The Nexus of Identity

IBM's submission for the W3C “Identity in the Browser” Workshop

by Maryann Hondo, Mary Ellen Zurko, Matthew Flaherty, Paula K. Austel, Sridhar Muppidi

IBM's customers commonly look to us to provide guidance and practice as they establish or extend their web presence. It is not uncommon that many of their questions are around the best way to secure, manage, and share federated identities. From RACF to web & user security, followed by Java, web services, SOA, and cloud IBM has been committed to its customers and to the community (this is our 100th year) throughout the industry technology cycles. Our Tivoli brand provides a comprehensive set of security products for customers with a range of identity federation challenges but every area of IBM is aware of “digital identity”. It is one of the basic elements of internet access today. It's part of our DNA.

Reliance on web technology, both inside and outside the enterprise has moved beyond a set of "pages" that collect a few pieces of information to sophisticated applications that collect, store and share detailed information on individuals. HTML 5, CSS and JavaScript demonstrate exciting futures for application development. Helping our customers establish clear terms of service, and define and manage polices to protect user and data privacy is as important as helping them choose the technology to share the information. Protecting a range of information while still enabling the exchange of that information across many topologies is indeed a challenge and an ongoing one with new social media and cloud strategies. To complicate matters, identity information can be required across company boundaries or prohibited across boundaries.

Businesses want to safely leverage existing identities and take advantage of social network data but to enable this they need technologies that allow them to establish trust relationships with other identity providers as well as establish a range of trust within company applications. Enterprises have policies on everything, from authenticators, to sessions, to access of company owned data. They have a notion of identities "inside" the enterprise, built up around years of managing those identities. Those identities each have a contractual commitment to the enterprise. Enterprises also have a notion of identities "outside" of the enterprise, and sometimes some interesting "in between" cases. For example, contractors have a contractual relationship; sometimes enterprises want to treat them like employees, and sometimes they have special conditions and policies that need to be applied. In some cases, the relationship, such as contractor, is baked into the naming scheme, as a form of rudimentary transparency. Partners also have a contractual relationship, though they are certainly primarily affiliated with their own places of business. Sometimes enterprises want to control the identities of people from other enterprises (or companies, or organizations) when they work together—or at least know more about those identities (who vouches for them, with what assurance).

Through all of this, two themes of concern emerge amongst our customers. One theme is ***control***. Who controls the format of an identity, authentication, and assurance. Another theme is ***transparency***. How do the attributes of an identity relate to the requirements of the context it's used in (the enterprise, the web site, the other identities it interacts with). The identity indicator needs to be in a format that is user understandable and meaningful. Time and again there is the question of uniqueness, understandability and disambiguation. While there is (still) only one "Mary Ellen Zurko" on the web, there are several Thomas Roessler's (and two of them actually work in security). Only one person uses mzurko@us.ibm.com for email, although there are attack scenarios and spoofing attempts. Other attributes can be of interest to users and enterprises digitally interacting with identities, including identity provider, authentication strength (including multi-factor).

Our practical approach to user authentication within IBM Lotus works between the enterprise and the LotusLive web sites. We are using SAML which is widely available in enterprise grade form, and can be made to work with a variety of user agents with minimal end user disruption and interactions. In use cases where a user wants to delegate authorization for one service provider to access data at another service provider we use the OAuth protocol for authentication and authorization. Our customers are interested in additional controls that increase security and do not increase the overhead on their users. IP range restrictions that can ensure authentication from within the enterprise are of interest, since they provide another defense against external phishing attacks.

A browser ***only*** solution, which the workshop posits, is problematic for an enterprise that uses multiple user agents. It is not clear what scenarios balance the competing interests (enterprise control of their authentication policy, their naming policy, a diversity of protocols and user agents) with the change that an additional approach to digital identity management would bring.

Yet, the success of the browser on mobile devices and smart phones highlights an ever increasing need to have working solutions that include the browser as a user agent providing both transparency and control. While a simple goal, the coexistence of different compliance and policy requirements as well as a heterogeneous set of platforms makes the problem not always so simple to execute. Mobile devices are not simply extensions of the traditional desktop. These devices are shared among multiple users, are easily lost, stolen, or can have stored credentials extracted from the device. While these are not new issues, the anecdotal evidence is that enterprises have greater concerns than other environments about mobile web security, including identity. One important issue is around authentication and of the identify of the person using the mobile device. Traditional authentication to an application typically takes place once at the start of a session. With mobile web, there is certainty that the authenticating user is the the same person during the lifetime of a session.

As IBM developers and web users ourselves, we can observe that people often have multiple identities (personas) across their enterprise and the Internet, for example an identity and profile within their company that is meaningful and intended for colleagues and business partners, an identity and profile for their business personality in a business oriented social network, and an identity and profile for interaction with friends and family. Furthermore, in some countries citizens can obtain personal IDs from the government that represent their identity as a citizen of the country with attributes as in their passport, with the ability to provide smart card based authentication and signatures strong enough for elections and contracts. Being a global company we understand many of the challenges companies face, because we face them ourselves.

We believe that as a result of people having multiple identities, it will become increasingly important to not only enable leveraging existing identities to sign on to web sites and services, but to also advance how to select aspects of a profile in a particular social network (e.g. from in-company social network) so that they may be surfaced outside of that network (e.g. to a business oriented external social network), how to associate and correlate multiple identities of the same person when desired, how to allow people to express which of their identities they intend to use as primary personal vs primary business point of contact, etc.

The reality of the environment in which digital identities are exchanged, is that there are intermediaries or proxies which play a critical role in establishing various types of trust relationships between the front end client and the back end service. Our customers may want to authenticate an identify, perform a credential mapping, or they may augment the identity requiring us to support both identify propagation and delegation. To establish trust, the browser acting as the *user agent* needs to securely protect the user identity when the identity is in store or in flight. To do this well an intermediary needs to process many protocol patterns and token types (ex. client certificate over SSL, basic authentication, or token based identity in the http header). Tokens in http headers are used extensively to provided SSO and standards like SAML provide a step toward interoperability but sometimes the size of the token and the processing required can impact network performance.

We support the W3C initiating a workshop around digital identity and think it is a good time for the industry to take stock of the existing patterns to determine if this provides us a sufficient platform for the next decade.

We would love to bring our wide range of experience (from policy to research to product development) to the workshop and work with our customers and the community at large to make sure we all have the right technology to face the challenges ahead.