

Improving Digital Credential Management in Browsers

David L. Wasley, Editor¹

July 2000

Use of the web for trustworthy business applications requires use of strong digital credentials² supporting the identity³ and authorizations⁴ of the individual transactor. Such credentials can be created and validated by means of a well designed and deployed Public Key Infrastructure (PKI). This is the strategy chosen by both higher education and the Federal government for new business applications, both internally and with external partners.

Two predominant web browsers - Netscape Communicator and Microsoft Internet Explorer - will be the client tools used by most transactors. Each browser has some degree of support for PKI credentials but neither has sufficient flexibility and ease of use to be truly viable, especially for less technically adept individuals. Vendors of these browsers and developers of other newer web clients must incorporate far better support for sophisticated and user-friendly use of these digital credentials.

To that end, the higher education community through its professional IT activities, offers suggestions below for improving existing browser functionality. These are intended to be starting points for discussion, not specific requirements. They are intended to lead towards desktop clients and more sophisticated servers that can support “user friendly” and secure, robust electronic commerce environments.

Assumptions

A number of practical considerations underlie the recommendations to follow. These include:

- 1) *Most individuals will require or desire a number of different digital credentials. Just as we have different forms of identity and authorization in traditional commerce (authorization lists, drivers' licenses, credit cards, passports, checks, employee ID, ...) we will have multiple digital credentials. For example, one might have a personal credential, one or more functional or role based credentials, short lived temporary credentials, group membership credentials, and so on.*

Specific industries might well have need for unique credentials. For example Banking might have different requirements than Health Care. Libraries will want credentials that protect intellectual freedom by not mapping (directly) to a known individual. Other on-line resource providers may want credentials that assert eligibility of the holder under terms of a specific contract.

Equally important, individuals may not be willing to tolerate the potential invasion of privacy that a single credential would constitute if it were to be used for identification in all network-based activities in which the holder engaged. However, some individuals may prefer fewer credentials and be willing to aggregate functionality at the cost of privacy. This scenario must also be recognized.

- 2) *Most users of the browsers will be non-technical or at least technology-challenged. The typical user must not be confronted with technical jargon or complex choices to make with respect to security issues. The user security interface in browsers today is a serious impediment to wide spread acceptance and use of PKI credentials.*
- 3) *The PKI credential will contain very little of the holder's actual identity attributes. The primary personal digital credential is intended to be a long lived binding of a digital token to a known individual. If it were to contain very much about an individual, it would become invalid relatively quickly. Furthermore, the attributes of "identity" that are useful will vary from application to application. "Common name" for example indicates nothing of a person's roles or authority. However, roles and authority change often enough that they should not be included in a long lived credential. Therefore, the vast majority of identity attributes will be held in associated secure LDAP databases and/or directories.*
- 4) *Access to resources will be based on algorithmic business rules. In a complex environment, a user may need to gain access to a wide variety of resources, each with its own access management rules. These should be automated so that the user need never be bothered with excessive login prompts, choices, verifications, etc. The PKI credential environment will include secure LDAP directories that will hold the individual attributes necessary to support the algorithmic automation of access rules.*
- 5) *Browsers will not be the only client application using PKI certificates. Secure email, for example, or SSH applications or even multiple browsers should be able to use the same set of PKI credentials. These credentials represent a collection of important authentication and other information that is not specific to an application but to the individual to whom they apply. PKI credentials should be cached at the system level and be available to all applications through a standard API.*
- 6) *Users will need to invoke the same digital credential on more than one computer, and some of these will be multi-user or public computers. Students using public computer labs or special purpose systems will need to be able to use their PKI digital credentials without the risk of leaving*

traces of them on the platform after they depart. Individuals who use multiple computers routinely, for example telecommuters, should be able to make use of the same set of credentials with all their platforms.

- 7) *Trust structures will evolve and become far more complex. Simple two level, single root trust hierarchies simply do not reflect how trust works in the “real world”. Instead there will be lots of smaller, multilevel trust hierarchies with bilateral agreements both with bridge CAs and directly with other hierarchies. Policy reconciliation at the bridges is an unsolved problem but must be accommodated, ultimately.*
- 8) *Networks will continue to become more robust, ubiquitous, and of higher capacity. Today 100 Mb/s to the desktop is taken for granted and inter-institutional data capacities range from 45 to 600 Mb/s and continue to grow. Critical networks are made robust through redundant connectivity and engineering to avoid single points of failure. Therefore, distributed systems, using resources at several different network nodes, can be relied upon for mission critical applications.*

Finally, we assume that not all conventions, protocols, and hardware exist today to achieve the level of convenience and flexibility that we believe is required. We are anxious to participate in development of these added standards. However they are developed, it is essential that they be freely available to the industry and not be proprietary to any one developer.

Recommendations

The following set of recommendations is neither exhaustive nor diagnostic. Additional functionality may be required once more experience has been gained with this technology. Additional misfeatures may be found as well.

The challenge for browser developers is to provide ways for the technology-challenged user to manage a fairly complex set of security and authentication features without undermining their control or the flexibility of these features.

These recommendations are not in any particular order.

1. Thoroughly document the browser’s use of PKI credentials

Documentation, if any, is sketchy and often difficult for non-technical users to understand. The security and trust models and related assumptions must be clearly documented. There should be both easy to understand, basic descriptions and detailed documentation for technical support staff and adept users.

2. Clarify dialog boxes to avoid confusion

Some browsers or browser dialogues are sufficiently ambiguous that novice users can become confused. Passwords have been known to be entered in “userID” boxes, for example, because at a quick glance the user was not certain what was being asked for.

3. SSL negotiation and user privacy

During the SSL handshake, the server may be configured to send a certificate request to the client. This is only a request; the SSL connection might still be established if the client has no certificates. However, if the client does have one or more certificates, either one is sent automatically or the user is asked to select one. The user is not given the option of declining to send any. In order to allow the user to defend her or his privacy, there should be a “none of the above” option which tells the browser to send no certificate even if one is requested. In the same vein, there should be a way for a user to specify specifically which of their certificates may be used for SSL purposes (see also item 8 below).

4. Accepting Server certificates

When a browser doesn't already “trust” a server's CA, it asks the user to approve use of the server's certificate. (Some browsers simply don't allow unrecognized certificates to be accepted at all.) This dialogue could be avoided if the browser was able to discover a trust path or chain between the CA that issued the certificate and at least one of the CAs in the user's personal certificates. The credential is merely an assertion of authenticity vouched for by a trusted authority. Acceptance of a server's credential does not imply trust of the server itself but merely of its identity. The user has already indicated her or his desire to gain access to the server by selecting its URL.

If no trust chain can be created, the user should see a simple explanation of the choice she or he must make and be required to see no more than one or two screens in order to accept or deny the server's certificate. As a last resort, organizations that wish to do so must be able to add a server's root CA's certificate to the browser's cache as an external configuration file.

In the meantime, it must be easier for common root CAs to be configured into the browsers as distributed. For example, the Corporation for Research and Education Networking (CREN) is supporting a root CA for the higher education community in the US. It should be reasonably straight forward to add the CREN root CA to the list of trusted CAs that comes with the standard browser distributions.

5. Trust path discovery, including bridge CA's

Many trust hierarchies will be deeper than 2 levels and there will likely be branches through multiple bridge CAs and ad hoc bilateral agreements. Both servers and clients will need to validate PKI certificates using the X.509 open standards policy processing rules through moderately long trust chains, which may have policy-based restrictions along the way. Local CAs might mediate trust chain discovery for their clients, or special servers which focus on trust chain creation may be developed, but if neither is done, ultimately the client may have to discover the chain independently.

6. Dual certificates: signing vs. encryption

Both digital signatures and data encryption are essential for a robust business environment. By the same token, the decryption (private) key that must be used to recover an encrypted document often must be escrowed in order to ensure business continuity, while the digital signature (private) key of a transactor must never leave control of the individual holder or else any digitally signed transactions made using that key become repudiable. Therefore, browsers must support dual certificates and respect the appropriate use of signing versus encryption key pairs. Again, this requirement is in conformance with the X.509 standard processing rules.

7. Signing and encryption

Browsers must make the user interface for signing and/or encrypting information simple and intuitive. For example, browsers used for financial applications or e-commerce should have a "button" for signing the information input by the user so that any resulting transaction can be validated at the time as well as audited later. Encryption at the transport layer (SSL) may not be adequate to secure sensitive information once it is received and auditability requires archival verification of the identity of the specific transactor.

8. Managing multiple PKI credentials (private keys and certificates)

One of the most difficult problems to be solved is making the use of multiple PKI credentials held by the user as transparent as possible to that user. It is absolutely inevitable that users will have more than one PKI credential (possibly a great many more) and it is critical to minimize the user overhead in deciding which to use in a given situation.

A new set of standards is needed whereby a server can ask for a certificate of a certain type or with certain attributes, but only for certificates which are

allowable by some overall aegis. For example, an on-line publisher's server should be able to ask for a certificate that asserts the holder's eligibility to gain access to their holdings but should not be able to ask for a certificate that asserts eligibility to gain access to a different publisher's content.

This could be accomplished with a combination of globally defined certificate extensions and/or external information in the user's certificate caches so that the appropriate credential can be chosen automatically by the browser. In the near term, user's might manually define preferences at the time a certificate is issued. The ultimate goal is to minimize the user's involvement in every access control dialogue while retaining appropriate levels of privacy, security, and trust decisions.

9. Certificates on demand (e.g. anonymous)

Many servers and applications need only associate a credential with a class of individuals eligible to make use of the service. For example, access to on-line library materials may be granted to any member of a campus community and the specific identity of the individual may be inappropriate to reveal during the process of granting access. Therefore, users will have (or be able to get) "anonymous" or other special purpose certificates "on demand" from their primary CA (the CA that issued their personal identity credential).

This requirement is strongly linked to the issue of managing multiple certificates. In a scenario where a server asks for a certificate type that the user does not have, the browser should automatically ask the user's primary CA or perhaps a special purpose authorization server for such a certificate and add it to the user's cache.

For example, a McGraw-Hill server might ask for a certificate asserting eligibility under the terms of the holder's campus contract with McGraw-Hill. If the user's cache does not contain such a certificate, the browser should ask the campus' CA or perhaps a campus library authorization server for one on behalf of the user, using the user's primary identity credential to validate the request. The CA or server would check the user's attributes and either issue the appropriate eligibility credential or deny the request.

Electronic voting is a special case of this. What the user sees is a web page with a description of the election issues, etc., and a button to select when the user is ready to vote. Then, behind the scene, the user's browser asks the ballot server for a copy of the ballot and a certificate asserting eligibility to "vote". The user is presented with the ballot and radio buttons to select their preferences. When complete, the browser uses the voting certificate to submit the completed and encrypted ballot to the "ballot box server". The

ballot box only allows one instance of that certificate to be used and has no way of (directly) associating the user with the ballot. What the user sees is simple and intuitive; what the browser does is fairly sophisticated.

10. Using private keys and certificates for multiple applications

The browser will not be the only application using private keys and PKI certificates. For example, secure email clients and secure “telnet” require the use of public/private key pairs and/or certificates. It can not be assumed that a single client package will or should provide the complete user interface to managed services.

Therefore, PKI credential caches (which may include both private keys and PKI certificates) should be managed at the operating system level with a standard API for all cache functions. In preparation for this, browsers should be designed to be able to easily take advantage of such system services when they exist.

For the near term, import and export of PKI credentials between the browser and a local encrypted cache should be simple and reliable.

11. Portability/Roaming: Smartcards, USB dongles, PKCS11+

People use more than one computer and in some cases use shared computers. The PKI credentials that they need are associated with the individual, not with the particular computer they happen to be using. Therefore, there must be a convenient and secure way for people to make their PKI credentials available on any platform they use.

Smartcard technology seems promising for this but there are several impediments to be overcome. First is the hardware interface; second is the capacity and functionality of the smartcard; third is cost. It is reasonable to assume that these can be overcome in the near future.

Related to the need for standard operating system support for PKI credentials, the cache managed by the operating system should be that which is stored on the user’s portable smartcard. Until this is widely implemented, the browser should do the same: look for and utilize a smartcard-like object if it is connected to the user’s platform. In no case may a browser ever copy a user’s private key from the smartcard into a local cache!

Until smartcard technology is widely available, browsers should be able to support standard, secure, PKI credential servers. Such servers hold the user’s private key, double encrypted, and the associated PKI certificate. Only identity

credentials, possibly to be used also for encryption, may be served in this way; digital signing credentials may not be served because to do so would compromise the veracity of the signing (private) key.

When the user needs her or his PKI credential, the stored object is downloaded (over SSL) to the user's current platform. In use, the private key is unlocked only at the time it is needed for certificate signing or data decryption. The certificate signing is done in the computer's memory, and then the unlocked private key copy is discarded leaving only the encrypted copy on the platform. Neither the private key nor the PKI certificate is ever stored on disk. Both should be discarded entirely from the browser's memory after a suitable timeout period. This approach allows roaming without smartcards, but is obviously not as secure as smartcards since the crypto functions are done in memory rather than on the card.

12. Servers must recognize Attribute Certificates

Finally, servers must be smarter as well in their use of PKI digital credentials. Digital credentials will not include all possible identity attributes. Instead, they will refer to directories or databases where such attributes can be found. There should be standards for commonly needed attribute object classes and configurable rules based processes for using these attributes in access management decisions. Careful adherence to such standards will help minimize the unnecessary proliferation of the number of certificates a user must hold.

Getting it Done

We realize that the requirements above are non-trivial to implement (although some may already be in the pipeline). The higher education community is both ready to help with this development and a natural proving ground for the results. We look forward to working closely with industry to achieve this important next step towards a robust and easy to use e-business environment.

¹ Contributors to this document included

Judith Boettcher, CREN

Michael Gettes, Georgetown University

Clair Goldsmith, University of Alabama and Chair, EDUCAUSE Net@EDU PKI Work Group

Richard Guida, Department of the Treasury

Ken Klingenstein, University of Colorado and Chair, Internet2 Middleware Work Group

Clifford Lynch, CNI

Eric Norman, University of Wisconsin

David Wasley, University of California, Office of the President

² “Digital credential” refers to the combination of PKI certificate and private key. The credential is something that can be validated independently by a receiver. The content of a credential certificate may include a digital token that is unique to the holder of the credential. It is this token that allows retrieval of additional attributes relevant to the holder. This form of credential is therefore a personal or identity credential. Other forms of digital credentials may assert specific eligibility for services or membership in a class of non-specific individuals.

³ “Identity” is the set of attributes pertaining to an individual entity. These include such things as common name, addresses, roles, status within a context, authority, etc.

⁴ “Authorization” is a process followed by a service provider to determine whether a requestor is to be allowed to complete a transaction or gain access to a resource. Authorization may be based on some set of the requestor’s attributes, location of the access node, time of day, availability of resource capacity, or any combination of these and other things.