

Federated Authorization Problems and Models

What follows is the initial post and a [summary of comments](#) of a [thread](#) on MACE-Paccman and REFEDS about authorization in federated environments.

To paraphrase Roland Hedberg, it is high time to seriously address authorization as we work on our (inter-)federation identity and access management (IAM) infrastructures. Two patterns are commonly found today, depending on whether the locus of authorization evaluation is at the IdP or RP and I would argue that there is a third alternative that is worthy of consideration.

- 1) If authorization decisions are to be made at the IdP then they are typically expressed as entitlement attribute/values back to the RP. The RP needs to share the business logic for computing the entitlement with the IdP (out of band), and the IdP needs to be willing to do the work to compute authorization using the RP business logic. This scales well for a small set of common entitlements across a(n) (inter-)federation.
- 2) If authorization decisions are to be made at the RP then the IdP needs to release relevant attributes about the authenticating subject for use by the RP access control system. Group/Role memberships of the subject are among those relevant attributes. In this case, the RP can conceal its business logic, but it may need to obtain consent from the user and/or the user's IdP as a condition for obtaining those attributes.
- 3) Warning, some concepts that follow are borrowed from the XACML conceptual model. Consider that in simple situations the "business logic" referred to in 1) and 2) above might be expressed in a computable 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: Subject S shows up at the RP and requests to do action A1 on resource R1. The RP (somehow, possibly out of band) specifies or references a policy rule as above and asks a "Policy Decision Point (PDP)" for a boolean-valued response whose semantics is T=>Allow, F=>Deny. VERY few of our existing infrastructures include anything like a PDP.

Note that the Role/Group memberships of subjects might be carried in a VO attribute authority (AA) to which the RP belongs (think Policy Information Point, PIP). Those Role/Group memberships might also be MANAGED by VO members with suitable delegated admin rights. One of the challenging 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 they are managing. If those subject identifiers can be obtained from an IdP in a SAML assertion, then an "undecorated user identifier" would be all that the SP would need from the subject's IdP to enforce the authorization decision. Note that this model locates the Policy Enforcement Point PEP at the RP. One further extension would be useful: The VO could offer a service by which an RP admin could define and persist their RP-specific computable access policy rules (think Policy Administration Point, PAP).

We have not generally thought of our (inter-)federation IAM infrastructures as containing PIPs (well, we have started to talk about AAs) or PAPs in addition to IdPs and RPs, but if model 3) is interesting to us, we will need to think those thoughts. --Keith Hazelton (6 September 2012)
