Internet Engineering Task Force

Internet-Draft

Intended status: Informational

Expires: June 30, 2013

K.Klingenstein, Editor

Internet2

Contributors:

Keith Hazelton, Univ of Wisconsin, Madison; John Bradley, PingId; Leif Johannson, Nordunet; Your name here;

Dec 31, 2012

Attribute Design Considerations

draft-klingenstein-attr-design-issues-00

Abstract

This document is intended to present and discuss some of the key issues in designing and using attributes within, and between, trust federations. It is intended as a companion piece to work discussing good application design in a federated world (see Cantor et al.). The document is intended to inform those planning to build multilateral federations within a vertical or community. Such groups need a shared set of attributes to serve as payloads for the exchanges.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as “work in progress”.

This Internet-Draft will expire on February, 2013.

Copyright Notice

Table of Contents

1.0 Level Set

1.0 Scope of document and intended audience

This document is intended to illuminate the areas of issues in attribute exchanges between autonomous organizations. It assumes that the organizations have one or more trust fabrics to use in order to have assurance about the validity, on several levels, of the attribute values being exchanged. These attributes are often used for real-time access control decisions but may be stored for reuse. It is intended for use by verticals or affinity groups that are intending to use a multi-lateral federation to exchange common attributes.

It does not deal with attributes that are used within an organization, for example, by a Payroll system, or solely by local services. Nor does it address instances where attribute exchange is incidental or contractually managed.

1..1 Terms:

Identity provider: An entity (usually an organization) that is responsible for establishing, maintaining, and securing the identity associated with individuals.

Relying party: An entity that manages access to some resource. Security mechanisms allow the relying party to delegate aspects of identity management to an identity provider. This delegation requires protocol exchanges, trust, and

a common understanding of semantics of information exchanged between the relying party and the identity provider.

Identity: Any subset of an individual's attributes that identifies the individual within a given context. Individuals usually have multiple identities for use in different roles or contexts.

Federation: a set of organizations that agrees to exchange a set of shared attributes using common protocols and to follow identified procedures and policies that provide trust in the validity of those exchanges. Often, particularly in multi-lateral federations, there is a federation operator that manages metadata by aggregating, standardizing and publishing it.

Attribute authority: An entity (frequently an organization) that can authoritatively assert certain attributes about particular identities. An attribute provider describes an organization, either an authority or their delegated agent, that provides attribute values to relying parties.

Schema: A set of attributes in an explicit or implicit context that are grouped together, for purpose of categorization, user interface customization, standards-based scoping, or other defined purposes.

1.2 Common attributes and schema

There are several standards-based schema that provide attributes of general interest already. In some of these instances, the schemas themselves are so large (e.g. Inetorgperson, orgperson, person) that these "well-known" attributes are cited as a reference but the explicit reuse of such an attribute at a federation-wide level is widely done (e.g. email address, which is defined in ISO ????), with each federated participant identifying their local equivalent. In some cases, there are multiple standards options for a general-interest attribute, and the federation can serve to identify particular common ones for shared use.

Federation level attributes typically represent the domains that the federation serves. Within the R&E community, eduPerson has enjoyed widespread adoption, and its parsimonious use of attributes (less than 10) and values is likely a key reason for its broad use. At the other end of the spectrum, specialized federations in the US government interagency space have defined schema with several hundred attributes.

Schema are often embedded in protocols (e.g. SCIM) and products (e.g. vCARD). Conservation of schema is a good thing but difficult to achieve in practice.

2.0 Basic tradeoffs and tussles

This section identifies a set of high-level principles and tradeoffs in attribute design for multi-lateral communities.

2.1 The bigger the community, the more basic the maximally practically achievable schema.

Perhaps the most important tradeoff to recognize is that as the size and diversity of a community sharing a set of attributes increases, the size of the subset of attributes to populate and use in a federation decreases. In some sense, this appears to be an anti-scaling consideration; in another sense it is a recognition that minimal constructions have the potential for the broadest adoption.

2.2 Authorization decisions can be made at either IdP, RP, or a dedicated third-party decision point and the choice has bearing on which types of attributes will be relevant.

In the attribute ecosystem, many exchanges of attributes are for authorization and access control. Two patterns exist today, depending on whether the locus for computing authorization is at the IdP or RP:

Authorization decisions can be made at the IdP, and typically are expressed as an entitlement attribute back to the RP. The RP needs to share the business logic for computing the entitlement with the IdP, and the IdP needs to be willing to do the work to compute authorization using the RP business logic. For scaling, a small set of common entitlements among federation members is good.

Authorization decisions are made at the RP, and the IdP releases relevant attributes for use by the RP access control system. The RP can conceal their business logic, but then may need to ask consent from the user to release the attributes. This approach may present other privacy issues as well.

There is a third alternative, attractive for its efficiencies, but not yet viable given the lack of incentives for application developers. Consider that the "business logic" may be expressed in a policy rule of the general form "subjects (S) carrying role (or group membership) G may perform actions within the set A on resources in class R". There would seem to be cases in which the desired process would be: SP (somehow) specifies or references a policy rule as above and asks a "Policy Decision Point" for a boolean-valued "attribute" whose semantics is T=>Allow, F=>Deny? Note that the Role/Group memberships of subjects might be carried in a VO attribute authority to which the SP belongs. Those Role/Group memberships might also be MANAGED by VO members with suitable delegated admin rights. One of the use case assumptions here is that those VO delegated admins will have available to them (possibly pseudonymous) identifiers for the subjects whose Role/Group memberships are being managed. If those subject identifiers can be mapped to one or more credentials, then an "undecorated user identifier" authN assertion would be all that the SP would need from the subject's IdP to make the authorization decision.

2.3 Attribute acquisition is complex and immature. In particular, the movement of attributes from source of authority to other locations, such as an IdP, is just being considered.

Attributes may be acquired by an RP in several ways. They can be asserted by the IdP at initial boarding time. They can be gathered statically at boarding time by the IdP from pre-determined attribute authorities. They can be asserted by the IdP dynamically in a real-time assertion. They can be provided, either statically or at run-time, by an attribute authority.

They may be self-asserted by a user on a web form run at the RP. With so many sources of attributes, a RP may acquire two or more different values for the same attribute and face a reconciliation decision. See the discussion below on sources of authority and LOA of attributes.

The understandings in this space are early. There are gaps in common protocols, for example to link attribute authorities to IdP's in persistent uni-lateral or bi-lateral manner. LOA issues are also not well understood.

2.4 On the wire attributes are different than stored attributes.

The focus of this discussion is on attributes "on the wire" – being exchanged between autonomous systems. In many cases these values may not ever be stored at the IdP but are calculated only at run-time from local attributes. For example, the authentication context fields of SAML assertions, which are used to specify level of assurance, are calculated from static values (e.g. how was identity vetted) and dynamic values (e.g. how was authentication most recently done). The nature of the identifiers and labels on the wire are often different than stored values as well.

2.5 Risk management

It is prudent to minimize the risk of failure. As at other levels of Internet-scale infrastructure, the principle "Be conservative in what you send and liberal in what you receive" is applicable to attribute exchanges.

2.6 Of identities, identifiers and personally identifiable set attributes

It is useful to distinguish the concepts of identity, identifiers, and personally identifiable sets of attributes. Context is an important aspect of these terms.

As noted above, Identity is any subset of an individual's attributes that identifies the individual within a given context. Individuals usually have multiple identities for use in different contexts.

Identifiers ((eg username, UUID, SSN, Subject Name, etc) are attributes that are specifically designed to distinguish one Subject from another. Authentication operations have traditionally involved the use of identifiers, so people tend to associate them with "identity", and obviously identifier attributes are often useful in any real identity system. An identity may have many identifiers. An identifier is unique to an identity within the scope of the identifier.

Personally identifiable information is any subset of attributes of an individual which identifies this individual within any set of individuals. For example, Set A is the set of all people in the building I'm in now. We're all employees of the University of Washington. So my attribute "employee of the UW" doesn't identify me within Set A, hence is not part of my identity by this definition. Set B is the set of people attending next week's bar bof. Only one of them, me, is a UW employee, so my attribute "employee of the UW" does distinguish me,

hence is part of my identity. So when I send "UW employee" on the wire, am I sending "identity information" or not?

Acknowledgments to RL Bob Morgan for the above.

2.7 Contexts

The term "context" has a number of uses in discussions around attributes, and it is useful to define the specific meaning when it is used. Context can refer to a set of coarse-grained roles, such as consumer, citizen, employee, etc. Context is also used to describe the fine-grain summation of conditions at both the IdP and RP that determine the legal basis for transactions. And context is also used to describe the expected use, and reuse, of information.

3.0 Design issues

3.1 Number and size of schema

Schema represent categorizations and organizing constructs for attribute providers, application developers and end-users. For attribute providers, a schema represents a set of attributes that they need to be able to provide for their users, but for which they may not have values. For application developers, a schema may be seen as an opportunity to get fresh values of as many relevant attributes as possible. For users a schema can be seen as a context within which to consider the release of personal information.

Existing schema may contain attributes desired by new schema. It should be noted that proliferation of attributes and/or schema has risks. Attributes need to be populated broadly to be effective; they may need to be understood and managed by people.

3.2 Attribute values and names

3.2.1 Namespace issues

Attribute names should have unique identifiers, such as a URI. Attribute values may also require unique identifiers.

3.2.2 A manageable controlled vocabulary

The specification of a controlled vocabulary (i.e. sets of permissible values for an attribute) creates the semantic precision of an asserted attribute. In general, the smaller the controlled vocabulary, the more useful the attribute will be for simplifying business logic, increasing applicability, easing the population of values by attribute authorities, etc. Perhaps the proper epithet is that the vocabulary should be "as small as necessary, and no smaller."

3.3 Semantics, sources of authority, and LOA of attributes

The related issues of semantics, sources of authority and LOA of attributes are difficult ones. They share a clear tradeoff between the rich complexity of

underlying issues and the challenges of populating and consuming such complexity in a distributed world. They relate in several ways:

The semantics of an attribute can be defined in several ways: by regulation, by a federated community of interest, by relying parties of such importance as to define attributes (e.g. national research computing resources). If there is a single source of authority, it can define its own semantics.

For example , the eduperson affiliation notes that "It is not feasible to attempt to reach broad-scale, precise and binding inter-institutional definitions of affiliations such as faculty and students. Organizations have a variety of business practices and institutional specific uses of common terms. Therefore each institution will decide the criteria for membership in each affiliation classification. What is desirable is that a reasonable person should find an institution's definition of the affiliation plausible. "
(http://middleware.internet2.edu/eduperson/docs/internet2-mace-dir-eduperson-201203.html#eduPersonAffiliation)

Sources of authority can include organizational (e.g. role in the organization), governmental (e.g. voting precinct or citizenship), peers (reputation-based attributes), self-asserted (e.g. preferred language), or a certifying authority (e.g. physicians and medical certification services).

A source of authority may issue attributes itself. Alternately, it may export some standard business logic to each IdP to compute individual users' attribute values. For example, a government may define the rules for capturing the value of citizenship by an organization. In this case, the asserted value may be issued under the authority of the IdP, rather than the attribute authority.

The level of confidence (or LOA) in the value assigned to an attribute is a difficult issue. It can be affected by the original assignment process, security of the attribute store, flaws in the process that moved the attribute from a source of authority to the the IdP, even the strength of the enabling authentication triggering attribute release, etc. At this point it seems wise to delay considering the LOA issues.

Linking of accounts can result in attributes associated with those accounts being exchanged. The process of linking accounts is a argely unexplored area in practice. A weak linking protocol can result in a lack of confidence in the values of attributes exchanged through the linking mechanism.

3.3.1 Validity period and revocation

Every credential and attribute exchange system must deal with the related issues of validity periods and revocation (dynamic end-of-trust for an assertion). Freshness of data can have economic value in a monetized attribute system.

For dynamic attribute exchanges, an immediate consumption approach avoids the need for revocations and validity periods. For longer term use cases, the validity period can be encoded into either the syntax or the semantics of the

values, using extensions attached to the attribute or conveyed in the metadata bundle.

3.4 Extensible attributes – the value of entitlements

One of the more valuable types of attributes are an extensible set of "entitlements" – values computed at the IdP in accordance with the RP's business logic, and asserted back to RP. Entitlements are a major way for an RP with many community members to simplify its application work by receiving authorization information from those members. A shared set of business rules, among the IdP customers, about access control permissions translates into a consistent assertion of appropriate entitlements. Creating an attribute that can convey entitlements is an important part of schema design.

Note that entitlements may reflect a community of interest consensus that may span multiple federations. On the other hand, the consensus may not even apply to all members of a single federation. Establishing an entitlement takes only a name space to anchor values and business logic to help community IdP's properly calculate the entitlement value to send.

3.5 Attribute mapping

Attribute mapping is inevitable as separate federations create different attributes for essentially the same characteristic. Even in those instances where separate federations share common attributes, different semantics on the same attribute may occur. (For example, eduperson primary affiliation provides distinct values of student, faculty and staff, among others. In the US, teachers and researchers are mapped into faculty, and administrative staff are mapped into staff. In the UK, they are all considered as "staff".). Further, how multiple surnames are handled is inconsistent.

Attribute mapping can be done in several ways. One can map by reference, where attributes x and y are declared mapped, and values assigned to x are consumed by SP expecting y. Alternately, a service can map by value, converting values of attribute x into values for attribute y.

3.6 Attribute metadata

A key design consideration is how to manage the metadata about attributes. Metadata may include expiration date, terms of use, provenance, etc. There are several possible mechanisms to manage the metadata.

In the federated metadata, as tags and information about various sources of authority and conditions on the attributes they issue

Another possible mechanism is to have extensions to attributes passed along with the attribute, via accompanying normative qualifiers such provenance, use restrictions, etc.

In the semantics of the attribute itself, by proscribing all of the metadata into the values and business logic of the attribute. This approach likely has limited value.

3.7 Query languages

As in other areas of information technology, common query languages could be useful. Beyond their values, they can resolve semantic inconsistencies among participants. For example, asking if a subject is "over legal age" faces problems with different jurisdictions setting legal age differently; a query language that can ask if the subject is over age X allows the RP to apply appropriate business logic.

3.8 Attribute Bundles

Attributes Bundles are a set of attributes, not necessarily in the same schema, that are grouped together in their shared pattern of use by applications.

One early observation from the attribute ecosystem is that attributes tend to travel in bundles, i.e. that certain sets of attributes are very commonly requested for certain categories of applications. Schema are one grouping, generally oriented around a standards group or a federation. Bundles are another, ones that cut across schema and represents a common "minimal attribute set" for a type of application. An example bundle is the R&S - Research and Scholarship - bundle in R&E networks that contain several attributes (name, email address, user identifier, targetedId, and user affiliation) that are typically necessary and sufficient for collaboration and research applications.

Such bundles provide a useful granularity at several levels. For the enterprise and the user, it provides a single association and greatly helps the user in providing informed consent for the release of personal information. For developers, it gives them a target category of attribute sets that are readily available for federated exchange. For interfederated use cases, however, it can be problematic if similar bundles are not composed of exactly the same set of attributes.

3.9 Application Compliance

There are several dimensions to compliance issues, including general use and disposal of attributes, qualifying for certain privacy profiles, requirements for user consent, etc.

3.9.1 Appropriate Use and Disposal by the RP

Terms of use on attributes received in a transaction is an important consideration that has a technical and a policy dimension. The technical aspects can be encoded either into metadata (thus passing out-of-band) or conveyed in the transaction as extensions to the attributes themselves. The latter approach presents many of the issues considered in classical digital rights management.

While not a direct part of attribute design per se, the policy and trust framework around the attributes should address proper use and/or disposal of attributes

received at the SP. A representative policy can be found at https://refeds.terena.org/index.php/Code_of_Conduct_for_Service_Providers.

3.9.2 Certifying applications for attribute bundles

Corresponding to attribute bundles is the need for validating that applications are appropriate, in purpose, controls, etc. for release of a particular bundle by the IdP or the user. This certification can be self-asserted or vetted and audited by the federation or a third party. Services could attest that applications were compliant with "standard" privacy profiles, be they restricted to youth, restrictions against use by youth, citizen, etc.

3.10 Privacy Considerations

Several privacy themes are currently in active discussion:

What is personally identifiable information (PII)? Technically, PII is very contextual, as noted above. A set of attributes may personally identify a person within a certain set of individuals, but not within others. A user's expectation of privacy in interactions depends on context of the transactions, e.g. are commerce, citizen, or social contexts. Privacy policy is not sufficiently context sensitive at this point.

When is consent by the user needed to release attributes? Again, policy is inconsistent and under much discussion at this point. The EU Privacy directives play a significant role, and other national privacy policy controls are emerging.

4. Acknowledgements

This work was inspired by discussions at the ISOC identity ecosystem workshops held in Amsterdam and Gaithersburg MD in 2011 and 2012 and subsequent IETF meetings.

5. Contributors

Main contributors for this work has been

6. IANA Considerations

This memo includes no request to IANA.

7. Security Considerations

Attributes are often used to carry sensitive information as part of claims-based protocols. It is common for claims to contain attribute values that are used to allow or deny access to a protected resource. Some attributes carry identifiers as values. A discussion of the security implications of handling identifiers can be found in draft-iab-identifier-comparison.

8. Normative References

Authors' Addresses

Ken Klingenstein Klingenstein, Kenneth, Internet21000 Oakbrook Drive Suite 300Ann Arbor, MI 48104USPhone: +1-360-562-0319EMail: kjk@internet2.edu