Achieving the Privacy Goals of NSTIC in the Short Term

Francisco Corella
Pomcor

Karen P. Lewison
Pomcor

Revised May 4, 2011

Executive Summary

NSTIC has challenging privacy goals. In particular, the identity of the relying party must not be revealed to the identity or attribute provider, the knowledge provided to the relying party must be minimized, and the knowledge provided to relying parties in separate transactions must not be linkable. In the long term, zero-knowledge proof technologies will provide elegant solutions to these challenges. A shorter term solution, however, is urgently needed. In this white paper we outline an approach to achieving practically all the privacy goals of NSTIC in the short term. The approach will preserve and leverage much of the investment made by the US government and the private sector in existing identity technologies, and will provide a smooth transition to longer term zero-knowledge solutions.

1 Privacy Goals of NSTIC

The US National Strategy for Trusted Identities in Cyberspace (NSTIC) [1, 5] was officially launched on April 15 [4]. Like other Internet identity initiatives, NSTIC aims at obsoleting passwords and providing a Web single sign-on soluction. But in contrast to other Internet identity initiatives, NSTIC also aims at providing unprecedented levels of user privacy. This was emphasized in a recent blog post [8] by Howard Schmidt, White House Cybersecurity Coordinator. He said:

> We have an opportunity to design privacy directly into the fabric of the Identity Ecosystem. In fact, there are innovative technologies that will do just that. Consider this: in the physical world, people who show a drivers license to prove their age also reveal their name, address, height, eye color, and other unnecessary information. Online, however, you could have a credential that uses a privacy-enhancing technology that would not reveal any extraneous information.... Identity Ecosystem credentials can also keep your credential provider from tracking all of the websites you surf, or log into. And unlike your drivers license, which has a single identifying number on it, these privacy-enhancing technologies can enable you to log in with a different identifier for each website.

Anonymous credentials based on zero-knowledge proofs [3] hold the promise to satisfy these privacy goals, and may very well be the privacy-enhancing credentials that Howard Schmidt had in mind. However, it seems fair to say that these innovative technologies are still at the research stage, and it may take years until they mature into broadly deployable solutions.

In the meantime, though, there are immediate needs that call for a shorter term solution.

2 Need for a Short Term Solution

There has been several efforts in recent years to provide Internet identity and Web single sign-on solutions. OpenID [7], in particular, has achieved relative success. But a dramatic development is now taking place. Using OAuth [6], a few social and multi-purpose sites that cumulatively count as their users almost the totality of Web users, viz. Facebook, Google, Yahoo, Myspace, LinkedIn and Twitter, allow other Web sites and Web applications to delegate user authentication to them and, as a side effect, to gain access to user data and the user's social context, by a process that has been called "social login" [2]. More and more sites and applications are offering social login to their users. If this trend continues, one of the goals of NSTIC, Web single sign-on, may be realized to a large extent, in the short term, by social login.

But the current approach to social login has serious drawbacks and dangers.

First, social login reduces privacy rather than increasing it, because the social site is aware of the user's login to relying parties (Web sites and Web applications), and, often, of the user's activity at those relying parties.

Second, OAuth requires prior registration of the relying party with the social site. If social login becomes a de-facto user authentication standard, most Web sites and applications will have to register with the dominant social site, currently Facebook, just to be able to authenticate their users. The dominant social site will then have the unchecked power to disable almost any Web site or application by revoking its registration (in addition to the unchecked power of disrupting the online activity of any of its users by revoking his or her personal registration). This would eventually lead to calls for government regulation of the dominant social site in the US and other countries.

Third, current implementations of social login often provide a low level of assurance.

It is urgent to provide an alternative that avoids these drawbacks and dangers. NSTIC can be that alternative, but only if it can be deployed quickly.

3 Proposed Short Term Approach

The approach that we are proposing uses two ingredients:

One-time key pairs. To assert a user attribute, an attribute provider signs a certificate that binds the attribute to the public key component of a key pair generated by the user's browser that is used only once; key generation has no performance impact if the browser generates key pairs while idle and keeps them on hand. A one-time key pair is also used for delegating authorization.

Enhanced double redirection. Many existing protocols for delegating authentication use a mechanism where the relying party redirects the user's browser to the identity provider, which authenticates the user and redirects the browser back to the relying party. We propose to use an enhanced double-redirection mechanism for communication between an attribute provider and a relying party, where the broswer takes an active role, hiding the identity of the relying party from the attribute provider, generating one-time key pairs, and interacting with the user

to request permission; this will require an extension of the HTTP protocol. Enhanced double redirection will also be used for delegatied authorization and social login.

The approach encompasses attribute verification, delegated authorization, and social login.

Attribute verification is carried out by the following steps, where TLS (a.k.a. SSL) is used on all connections (PowerPoint slides illustrating these steps will be posted soon at
http://pomcor.com/documents/NSTICProtocolSteps.ppt):

Step 1. The relying party acts as initiator of an enhanced double redirection, redirecting the browser to the attribute provider and passing to the browser an attribute request and a callback URL.

Step 2. The browser forwards the attribute request to the attribute provider along with the public key component of a browser-generated one-time key pair. The browser retains the one-time private key and the callback URL. If the browser has a long term user certificate issued by the attribute provider, it uses it as TLS client certificate when connecting to the attribute provider.

Step 3. The attribute provider authenticates the user (e.g. using the long term TLS client certificate presented by the browser), signs a one-time certificate that binds the one-time public key to the requested attribute, and sends the one-time certificate to the browser, acting as responder in the enhanced double redirection.

Step 4. The browser asks the user for permission to communicate the attribute to the relying party.

Step 5. The browser completes the enhanced double-redirection by sending an HTTP request to the callback URL that it retained in step 2, using the one-time certificate obtained from the attribute provider as TLS client certificate, and using the corresponding one-time private key that it also retained in step 2.

The above protocol can be modified to provide delegated authorization instead of attribute verification. In this variation, the initiator of the enhanced double redirection is a Web application and the responder is a Web

site where the user has an account. The application requests access to the user's account. The application sends an access request and the public key component of a one-time key pair generated by the application. The site returns a one-time certificate that binds the one-time public key to a grant of access. As in the attribute verification case, the identity of the initiator, in this case the application, is not revealed to the responder, in this case the site.

The protocol can be further modified to implement social login, which is a combination of attribute verification and delegated authorization. The initiator is again an application, and the responder is a social site. The application sends both an attribute request and an access request. The application supplies a one-time public key that the social site binds to an access grant, and the browser supplies another one-time public key that the social site binds to an attribute. The attribute is a composite attribute that provides the application with the identity of the user relative to the social site and related data such as the user's full name, birthdate, avatar and/or email address. The access grant gives the application access to the user's account, so that it can, for example, issue updates on behalf of the user. As in the previous two cases, the identity of the initiator, in this case the application, is not revealed to the responder, in this case the social site.

4 Benefits and Shortcomings

The proposed approach achieves the privacy features of the Identity Ecosystem envisioned in [8]:

- In an attribute verification scenario, we have seen how the attribute provider does not learn the identity of the relying party to which the attribute is provided.
- The relying party asks for a specific attribute, and the attribute provider provides only that attribute, bound on-the-fly to a one-time public key in a one-time certificate, without revealing any extraneous information.
- Multiple one-time certificates cannot be linked, since they certify different one-time public keys.

The proposed approach also provides the following additional benefits:

- An identity provider existing today can be enhanced to provide one-time certificates of specific attributes derived from an existing multi-attribute certificate, or obtained from a database or directory. This makes it possible to leverage the identity investments made by the US government (CAC, PIV-I) and the private sector (OpenID providers).

- The user is authenticated by the attribute provider rather than by the relying party. This means that any relying party will be able to take advantage of the entire range of authentication technologies, including, e.g., multifactor authentication with biometrics, without investing in authentication technology.

- The approach encompasses not only attribute verification, but also delegated authorization and social login.

One shortcoming of the proposed approach is that the attribute provider in the attribute verification case has to be online. The browser can mitigate this drawback by requesting one-time certificates ahead of time, but that is a considerable complication.

From the point of view of privacy, the proposed approach has the following shortcoming: if the redirection initiator and responder collude, they can identify the user because both see the same one-time public key.[1] This means that, in the attribute verification case, the relying party (the initiator) can find out the identity of the user as known to the identity provider (the responder), and the identity provider can find out what transaction the user conducted with the relying party.

Notice, however, that this shortcoming does not matter in the social login case, since one goal of social login is to communicate to the Web application (the initiator) the identity of the user within the social site (the responder).

These drawbacks make us consider the proposed approach as a short-term interim approach, except perhaps for social login.

5 From the Short Term to the Long Term

La Macchia [3, slide 7] contemplates using ephemeral (one-time) public keys as the last link in a chain of zero-knowledge credentials, the motivation being

[1] This was pointed out to us by Kim Cameron at the Internet Identity Workshop.

that zero-knowldege technology is slower than ordinary public key cryptography.

This suggests that it may be possible to incorporate zero-knowledge credentials in a future version of TLS without modifying the core of the TLS handshake. If so there will be an easy upgrade path from the short term approach we are proposing to a long term approach based on zero-knowledge proofs: relying parties will be able to upgrade simply by using an upgraded version of their TLS libraries.

6 Conclusion

We have proposed an approach to achieving practically all the privacy goals of NSTIC in the short term. We have outlined a protocol that allows a user to submit to a relying party a certificate binding an attribute to a one-time public key, without revealing any extraneous knowledge to the relying party, and without revealing the identity of the relying party to the attribute provider. Different certificates cannot be linked to track the user because they are bound to different one-time public keys. Collusion between the relying party and the identity provider can reveal the user's identity to both; but an important extension of the protocol for social login is not affected by this shortcoming.

The approach can preserve and leverage the existing investments in Internet identity by the US government and the private sector, providing a smooth transition from existing technology to privacy-enhanced technology. The proposed approach may also facilitate a smooth transition in the future to a long term solution based on zero-knowledge proof technology.

References

[1] Jeremy Grant. National Strategy for Trusted Identities in Cyberspace, April 6, 2011. Presentation at NIST IDtrust 2011, available at http://middleware.internet2.edu/idtrust/2011/slides/03-national-strategy-trusted-identities-cyberspace-nstic-grant.pptx.

[2] Janrain (we believe that the term social login was coined by Janrain). Available at http://www.janrain.com/products/engage/social-login.

[3] Brian A. LaMacchia. New results using anonymous credentials: Constrained delegation and revocation, April 2011. Presentation at NIST IDtrust 2011, available at http://middleware.internet2.edu/idtrust/2011/slides/02-privacy-lamacchia.pptx.

[4] NIST. Administration Releases Strategy to Protect Online Consumers and Support Innovation and Fact Sheet on National Strategy for Trusted Identities in Cyberspace. Available at http://www.nist.gov/public_affairs/releases/whitehouse_nstic.cfm.

[5] NIST. Web site of the National Strategy for Trusted Identities in Cyberspace. Available at http://www.nist.gov/nstic/.

[6] OAuth Working Group. Web Page of the OAuth Working Group of the IETF. Available at http://datatracker.ietf.org/wg/oauth/charter/.

[7] OpenID Foundation. OpenID Authentication 2.0 Final, December 5, 2007. Available at http://openid.net/specs/openid-authentication-2_0.html.

[8] Howard A. Schmidt. The National Strategy for Trusted Identities in Cyberspace and Your Privacy, April 26, 2011. White House blog post, available at http://www.whitehouse.gov/blog/2011/04/26/national-strategy-trusted-identities-cyberspace-and-your-privacy.