Federation at Scale

Scaling the Federation

The federated approach to identity management is generally more usable and significantly more secure than the non-federated approach. On the other hand, federation is more difficult to deploy but perhaps not significantly so, especially if done with care. In any case, federation as a general phenomenon has seen only modest growth at best.

The reasons for the lackluster success of federation are varied and not well understood. Landau and Moore [1] offer a plausible explanation for the failure of federation on the open Internet. They conclude that there is little incentive for RPs to federate, especially since IdPs derive most of the benefit from an RP's decision to participate in cross-domain federation.

Growth in the InCommon Federation

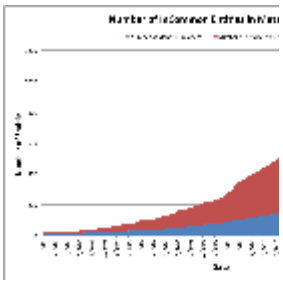

The number of entities (identity providers and service providers) in the InCommon Federation is apparently growing by leaps and bounds, but breaking this down, the number of service providers (SPs) is growing exponentially while the number of identity providers (IdPs) enjoys only a modest growth rate. Even though the number of SPs is actually underestimated (since some entity descriptors in metadata have literally dozens of distinct endpoints) we claim the number of SPs is **not** a good indicator of overall Federation growth. To understand the dynamics of the InCommon Federation, one should follow the growth of IdPs over time.

Why is this? Well, more and more organizations are introducing non-federated web applications into metadata (i.e., "non-federating" from a cross-domain point of view). In other words, the Federation Manager is becoming a software-as-a-service application that organizations use to manage their local SAML federation. We conclude from this that the number of SPs in metadata indicates that SAML Web Browser SSO is catching on in a big way, precisely because it *does* address cross-domain federation in addition to SSO. For most campuses, that translates into numerous (often hundreds of) local SAML deployments.

The upshot is that *cross-domain web browser SSO* currently enjoys only modest growth. This article exposes perceived barriers to growth and identifies potential solutions that purport to remove those barriers, thereby raising federation to a new level, that is, federation at scale.

The PAP Test

In what follows, we focus on three aspects of federation that figure prominently in an entity's decision whether or not to federate:

1. Penetration
2. Assurance
3. Privacy

We examine each of these concepts carefully, both in principle and in practice. In other words, we perform a kind of *federation PAP test* (where "PAP" of course refers to Penetration, Assurance, and Privacy) thus arriving at a health assessment of federation in the large.

Penetration

Penetration is the probability that an arbitrary user is able to provide an acceptable authentication context to the service provider (SP) on demand. Usually this just means that the user possesses a federated identity that the SP will accept in lieu of local authentication (which is basically what federation is all about). A phrase that captures this scenario is "Bring Your Own Identity" (BYOI). The implication of BYOI is that every time the user leverages an existing password, the world becomes just a little bit better place.

For the purposes of discussion, we define the following terms:

- *penetration filter threshold*: below this threshold, there are perceived barriers to federation
- *penetration pump threshold*: above this threshold, federation proceeds unhindered

AFAIK there is no available research that suggests what these threshold values might be in practice (but see the theory of Diffusion of Innovations). In any case, there's absolutely no reason to believe that the filter threshold and the pump threshold are the same. What is more likely is that between these two threshold values, a federation waivers in a kind of "wait and see" mode, with no significant barriers or enhancers.

In the InCommon Federation, we have not yet reached the filter threshold; that is, there are barriers to Federation growth that RPs actually take into account when deciding whether or not to federate.

In contrast, on the open Internet, the pump threshold has already been exceeded. Yet federation still flounders on the open Internet, and we wonder why. Apparently there are other factors at play here, some of which are discussed by Landau and Moore. [1]

Assurance

By definition, federation introduces an untrusted third party, namely, the identity provider (IdP). Thus the service provider (SP) quite naturally wonders about the level of assurance (LoA) associated with the user's IdP. Assurance requirements will vary from SP to SP of course, but whatever authentication context is required, it **must** be requested at runtime. Today the SP most likely does not care about LoA, either because the service itself is low risk (e.g., a wiki) or the transaction to be authorized simply does not require a high LoA.

In the InCommon Federation, IdPs are labeled as Bronze or Silver, indicating LoA-1 or LoA-2 (roughly). In principle, this satisfies the Federation's need for assurance. In practice, however, Bronze and Silver are expensive, for both the IdP and the Federation Operator.

On the open Internet, assurance is still mostly an unsolved problem. There are a handful of FICAM approved IdPs on the open Internet (including Google, PayPal, and VeriSign), but in practice the LoA of an OpenID assertion is neither sought after nor consumed at the RP. Today, LoA on the open Internet is mostly a solution looking for a problem to solve.

Privacy

Privacy functions as a showstopper in nearly every federation in existence today. In the InCommon Federation, privacy concerns at the IdP (fueled by FERPA) lead to conservative attribute release policies whereby an unsolicited SP is initially denied the identity attributes it requires to function. This forces bilateral agreements to become the norm and detracts from the overall user experience.

In response to this unfortunate state of affairs, InCommon has launched the Research and Scholarship (R&S) category of service providers. In principle, an R&S SP is more trustworthy than an arbitrary SP chosen at random, and therefore the IdP is induced to release attributes to R&S SPs "sight unseen." The mechanism for accomplishing this is attribute-based policy configuration at the IdP. The attributes used for this purpose are called entity attributes, which promise to dramatically reduce the configuration overhead associated with attribute release policy at the IdP.

R&S is a baby step in the right direction. Currently there are eight (8) R&S SPs while more than 30 IdPs have self-asserted their support of R&S. If R&S takes off, it could have a dramatic effect on attribute release (and hence on privacy) in the InCommon Federation.

The effect of service categories on other federations is unknown. In the EU, for example, privacy is so tightly controlled (via legislation), it's doubtful an R&S-like scenario could take hold. There are, however, rumblings in the EU even as we speak, so only time will tell.

On the open Internet, privacy is the rallying cry for a large group of constituents poised to derail the NSTIC effort. A major challenge to NSTIC is how to embrace privacy with political and technical correctness. A prerequisite (it seems) is that the IdP should have zero knowledge of the service ultimately visited by the user. This appears to be a technical challenge since the lack of this knowledge at the IdP is itself a privacy concern. A practical solution to this apparent contradiction is unknown (to this author, at least). We look to the NSTIC process to propose a solution that solves the privacy problem at both ends of the federated transaction simultaneously.

Stimulating Growth

Clearly there are numerous barriers to successful federation at scale. In what follows, we make some relevant observations for the InCommon Federation only.

If every SAML-enabled SP in the InCommon Federation were also OpenID-enabled (specifically, OpenID 2.0), these SPs could leverage the OpenID IdPs on the open Internet. We claim there is a huge advantage in doing so. Not only would the penetration threshold significantly increase but users would be encouraged to BYOI. For students, parents, and friends of the university in particular, the benefits would be significant.

As mentioned earlier, not all SPs care about LoA. Those that do can deploy mobile-based two-factor authentication (2FA) *at the service provider*. The technology exists today to do this safely, easily, and at low cost per user. Combined with a "something you know" first factor at the IdP (i.e., a federated password credential), mobile-based 2FA can provide sufficiently strong LoA for the vast majority of high-risk SPs. In effect, the Federation could realize the benefits of Bronze and Silver without a corresponding high cost of deployment.

To OpenID-enable every cross-domain federated SP in the InCommon Federation, and to simultaneously 2FA-enable the high-risk SPs (federated or otherwise), is of course an expensive proposition. To mitigate this cost, a centralized gateway might be deployed (at the Federation level) to realize certain economies of scale. The gateway would, in effect, put the theory to the test. If sufficient benefits accrue (metrics to be determined), the technology would ultimately be pushed to the fringes of the Federation.

This strategy has been successfully deployed before, that is, deploying a centralized service that temporarily gives RPs the benefit of new technology without the corresponding high cost. As the technology proves itself (or not, as the case may be), the natural tendency of federation is to push the technology to the edges.

This brings us to the third and most challenging component of the "PAP" test: privacy. Unfortunately, AFAIK there is no "silver bullet" that addresses the privacy issue. It is critically important that we find such a solution, however.

The reader will have already observed that successful penetration and adequate assurance give rise to "universal, trusted authentication," whereas privacy crosses squarely into the realm of authorization. Without privacy (or attribute release, depending on which side the coin falls), the overall federation problem is only half solved.

We believe *intra*-federation has a fighting chance with respect to privacy. In the InCommon Federation, for instance, service categories such as Research and Scholarship, in conjunction with standard SAML flows and generalized policy configuration based on entity attributes, will significantly improve the federated user experience, which in turn will increase the demand for federation.

There's nothing special about SAML in this hypothetical scenario. An equivalent endgame could be realized using OAuth2 technology. Indeed, there is some very interesting work going on right now that leverages OAuth2 for user consent-based attribute release. The latter appears to be a prerequisite for *inter*-federation in particular since EU federations (for instance) require user consent for all intents and purposes.

Federation has many moving parts, with only partially understood dynamics. Periodic game changers are expected and invited. At this point in time, it seems the InCommon Federation might benefit from the following innovations:

- a centralized gateway that bridges the open Internet and provides ubiquitous strong authentication to SPs that need it
- usable service categories and attribute-based policy configurations at the IdP
- an improved user experience by virtue of all of the above and related efforts such as enhanced discovery and federated error handling
- a simplified user consent mechanism

Are we there yet? No, not quite.

References

[1] Susan Landau and Tyler Moore. *Economic tussles in federated identity management*. http://firstmonday.org/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/4254/3340