administrators must communicate to work out the details. It helps if the SP organization has a document describing its requirements, but it may not. Even if it does, the SP may not state requirements precisely, preferring to be flexible to accommodate varying IdP capabilities. The IdP administrator at some point must apply configuration changes specific to the SP, changes that reflect both SP requirements and IdP organization policy. For example, some IdP organizations may be content with a loose, simple policy ("release all attributes for all users to this SP") while others may be more strict ("release only required attributes for designated users to this SP"). Policies might be based on characteristics of SPs (for example, "loose policy OK for higher-ed SPs but not commercial ones"), but even so these SP characteristics have to be discovered and applied manually. Attribute release decisions might also be delegated to users in some cases ("release attributes if the user says it's okay"), but the means to ask users is not commonly deployed by IdPs in InCommon today.

All of this creates a substantial barrier and delay in getting from "it would be interesting to federate" to "users are logging in." Sometimes a thoughtful, lengthy process is necessary if an SP has valuable or sensitive resources, or its access policy is complex and customized, but in many cases policy setup could be largely or completely automated if the right information and infrastructure were available. In identity management we see many instances where a combination of technology and policy is needed for methods to scale (e.g., role-based access management). Attribute release is another such instance. It is particularly important to lower the burden of federation startup for those many sites for which the effort and delay of manual attribute setup can't be justified.

The rest of this note describes a number of infrastructure elements InCommon intends to introduce and encourage participants to adopt, and suggests some of the discussions and agreements federation participants will need to engage in.

User consent and default policy

We believe that mechanisms for user consent are important in improving attribute management. The basic notion of user consent is simple: at the time the user is logging in via an IdP, before attributes are sent to the SP, the user is asked, via a web form, to approve the release of his or her information to the SP. If the user approves, login continues as normal and approved attributes are sent to the SP for access to the remote service. Having user consent available at the IdP helps the IdP administrator by supporting a default policy that can apply to many SPs. This policy is: "for attributes in a standard set, release what the SP asks for, if the user agrees." This policy clearly does not satisfy all needs, for example SPs with attributes outside the standard set. But we believe that for many IdPs this may be a reasonable default policy, that is, it can be applied to all SPs in the federation other than explicitly-managed exceptions. For all of those SPs, the burden of attribute release management at the IdP is completely removed. If the right information is available, this can reduce the "time to successful login" to zero, and eliminate (for many SPs) the bad user experience caused when the correct attributes are not provided.

In practice user consent has many issues that make it complex. For example, an SP might like to indicate that some attributes are required and others optional, so the user could decide not to provide the optional ones. Some attributes are inherently hard to understand (e.g., entitlements), raising significant usability questions. Users should be able to save and edit their consent decisions. SP attribute requirements may change from time to time and cause policy changes at the IdP.

There are of course many legal and policy issues associated with the release of personal information, for example, whether a site is "necessary to use for a business or academic purpose." Consent for release of information in a federation context may need to be integrated with existing IdP organization support for users controlling publication of their personal information (e.g., FERPA controls). Some institutions may think that user consent is at odds with their obligation to protect the privacy of institutional data. Many have concerns that users may not really be understanding "consent" and will just click on whatever is displayed in order to move on; hence consent is more like coercion from one point of view.

User consent at the IdP is not a panacea, but we do believe it is an important tool in the toolkit.

The vision, in practice

We seek to enable the following scenario.

In addition to the information registered in metadata now, an SP admin registers information about the attributes their site requires (or would like to obtain). In addition, a user-meaningful name for the SP (e.g., "Bob's Excellent Article Database") is registered so that users don't have to decipher URLs when deciding whether or not to release their information (in fact, a major purpose of identifying SPs using metadata is to associate a site's name and description with its various and changing URLs). SPs might also register other meta-characteristics (e.g., "operated by university", "test site") that could be relevant to policy decisions. This information is published in InCommon metadata alongside other SP metadata that is published today.

IdP admins install enhanced IdP software that consumes the SP's requested attributes and related elements from Federation metadata, making this new information available to the software's attribute release mechanisms. More importantly, IdP institutions develop attribute release policies that take advantage of the new user-consent mechanisms installed in their IdPs.

Now, when a user accesses an SP for the first time, the attribute and consent machinery at the IdP comes into play. If the IdP default policy is satisfied, the software determines the requested attributes and asks the user to consent to their release, storing the decision for later use. If the user consents, attributes are released and the login succeeds (subject to SP access controls). If release to this SP requires a human policy decision, this event can start a workflow to alert the policy decision-makers.

Steps to making it happen

This vision is appealing but it won't happen all at once. As with most innovations within the Federation, multiple stakeholders will contribute to the success of this effort.

The first step is to provide support for requested attributes in InCommon metadata. An initial implementation of an administrative user interface was debuted at the Spring 2011 Internet2 Member Meeting. Production deployment of this interface will commence in June 2011 and ramp up as we approach the 2011–2012 academic year.

The second step is to provide support for user consent in the IdP software. SWITCH, the Swiss federation, has led the way by developing a Shibboleth add-on called [uApprove](#) that is widely deployed within the SWITCH federation. InCommon IdPs can use uApprove today; a few [InCommon early-adopter IdPs](#) have already done so.

The Shibboleth team is working to integrate uApprove more closely into [Shibboleth IdP v3.0](#), which is scheduled to be released in 2012. When this happens we will encourage all Shibboleth-using InCommon IdPs to deploy IdP v3.0, subject to organizational policies and procedures of course. In the mean time, participants should seriously consider whether to deploy uApprove as an add-on.

While different IdPs will have different policies regarding attribute management and release, it makes sense for InCommon participants to develop consensus on standard policy options. We hope federation participants, both IdPs and SPs, will be interested in contributing to these discussions as strive make this vision a reality within the Federation.