Goals, Constraints, and Issues for Identity in the Browser

John Linn, Sr. Technologist, Office of the CTO, RSA, The Security Division of EMC, Bedford, MA, US

john.linn@rsa.com

Position paper for W3C Identity in the Browser Workshop, Mountain View, CA, May 2011

8 April 2011

Introduction and Scope

As a provider of authentication and identity management technologies, RSA, The Security Division of EMC is interested in supporting the security and privacy of browser users, present and future. We recognize the relevance of browser-based enhancements to this goal, and encourage their development. Rather than recommending particular methods to be implemented by browser providers and their peer sites, however, we focus here on suggesting common characteristics for consideration as methods are designed. To this end, this paper proposes candidate constraints and goals for browser-based identity mechanisms, in the interests of encouraging their utility, generality, and value.

Today, browsers act primarily as conduits for users' interactions with the Internet rather than as active agent entities serving on behalf of their users. From a security perspective, they support protected TLS channels to encapsulate traffic, maintain cookies that may manage authentication state, participate as peers in HTTP authentication when used, and can cache password credentials for convenience, but rarely take more active or comprehensive roles in representing and securing user identities. If and as users can delegate aspects of their digital identities to be managed by trustworthy browser-based modules, new modes of interaction can be supported, while maintaining or enhancing privacy and security. Browsers can participate in authentication and other protocols in ways that human principals cannot, can inform their users about characteristics of the sites they access, and can enforce policies on behalf of those users. On the other hand, delegation to untrustworthy modules can introduce new risks and/or amplify those already present. Browser design characteristics may restrict their users in unnecessary or undesirable ways.

When a browser serves as a user's agent, representing the user to the Internet and representing the Internet to the user, it exposes a user-facing human interface on one side and a network-facing protocol interface on the other. Further, it makes use of interfaces to its underlying platform and the services it provides, and may invoke other distributed services. Security-relevant prospects and potential concerns arise across each of these interfaces. Method designs must enable implementations that are flexible, secure, deployable, and usable. Subsequent sections will discuss aspects related to these areas.

User-facing concerns

Users must be able to direct what aspects of their identities are to be shared, and with which recipients. Privacy controls, including anonymity and/or pseudonymity support as appropriate, are relevant. As an important implication and special case, users must be able to control where and how their credentials

are to be applied, either on a per-operation basis or by establishing a policy to enable and govern autonomous credential use. Trusted path facilities are needed, to prevent users from being induced to provide credentials to attackers. In designing the management interfaces that are to be presented, usability considerations will be particularly important. Policy controls have often been difficult for users to understand or manage, and the security implications of delegation to trusted browser modules are significant. Similarly, the security and privacy views that browsers present to their users to describe site attributes should be both informative and usable.

Browser-based capabilities offer valuable potential to strengthen Internet authentication practice. As a goal, we recommend that the human-facing authentication interface and method (where human factors are primary concerns) should be appropriately decoupled from the network-facing authentication transactions that the browser enables. To apply browser-based credentials securely across the Internet, the browser user must first be authenticated locally, either by the browser itself or by its underlying platform. For either case, it's important for a browser-based method's design to minimize dependencies and assumptions that tie it to particular user authentication mechanisms, leveraging the strengths of those mechanisms rather than constraining their selection. The US NIST's Draft SP800-63-1, Electronic Authentication Guideline, identifies a set of authentication token types with different capabilities, attributes, and presenting different browser interface characteristics. We suggest that this set be considered for support as browser-based methods are designed.

Network-facing concerns

Today's World Wide Web comprises a uniquely pervasive base of installed technology. If and as use of new browser-based capabilities requires corresponding enhancements to servers, the changes will become useful only to the extent that they are deployed on both sides. Unless any protocol enhancements are either (a) beneficial to server operators, (b) strongly desired by users, to the extent that their absence may comprise a popular disincentive to access a particular site, and/or (c) mandated through regulation, their adoption should not be assumed with confidence. In any case, new features are likely to be supported gradually rather than becoming pervasive within a short time. It is desirable, therefore, for new browser-based methods to be able to offer useful added value within the context of existing protocols, even if greater benefit can be offered when and where extended protocol support becomes available.

Profusion of large numbers of different, incompatible protocol features can lead to complexity, confusion, and interoperability limitations. Consensus-based standardization will be important as a means to incorporate new capabilities gracefully and compatibly into the protocol base.

Browser and platform concerns

As a primary security goal, browser-based identity mechanisms should not increase the likelihood that a user's credentials will be intercepted by an attacker, or applied in a manner that is inconsistent with the user's intent. To achieve this goal, it is important for browsers, their underlying platforms, and any distributed services that may impact browser-based security functions to operate in a trustworthy fashion on behalf of the users they serve. As persistent storage of users' credentials can increase their

exposure and vulnerability, it is useful for implementations to support means to ensure that the credentials are not accessible in usable form except when a user's session is active.

Conflicts may arise among browser-maintained credentials, dependencies on underlying platform and hardware characteristics, and cross-client mobility. Desirably, users should be able to move easily and effectively among browser and platform instances. This motivates the potential value of storing credential data and other user information on servers which can be accessed conveniently from multiple clients. As noted above, however, such an approach may introduce requirements for trustworthy server design and operation, broadening the associated security perimeter beyond the browser itself. Design strategies to minimize the necessary level of server trust should be considered. As another possible alternative, though perhaps less conveniently, user data objects may be transferable directly between clients via user-mediated operations.

Conclusions

Browser-based identity mechanisms can aid users by supporting Web interactions that are highly usable while strengthening security and privacy, if the mechanisms are suitably trustworthy. On the other hand, mechanisms that cannot be trusted or that restrict user behavior can limit users' security or privacy and/or may adversely impact the experience they receive. New methods must coexist with existing practice, and must demonstrate significant value in order to motivate general adoption. We look forward to workshop discussion to clarify and refine the capabilities and characteristics that different approaches can offer.