Introduction to Federated Identity and the FedID CG

Abstract: The audience for this post is people who are unfamiliar with how privacy concerns may impact federated identity. That likely includes people from one of two groups: 1) people who are interested in privacy, but are new to the concept of federated identity, or 2) people who are familiar with federated identity, but are unaware of the changes being made to browsers because of the privacy concerns. The goal of the post is to provide an introduction to federated identity and why it matters, what the privacy changes are and their potential impact, and then describe how FedID CG is working to preserve federated identity in light of the privacy-related changes.

Federated Identity encompasses the technologies, standards, and use cases in which the user identification and user authentication services are separated from the service providing the resource a user is trying to access. The organizations providing the user identification/authentication services are generally referred to as Identity Providers, and the organizations that utilize their services are often referred to as Relying Parties. Federated identity makes it possible for a website, app, and/or API to outsource authentication to an external entity. In practice, users with an account with entity A can gain access to web apps B and C without having to create new usernames and passwords, if B and C outsource authentication. Sometimes referred to as Single Sign-On, or SSO, there is a distinction to be made between SSO and federated identity. SSO is a property of federated identity that makes it possible for a user to gain access to distinct web apps or API without having to reenter credentials. The broader use of federated identity is when the resources involved are located in different security domains and are owned by different organizations.

The types of organizations that use federated identity are as varied as the internet. It's a common practice to use federated identity to streamline account management and access by allowing users to log in with an identity provider account (those "Log in with Facebook", "Log in with Google", "Log in with …" buttons.). It's also commonly used by businesses to manage their emplo yees' access to company resources. Universities use federated identity to offer students multi-institutional academic programs, to provide shared access to educational resources, and for research collaboration. Federal institutions use federated

identity to manage access to federal resources too - as do financial institutions. And for one final example, it's also frequently used in Software-as-a-Service business models as well.

Federated identity significantly reduces the burden on users by limiting account proliferation. It streamlines the user experience, lowers the security risks associated with password re-use (e.g., credential stuffing attacks), decreases the raw number of access credentials that a user has to remember and manage, and facilitates inter-organizational relationships and management.

However, linking a user's identity across systems also raises privacy concerns, especially when done across organizations/entities (and to a much lesser extent even when the resources are all owned by the same organization). While the objective of federated identity systems is to facilitate a user's access to resources online, it was originally designed on top of web primitives (e.g. third-party cookies, top-level navigations, etc). These primitives can and are being abused to track users without their consent or full understanding.

In response to these concerns, user-agents are making changes to how they work with some of the fundamental primitives of the web to prevent the uncontrolled, hidden tracking of users. Since federated identity often utilizes these same primitives to exchange necessary information to complete authentication flows, we need to develop solutions that address these privacy concerns without breaking federated identity.

There are a variety of privacy-related interventions that user-agents are exploring. These include deprecating third-party cookies, controlling access to the client's web storage, removing certain parameters from links (often referred to as link decoration), and restricting the capabilities of navigational redirects. Federated identity often relies on these same mechanisms, and so the changes being made to improve support for end-user privacy are having an effect on federated identity systems. Since the most immediate change is the deprecation of third-party cookies (having already been deployed in Safari and Firefox, and publicly planned for Chrome in late 2023), the Federated Identity Community Group (FedID CG) is currently focusing most of its attention on the impact of that change. The group is working to preserve federation when third-party cookies are deprecated.

The FedID CG meets every week to provide feedback on the proposals that are relevant to federated identity. The full charter for the group can be found here. If you're interested in learning more, we are currently working on a draft report that will be published shortly. If you'd like to participate in the group, you can join FedID CG here. Please note that while a

W3C account is required to join, you do not need to be a member of the W3C. If you don't have a W3C account, you can sign up for one on the W3C account request page.