AttributeNaming

Attribute Names and Shibboleth

Shibboleth relies on consistent attribute naming to deliver information about browser users in a mutually understood way between the IdP and SP. Every attribute which has a definition and semantics that differ from others must have its own unique representation in a SAML attribute assertion to ensure that there are no misinterpretations or communication failures. This name must be expected and handled by relying parties. Values, vocabularies, and their meaning should be discussed as well, but are outside the scope of this document.

Much of the power of federated authentication is derived from the economies of scale accomplished by large numbers of providers speaking a lingua franca. Attributes are the language in which access control and release policies are written and are the piece of the infrastructure for which avoiding unnecessary proliferation of names is most important. Standards bodies have traditionally defined common attribute names and semantics(e.g. X.520, eduPerson, etc.) for LDAP and other information repositories. Some of these now define XML representations as well. Federations also can serve as locuses for attribute convergence.

The names for attributes in back-end data stores and consuming applications is decoupled from the expression of attributes on the wire. This allows for arbitrary local naming as long as the SAML expression is common. The mapping from data stores to SAML representations at the identity provider is performed using resolver.xml. These SAML representations are then made available to the web server and web applications in raw XML or through mappings performed using AttributeAcceptancePolicy files.

SAML Naming Conventions

The name of an attribute is expressed as `AttributeName="URI"` in a SAML 1.1 attribute assertion:

```
<Attribute AttributeName="urn:mace:dir:attribute-def:eduPersonScopedAffiliation"
     AttributeNamespace="urn:mace:shibboleth:1.0:attributeNamespace:uri">
    <AttributeValue Scope="supervillain.edu">Faculty</AttributeValue>
</Attribute>
```

While the Shibboleth software treats every attribute name as a string, it's recommended that URIs be used for attribute naming because of the uniqueness and namespace control they provide. The `AttributeNamespace` value shown above is just a convention used as a default in the Shibboleth software (and in various attribute profiles) to indicate that the entire attribute name is contained in the other field. Any attribute whose name is a URI is free to define itself using that namespace value.

The following steps should be followed when naming a new attribute:

1. Is this attribute standardized or defined by any organization which has already assigned it a unique identifier? If so, it should be used if at all possible.
2. If the attribute is defined through an LDAP object class, there is probably already an OID assigned. When possible, leverage the existing urn:oid namespace.
3. If no suitable name yet exists for this attribute, consider creating one preferably through constructing a proper URL, or instead using a delegated urn:mace namespace.

URL Naming

The most favored way of naming a SAML attribute is through URL naming. The creation and meaning of URLs is generally well understood by many people, and the DNS namespace is already extremely structured. Define new URLs only in namespaces you control and do your part to prevent attribute proliferation.

To create a URL name for an attribute, design a URL to be used as the identifier. If this attribute will be shared by a community, consider a URL that is common, e.g. `https://supervillain.edu/attributes/evilPersonUniqueID` for a campus-wide identifier.

URL attribute names may even be resolvable into documentation, providing helpful information for unwitting relying parties.

`urn:oid`

Section 8.2 of the SAML 2.0 Profiles suggests that LDAP attributes name themselves by utilizing the `urn:oid` namespace. These names are simply constructed using `urn:oid` followed by a standard OID. For example, DN should be expressed as `urn:oid:1.3.6.1.4.1.1466.115.121.1.12` .

`urn:mace`

The urn:mace namespace is a controlled namespace that is registered with the IETF and IANA for MACE working groups and organizations it works with. The namespace is intended to be delegated to individual organizations through registration with MACE. Once a subspace of `urn:mace` has been delegated to another organization(e.g. `urn:mace:switch.ch` that organization is responsible for any naming and resolution within that subspace. However, it's not permissible to arbitrarily define new attributes within the `urn:mace` namespace, or in any subtree you have not been granted.

Use this form to request a `urn:mace` namespace.