

Anonymous Credentials - A Report from the Scalable Privacy work

Anonymous Credentials:
A Report from the Scalable Privacy work

Background:

The NSTIC Scalable Privacy activity at Internet2 committed to an investigation of anonymous credentials. Anonymous credentials represent a small set of small intellectual properties, across a variety of copyrights, that provide unique capabilities identified in the NSTIC guiding principles – unobservability and the ability to provide assertions without revealing PII. Unobservability refers to the capability of a credential being used without the issuing authority knowing of its use. Providing assertions without revealing PII is exemplified by the ability of anonymous credentials to respond to assertions (e.g. “user is over 21”) without revealing the actual age. While the anonymous credential technology goes back over ten years to a Ph. D. thesis by Stefan Brands under the leadership of Ron Rivest of MIT, a number of factors have stopped the widespread deployment. This analysis highlights those factors and our efforts to address them.

Methodology:

The project leads, Ken Klingenstein and Steven Carmody, have a combined forty years of experience in PKI efforts within the leading edge R&E community. The project engaged with Anna Lysyanskaya of the Brown University Computer Science Dept. on the state of the art in anonymous credentials. That led to conversations with the two principal owners of the intellectual property Microsoft and IBM. (The IBM work involved the Alexandra Institutet, A/S of Denmark.) The conversations with Microsoft, related to their CPIM implementation, found common interest around the PrivacyLens UX we have been developing as part of the NSTIC grant, but never gained traction due to the complexity of the Microsoft organization.

With IBM the engagement has been sustained, and productive, to some degree. The IBM code base has been provided, generally in alpha state of a new version. New versions are delivered on an irregular schedule, and associated libraries follow later. We have been able to install a working system recently, to generate credentials that can respond to queries. The current version could best be described as an early prototype. It can only be used from the command line; you run a script that does a series of about a dozen queries to various services, and eventually types out a “yes” or “no”. There are no tools for creating or managing credentials. It currently cannot be used with a web browser (a requirement for most users). It is unfortunately far from a widely deployable product able to support a variety of needs, with facile credential creation, etc.

In addition, we've engaged in conversations with the principal EC effort – ABC4 trust. (www.abc4trust.eu) They are in the process of a yearlong study about what are the barriers to anonymous credential deployment. They have reviewed this paper and will include it in their analysis.

Findings:

1. “Classic” anonymous credentials still have a number of barriers to use. The marketplace is compounded by the intellectual property rights issues and the fact that the two companies owning it have much broader interests than identity management. It is a chicken and egg problem. The issues involved in training users how to manage certificate-based technologies have traditionally been a huge impediment as well. The companies respond to markets and there is none.
2. New requirements for credentials may present challenges in preserving unobservability. Requirements include informed consent, accessibility. Mobility, a variety of devices and operating systems, etc.
3. The license regime around anonymous credentials is an impediment to use. The abc4trust intent is to provide an open source implementation, which would be welcome. But then the code is only useful if it gets bundled with product (i.e. browsers); users won't download and install it themselves. so the adoption problems are challenging.
4. Unobservability is a difficult requirement, and the lack of anonymous credentials has led to workarounds that are problematic. Double blind portals and gateways are difficult to do right and are one-offs due to a lack of a marketplace as well.
5. Revocation remains a difficult issue for any static credential. That said, there are use cases for which revocation is so unlikely as to be unimportant.
6. It is likely that a deployable environment will need to use a hybrid architecture where some of the unobservability needs are met with a legal and policy approach rather than pure technology. This will be true with cloud-based approaches for mobility, for example.
7. Another hybrid approach that we evaluated was an enterprise-centric architecture. The enterprise would hold credentials for users rather than at the desktop. This would enable a host of capabilities, from mobility to support of informed consent, but at some cost of unobservability, which might need to be managed via policy rather than pure technology. Our goal was to implement this hybrid approach, but we did not get there. The code only provides very early stage functionality, and could never be given to an actual end user. And, for it to work at all and provide unobservability, a user would have to run a set of services within their desktop machine. That would have to be installed by whoever provided the desktop operating system. And the mobile world is moving so fast that desktop/laptop computers are almost irrelevant now.