Identity Hubs Capabilities Perspective

A White Paper from Rebooting the Web of Trust V

by Adrian Gropper, Drummond Reed, Mark S. Miller

ABSTRACT

Identity Hubs as currently proposed in the Decentralized Identity Foundation (DIF) are a subset of a general Decentralized Identifier (DID) based user-controlled agent, based on ACLs rather than an object-capabilities (ocap) architecture. The current approach has both security and scalability issues. Transitioning the Hubs design to an ocap model can be achieved by introducing an UMA authorization server as the control endpoint. This avoids creating confused-deputy security issues and expands scale by enabling the hub to delegate access to resources not stored in the hub itself.

SECURITY AND PRIVACY IMPACT ASSESSMENT

Control over personal data is becoming more important as networking and storage costs become negligible compared to the value of our data itself. A prominent example of this is the EU General Data Protection Regulations (GDPR) and Payment Services Directive (PSD2), which mandate sweeping new responsibilities for transparency and control to all enterprises that hold personal data.

In general, personal data fits into three distinct categories:

- Data generated by a service or device (e.g., a lab test of my anonymous blood sample or location streamed by my mobile device)
- Data aggregated by a service (e.g., a credit reporting enterprise)
- Data aggregated by a fiduciary or self-sovereign agent (e.g., a personal server that I own or a secure mobile wallet)

Privacy threads through all aspects of access to personal data for:

- the subject (we ignore guardians in this short paper) of the personal data; and
- the parties seeking access to the personal data about the subject.

A scalable personal data architecture will guide the implementation of secure and privacy-preserving services without the costs and risks of unwanted aggregation.

SELF-SOVEREIGN IDENTITY

The scale of a personal data architecture is limited by federation as a root of trust. Federations work best when applied to entities of similar size and scope. Banks can form successful federations of ATMs. Professional or semi-professional sports teams can federate, separately, to form leagues. But a federation related to identity could span the full range of human, machine, and institutional entity activity across countries, cultures, professions, and risks.

A standardized self-sovereign identifier (e.g., a DID) uses distributed public ledgers as a source of trust while retaining the option, on a person-by-person basis, to leverage federated identities or not.

Self-sovereign identifiers and verifiable credentials

are useful for both the subject and the requesting party. Both can benefit from self-sovereign mobile wallets with local biometric protections and from the ability to sign documents in a legally non-repudiable manner.

Self-sovereign identifiers serve as the root of trust for the subject's agent and for claims presented by a requesting party.

SELF-SOVEREIGN AGENT

The essential component of the agent is the ability to authorize access to the subject's resources by requesting parties. As such, it secures the subject's policies without actually sharing the policies. By way of analogy, cryptographers design secure elements to protect private keys from having to be shared while still being effective in the actions they control. The execution of arbitrary code as part of a secure element is core to the general capability to issue authorization decisions while also keeping secret the policies that led to the decision.

To the extent that an agent also hosts a database holding data about the subject that can be shared with requesting parties, access to that data is also under the control of the authorization server.

IMPLEMENTATION EXAMPLE

Preamble: Subject Alice goes to have four blood tests at a resource service Lab; they will be made available, in the future, under her control to various requesting parties. One of these requesting parties is Bob. Bob has verifiable credentials as a licensed MD and a member of his temple. Alice controls an authorization server (AS).

1. Alice presets a policy in her AS that any MD can access anything other than "Sensitive" tests. This happens only once and applies to many transactions with many service providers. This can be inherited along with the code for the AS from a source that Alice trusts.
2. Alice registers in-person at Lab and provides a blood sample and a pointer to her AS endpoint. The Lab registers four tests as four different resources with the AS. One of the resources has metadata STD. This happens only once regardless of how many times Alice gets lab tests at Lab. Time passes...
3. Bob, using out-of-band directory interactions or a referral from someone, is given a pointer to Alice@Lab resources. Bob may or may not know that there are four of them at Lab.
4. Bob goes to Lab to request the resources. He does not identify himself using any traditional forms of identification. Lab looks up Alice's AS endpoint and gives Bob, whoever he is, a pointer to the AS along with a ticket that captures the context of the transaction for later reference.
5. Bob authenticates to Alice's AS and presents the MD credentials but not his temple membership. The context ticket says he's asking for three of the four results.
6. Alice's AS consults the policies and the metadata associated with the context ticket (as registered, or dynamically at the time of the request) and determines that one of the three tests is Sensitive. In some cases, metadata shows up that the AS does not understand and Alice herself may need to be consulted; otherwise the process is automatic once Bob has submitted the ticket and the claims.
7. The AS issues to Bob's software client a token granting access to two of the three tests he requested. The token can be associated with OAuth refresh so that Bob will not need to re-authenticate for a year.
8. Bob's client presents the token to the Lab. The token can encode the permission for the two tests or the Lab can consult the AS to determine the scope of access associated with the token.
9. The Lab exposes two of the four tests to Bob's client.

- Note that Alice can choose to move some or all of the Lab data to an aggregator because she considers the Lab to be unreliable beyond a certain time. If she does that, Alice will need to ensure that the lab test retains provenance to Bob's satisfaction and that represents a significant added cost.
- Although delegating access control to an external AS is an added cost to the Lab, it is also a feature that attracts Alice as a customer and it reduces their legal liability under GDPR.
- Note that Alice and Bob both have the ability to attenuate the scope of disclosure about them.
- The only single point of failure is Alice's AS, which must be sufficiently reliable.
- Alice has a single point of control for all of her participating services, which is a major convenience.
- Alice can have as many AS endpoints as she wants separate personas. To avoid unwarranted correlation, the AS endpoint can go through a mixer or hub.

Token-based access control to RESTful APIs has been standardized as OAuth2. The UMA standard builds on top of OAuth2 to enable the separation of the authorization server from the resource server. The subject authenticates into the authorization server of the agent in order to create or modify policies and to manage the code that uses policies to respond to requesting parties. The requesting parties bring tickets to the authorization server that represent the context of a resource access request along with whatever credentials the requesting party chooses to expose to the authorization server. The requesting party then receives a limited scope access token to the resource regardless of whether the resource is coresident in the agent or hosted by an independent enterprise.

PROPOSAL TO DIF

We propose an update of the DIF Identity Hub to specify OAuth2 and UMA as the protection mechanisms of the hub. That would be followed by a privacy impact assessment of all aspects of the Identity Hub specification including specific reference to GDPR and PSD2 compliance.

REFERENCES

ACLs Don't

http://www.hpl.hp.com/techreports/2009/HPL-2009-20.pdf

UMA

https://kantarainitiative.org/confluence/display/uma/Home

DIF Hubs

https://github.com/decentralized-identity/hubs/blob/master/explainer.md

---

ADDITIONAL CREDITS

Authors: Adrian Gropper, Drummond Reed, Mark S. Miller

About Rebooting the Web of Trust

This paper was produced as part of the Rebooting the Web of Trust V design workshop. On October 3rd through October 5th, 2017, over 50 tech visionaries came together in Cambridge, Massachusetts to talk about the future of decentralized trust on the internet with the goal of writing 3-5 white papers and specs. This is one of them.

Preliminary Workshop Sponsors List: BigChainDB, Blockchain Lab, Digital Contract Design, IDEO, IPFS, Protocol Labs, Toni Lane Casserly

Workshop Producer: Christopher Allen

Workshop Facilitators: Christopher Allen, with additional paper editorial & layout by Shannon Appelcline.

What's Next?

The design workshop and this paper are just starting points for Rebooting the Web of Trust. If you have any comments, thoughts, or expansions on this paper, please post them to our GitHub issues page:

https://github.com/WebOfTrustInfo/rebooting-the-web-of-trust-fall2017/issues

The next Rebooting the Web of Trust design workshop is scheduled for early 2018 on the west coast of the USA. If you'd like to be involved or would like to help sponsor these events, email:

ChristopherA@LifeWithAlacrity.com