Conventions for representing Identity Providers in the Bamboo Trust Federation

Overview / Executive Summary

Project Bamboo's current (as of April 2013) IAM implementation is predicated on trusted client applications authenticating users, and asserting attributes about authenticated users to the Bamboo Services Platform when making service requests. Cf. the *Bamboo IAM from a client application's perspective* section of *Identity and Access Management - Authentication and Authorization* for both broader context and specific details.

One critical attribute of authenticated users is a reference to the Identity Provider (IdP) used to authenticate her/him. This information is asserted about a user by the client application in three distinct but interdependent contexts, enumerated below. In order for the IAM mechanism as a whole to function properly, all clients must represent IdPs to which users authenticate uniformly when associating user identities with login events.

Details about why these conventions are essential, and recommendation for standard representations to be used, are described on this page. It is worth noting that this convention is dependent upon proper (conventional) configuration of a Social/SAML gateway for social media identity providers (such as Google), as described on the page *Social/SAML Gateway to enable social media identity provision*.

Where are IdP identifiers used?

Identity providers authenticate users in the Bamboo Trust Federation. Trusted client applications are responsible for requiring and managing execution of authentication. In general, they do so by operating as a Shibboleth SP (Service Provider).

The identity provider used to authenticate a user in a given session is communicated or used in the following ways:

1. An IdP identifier (IdPID) is a component of the so-called ***SourcedId*** used as an input parameter to the Bamboo Person Service (cf. *Person* API).
    a. When constructing a SourcedId, clients are assumed to assert ONLY the IdP used to authenticate in the current session (the session within which a request to BSP is made)
    b. A client may assert any string as SourcedId.
    c. However, consistency is of paramount functional importance, such that an IdP is identified in the same way whenever a SourcedId is communicated to the BSP
    d. *Consistency across multiple clients* is necessary: this is what permits different client applications to match the correct BPId to a user no matter what application s/he uses to authenticate.
2. An IdP identifier may be asserted as an element of a ***scoped role*** passed in the X-Bamboo-Roles header of any request to the BSP made by a trusted client application
    a. The a **scoped role** format is assumed to be role@domain
    b. In the current implementation (which does not rely on N-tier AuthN), the only meaningful role that may be asserted with respect to an institution is affiliation, i.e., that *the authenticated user is affiliated with an institution represented by the domain part of the scoped role*; this affiliation is inferred from the user's ability to authenticate to the institution's IdP.
3. An IdP identifier may be asserted in a **Policy** that, for example, permits access to users who have any of a set of permitted *scoped roles* where the set of *scoped roles* represents institutional affiliations (e.g., *affiliates of Berkeley, Wisconsin, and Tufts may access resources governed by this policy*). Cf. *Authorization and Policy* for more detailed discussion.

Why is a convention needed?

1. **ScopedId values passed by *different* authenticating clients want to assert *the same* ScopedId value for any given user authenticating at a given IdP**. That is, Professor Jones authenticating to Acme University should be represented by the same SourcedId whether she has authenticated through Research Environment XYZ, Research Environment ABC, or the standalone *Account Services Module*.
2. **Affiliation asserted by a client in the form of a scoped role *must* match a scoped role allowed by a policy in order to successfully "pass" the policy's rule**. Clients must *either*
    a. know how affiliations are represented *differently* in the policies associated with various services to which the client directs request on an authenticated user's behalf (so that they can construct scoped roles appropriately for a given request); or,
    b. know how affiliations are represented *conventionally* in policies that apply to all services governed by BSP policies (so that they can assert affiliation in *each* request without needing to be aware of policies that may govern the resource being requested)
    c. b is easier for client developers than a!

Conventions for representing Identity Providers in the Bamboo Trust Federation

A user's fine-grained (detailed type of) affiliation is unknown in the currently-implemented IAM solution used by Bamboo, which does NOT rely on N-tier authentication and therefore *does not pass* fine-grained user attributes that may be asserted by an IdP.

A user's role at the institution running the IdP at which a user authenticated is therefore *undefined*.

Therefore, the recommendation is to use the convention undefined@domain to represent user affiliation inferred from the IdP used to authenticate in the current session.

Examples of user affiliation in scoped role format:

- undefined@berkeley.edu
- undefined@wisc.edu
- undefined@google.com

Examples of Identity Provider identifiers in IdPID format:

- http://berkeley.edu
- http://wisc.edu
- http://google.com

Client responsibilities for implementing the conventions

Client applications are responsible for extracting the conventionally-represented IdP identifier from attributes available to it following successful user authentication. As in the examples above, the appropriate string is expected to be the top-level domain and its immediate subdomain.

This string must be extracted by a client tool, application, or service from the Apache httpd environment variable *Shib-Identity-Provider* available to the client on successful user authentication. This environment variable is available to a client because such clients are required to act as a Shibbloleth Service Provider.

For example, a client would extract *wisc.edu* from the Shib-Identity-Provider value *https://logintest.wisc.edu/idp/shibboleth*.

As details of parsing and manipulating strings are particular to a given client's implementation language or platform, this documentation does not suggest how a client is to meet the responsibility for extracting the domain (such as *berkeley.edu*) and (depending on the context) either prepending *http://* or *undefined@* to construct the IdP identifier.