Federated Authorization

Keith Wessel
Beth Vanichtheeranont

What is Federated Authorization?

A few key terms:

IdP – Identity Provider
SP – Service Provider
SSO – Single Sign-On

What is Federated Authorization?

A few key terms:

Authn – AUTHeNtication – you are who you say you are

Authz – AUTHoriZation – this is what you can do (because of who you are)

What is Federated Authorization? A few key terms:

SAML – Security Assertion Markup Language
OIDC – OpenID Connect

What is Federated Authorization? A few key terms:

RBAC – Role based access control
ABAC – Attribute-based access control

Federated Authentication vs. Single Sign-On: what's the diff?

Single Sign-On (SSO) allows access to multiple SPs without having to reenter credentials from your originating organization

Fed authn allows access to multiple SPs within a federation (multiple different organizations) using your home organization's credentials

What is Federated Authorization?

Releasing information from your home organization about you to grant or deny access to SPs from outside organizations.

How is information released?

eduPersonAffiliation:
Specifies the person's relationship(s) to the institution in 8 broad categories

- faculty
- student
- staff
- alum
- member
- affiliate
- employee
- library-walk-in

Ex: HathiTrust only allows access to student faculty and staff affiliations

How is information released?

eduPersonEntitlement: URI (either URN or URL) that indicates a set of rights to specific resources

ex: College of Law is federated with a printing service for billing that is accessible based on college code. If their code is College of Law they're allowed to print

ex: our library releases the InCommon entitlement value to publishers that they've mapped to types allowed to access subscription services

urn:mace:dir:entitlement:common-lib-terms

https://sp2.dontsueus.com/entitlement/funds-uiuc-law

Use case: Federating REDCap at UIUC

The Current Landscape

- Researchers in current research project(s) from multiple institutions collaborating in UIUC REDCap instance
- All said institutions are part of the same federation (InCommon)
- The research data is held in UIUC RedCAP instance

How It Works Now

- Bi-laterally federated REDCap instance
- Guests need to take privacy training through our Knowbe4 instance (also bilaterally federated with Illinois)
- Guests from other institutions need guest accounts to access said training OR
- Guests have to provide proof of equivalent privacy/HIPAA training manually to PI
- PI in charge of verifying training/access eligibility every year - manually

The (Potential) Solution

- Federate our REDCap instance multi-laterally (with InCommon)
- Researchers home institution tracks training (via Grouper group or similar RBAC)
- Group membership would have expiration dates set based on HIPAA certification
- User's home org can automate verification of annual training via API call to training platform
- Share an attribute (ABAC) derived from the above group such as eduPersonEntitlement with UIUC REDCap
- User now logs into REDCap with home credentials and IdP has indicated that they have current training

Potential Pitfalls

- Privacy has issues with lack of comprehensive coverage of some third party training programs
- SP doesn’t want to evaluate every third party for content quality
- Would rather have ‘all or nothing.’

Potential Mitigations

- Could make eduPersonEntitlement value in our REDCap namespace ex:
  https://healthinstitute.Illinois.edu/shibboleth/hipaa-trained
- Every school has to map that attribute for every institution federating with us in their namespace with business rules about who meets this criteria for this requirement
- Have to repeat this process for every new member
- (This is current best practice)

Potential Mitigations

- Could have eduPersonEntitlement value in each school's namespace
- urn:mace:incommon:uiuc:edu:certification:hipaa BUT
- Every school would have a different value and that would cause scalability issue for our REDCap instance

Potential Mitigations

- Could have InCommon designate a value that specifically stands for ‘this person is currently HIPAA-trained’ and tell the rest of the members ‘this is how you share this data’ e.g.
  urn:mace:incommon:certification:hipaa
  BUT
- What if we didn’t feel institution X’s training was up to our standards?
- REDCap could flag as ignored any that came from Institution X but we’d have to eval every institution

What if…?

- Have a standard format that includes domain ex:
- urn:mace:incommon:certification:<yourdomainhere>:hipaa
- That way, instead of having to encode a different value for every IdP they can instead regexp to match against the domain and allow/deny based on that

What if…?

- Include name of training source (third party instructor) and evaluate based on what we know about their certification class content
- urn:mace:incommon:certification:<yourdomainhere>:hipaa:<trainingprovider>
- Majority of institutions outsource training to a limited list of vendors, evaluate based on that
- Especially nice for institutions who require different levels of training for different projects

Likely path

- We'll likely start with option one, as it's the current best practice, but any of these could work, its just a question of scalability/how widely we want it used

Future State

- We want to move towards more scalability and much less overhead for either SP OR IdP
- Express things about the org instead of individual: e.g. this person's institution is GDPR compliant

Broader Use Case: Solving challenges with Open Access

Open Access (OA)

- A set of principles and a range of practices through which research outputs are distributed online, free of access charges or other barriers

Open? Access?

- Federated identity and access control
- Discussion - looking for open access control issues and patterns
    - Time embargoes, geo-location embargoes
    - Institutional rules
    - Repository federation rules
    - Special collection rules
    - Privacy and PII considerations
- Engineering an approach that implements the rules… and the exceptions
- How do identity federations align with repository federations?

Solving Open Access Challenges with Federated Authorization (fed authz)

Federated Authz

An institution could release standard values of eduPersonEntitlement that could express any of a number of characteristics about an individual that would grant them access to restricted collections

Federated Authz

- Department
- Location
- Position in the Organization
- Any other credentials that would be used to determine if they were eligible to access a restricted collection

Federated Authz

- And do it all with limited information sharing about user, just enough to make a decision
- While preserving privacy

Discussion