Hybrids of Social and SAML

Perspectives on Handling Both Social and SAML Identities

- [Keith Hazelton, 20-June-2011]
- Relying party wants to make some federated resource accessible to people who have a Social Identity Provider (Twitter, Yahoo, Google, Windows Live)
    - On one hand they want to minimize changes to their Service Provider implementation and their application code
    - On the other, they want to know which ~~and what type of~~ IdP is handling any particular access instance
    - If they want identity attributes as well as an authentication assertion, they want those attributes to have consistent names and consistent value syntax. The same attribute name should not represent two different attributes nor should the value syntax vary for a given named attribute.
        - This applies whether a Social-to-SAML gateway is involved or not.

- The user expects to make a search-and-one-click selection of their IdP of choice
    - They would like to see a given social identity provider identified the same way regardless of their path to the SP
    - They might expect that they would be recognized as the same individual regardless of their choice of IdP, but in general this is not possible without some user mediated account linking on a per-SP basis

[Steve Carmody, 20-June-2011]

1) Over time, implementations of support for SAML and Social Identities will most certainly evolve and change. There are already several Social --> SAML gateway implementations in use. If the market ever identifies a "choice" of a preferred Social Protocol, we may see existing native SAML implementations extended to include native support for that protocol. (There is a consensus that OpenID (as used today) will not that be that choice; instead, people are watching openid-abc and oAuth v2.)

Because of the high degree of uncertainty in this space, a starting principal should be "The information provided to an Application when describing an authentication event should not be affected by the implementation of authentication support." (ie an application should see the same values being presented whether a gateway or native implementation is being used).

It was suggested that we should adopt an approach of defining what the infrastructure should deliver to the application; once there is agreement on that, then work backwards and define how the end-to-end and GW implementations would deliver this information.

2) As noted though, there is currently a high degree of uncertainty in this space, and initial implementations should probably be described as prototypes or experiments. All of the currently known implementations are gateways that map onto SAML. At one point, we thought that supporting all these protocols in software would be possible; but, they are all moving targets and have become proprietary.

That said, it is recognized that gateways are inevitably a bad compromise, and will be painful to disassemble. They distort all sorts of things (eg who is **really** speaking). Previous gateways with other protocols have all eventually foundered (eg email). The basic model of the internet is is end-to-end, and that continues to work well. So far, gateways have addressed problems in the short term, but have become problematic over time.

Understanding all that.. we understand why gateway implementations will be used in the short term.

However, we strongly recommend against a Federation standing up a temporary Social --> SAML gateway in the near term, because inevitably it will not be temporary. Instead, we expect to see a half-dozen campuses building gateways using the classic microsoft model (local gateway, mapping all the incoming protocols to what's supported locally). Hopefully, as a community we will learn from their experience; we will try to agree on standardization that will insulate applications from GW vs end-to-end architecture differences.

3) One of the reasons that SPs/Applications do not want to deal with multiple protocols is that application developers do not want to deal with the issues surrounding Discovery. They would rather hand this problem off to the gateway. While that approach greatly simplifies the problem for the developer, it also offloads part of the problem onto the user experience.

4) Applications do NOT want to deal with multiple protocols; they want the information about an authenticaiton event presented to them in a standard format, independent of the protocol that was used. That information may contain some information that is only relevant for some authentication mechanisms. However, teh "stadndard" information should all be presented in the same way, with the same labels and encoding.

5) The consensus is that it is extremely unlikely that the social identity providers would agree to issue SAML Assertions to SAML SPs.

The issue that has too many names: Invitation, Volunteering, Conscription and other ways of adding members to a CO

All collaborating organizations need collaborators. How do individuals get a CO membership? Process models proliferate to cover all the ways this needs to be done. We identify the basic modes and discuss their applicability, their strengths and some of their unresolved issues.

Invitation

Volunteering

Conscription