Do you know who I am?

For the W3C Identity Workshop, Mountain View, May 2011

David Singer

Multimedia and Software Standards, Apple Inc.

1 Do you know who I am?

We all log on to multiple web sites, often many times. Each web site asks our user-id; sometimes it is created by the service, sometimes we choose, and sometimes they ask for, and use, our email address. Each web site wants, or makes, a password. None of us can remember all those IDs and passwords, and we have various ways to cope.

This leads to a ‘splintered’ identity on the internet. Services – particularly advertising services – sometimes work hard to correlate user identifiers, as this enables them to know more about the users and therefore target advertisements better. Technology providers fret that asking users for these multiple identities is inconvenient and error-prone. Wouldn’t it be easier if they had ‘one identity’? OpenID is one response to this.

But is ‘one identity’ desirable either for the user or the service provider? For the service provider, it can lead to undesirable or unforeseen problems. A prominent one is that if I rely on another web-site to verify my users’ identities, then I am at the mercy of that site’s ability to verify – if it goes down, or the business-to-business link enabling that verification stops working, users can no longer log on to my site; that may affect my business. Similarly, users worry that if their identity becomes compromised or hacked, that means that all their accounts are now vulnerable.

2 Do you know to whom you are talking?

‘Phishing’ is a significant issue for online services. Rogue sites emulate real sites and lure users into logging on, and then use those credentials to compromise their accounts. We have thought for a while that there is not much we can do, apart from education, to protect users from this; if they are fooled by a plausible-looking site, what can be done?

However, mediating the discussion is a user-agent; an agent is expected to act on behalf of, and in the best interests, of its user. In this case, it’s worth asking whether the user-agent can be much more involved in assuring both that the site is talking to the real user who owns the account, and that the user is talking to the real site where they have an account.

It used to be the case that users logged on using HTTP authentication – ‘in the chrome’. That is getting rarer and rarer; services prefer to offer a page, with hints, context, and so on, to hand, for example. Unfortunately, this makes it harder for the user-agent to know when filling in a form is actually the user self-identifying, with an implied need for site identification/verification. The use of a password form input is a clue, as are the use of security credentials.

3 Mutual Identification

Do we need a new system that ‘replaces’ HTTP authentication, and enables the user-agent (a) to help the user self-identify without having to remember every facet of their fractured identity and (b) enables the users to trust that they are talking to the site they intend? Once the user-agent knows that a mutual identification is in process, there is much that they can do; “you seem to be logging on to a site for the first time from here, let’s verify that the site is the one you think it is”, for example.

Such an ability – not a single identity, but a way of noticing identity transactions – may allow the users to maintain their ‘splintered’ identity, and sites to retain their autonomy, while letting the user-agent ameliorate the problematic aspects – ‘who am I on this site’, ‘is this site the real site’, and so on.

4 Closing questions

Do we need ‘single identity’ or ‘identity management’, or do we need to manage ‘identity transactions’ better?