Web Authentication Position Paper
– Bank of America

In response to the Dec 15, 2005 Call for Participation in the March 2006 meeting on "Toward a More Secure Web - W3C Workshop on Transparency and Usability of Web Authentication," Bank of America offers the following comments.

Introduction

Bank of America has a strong interest in Web authentication. As one of the world's largest financial institutions, we already serve tens of millions of customers and business clients via the Internet, and we view the online environment as an important source of potential growth.

For us to realize that potential, however, the online environment will have to be more secure than it is today. To drive growth, we must build an online environment that earns the ongoing trust of customers and clients. Customer trust is a key driver of all businesses, especially financial services businesses.

Users expect online financial and commercial services to be secure. But they often do not understand the limitations in today's computing environment, and as a result they have unrealistic assumptions about the ability of any business to secure transactions and information. Their lack of awareness often also contributes to the security problems that adversely impact them.

We believe three issues stand out if we are to move towards higher levels of customer acceptance and use of the Internet as a transacting commercial medium, and thereby to realize its growth potential for us and all other commercial enterprises. The three issues are:

- Trustworthiness;
- Responsive human interfaces; and
- Reputation services.

Underlying all of these issues is the need for a broad industry adoption of solutions. Most companies doing business on the Internet rely on third parties to provide the network and software infrastructure that supports business-to-business, business-to-consumer, and consumer-to-consumer transactions. Since these transactions must be secure, it stands to reason that the standards should be universal.

This paper will look briefly at all three of issues, offering some comments and initial suggestions.

Trustworthiness

At one time, most users of e-mail trusted it implicitly. But events—and media sensationalizing—have combined to make e-mail far less trustworthy. Seasoned and inexperienced users alike can no longer be certain that a message in their inbox is what it claims to be, or is even from whom it claims to be from. We certainly know not to trust any file attachment. In short, e-mail has lost much of its trustworthiness, and thereby much of its effectiveness.

The same is increasingly true of other online environments. Are the correct applications running on your computer? Is your WiFi connected to a network you can trust? Is the Web site you are on *really* the Web site you're on? Users buy computers trusting the software installed by the manufacturer is good software. They expect that applications work as designed (and promoted).

In the physical world, we trust what we see because of our experiences, the rules of law and physics. In cyberspace, our experience is mediated by code, and things don't always behave as we expect. The code is simply not as trustworthy as the laws of physics or the regulations of everyday life.

Yet trust is implied across today's computing environment, with little basis for that trust (see Ken Thompson's comments at http://www.acm.org/classics/sep95/).Viruses, Trojan Horses, Phishing attacks, DNS attacks, Bot networks, and a plethora of other malicious code demonstrate that users rarely know exactly what programs are running on their machines. Today, there are few limits on what code can do, and who has the right to create and execute code on a user's machine. Malicious software routinely installs itself without asking permission.

But even code that asks permission can be bad for the user if the user lacks the knowledge to make a wise, informed decision. Again the user is reduced to trust.

To cope strategically, we need systems that allow multiple values of trust—trustworthiness systems that parallel how we manage human and business relationships every day. This multi-value trust mechanism should be automated within both application and system, able to assess and evaluate the trustworthiness of data (and code), to enforce the separation of fully trusted data (and code) from semi- or untrusted data (and code), and to deal with data (and code) by referring to its trustworthiness level. Systems need to be able to execute fully trusted programs, while running untrusted applications in sandboxes like Java.

Today, we are already seeing the beginnings of "trusted" and "untrusted" in this sense; policy management for sandboxed Java applications and signed ActiveX controls are examples. But ultimately we will need a more refined multi-value system. Near term, we need to improve consumer understanding of their systems. However, relying on user education alone will be a slow, tedious, and probably unsatisfactory process. Users simply cannot compete with hackers' speed and agility at coming up with new attacks. *Systems and applications must automate complex trust decisions* and be resilient to poor user choices.

For example, unsigned ActiveX controls from www.badguysr0x.com should, by default, not have access to any financial or commercial applications and directories. So even if a user downloads bad code, the system defaults to safety rather than breach. Under such a model, we can live with the users' decisions, and we know their decisions will improve as they are better able to survive the bad decisions. Think of this as computing Darwinism: If the first attack kills you, then you are not suited for the wild. If you survive, you learn and get stronger.

How does this multi-value trustworthiness model get implemented? Industry groups and responsible companies need to drive solutions, creating products with built-in mechanisms to assess, determine, and enforce trust. These responsible designs must help users make good trust decisions—without telling explicitly telling users whom to trust. Software development teams must be managed to ensure the trustworthiness of their development process and the final code they produce.

Users, meanwhile, will need systems that verify (to some degree) the source of all software and allow them to make reasonable judgments about the software running on their machines. When users see a Web site presented in a browser display, they should be able to verify that what they see is really what they've got. They should know the real owner of the Web site. They should know whether data they enter in a form will remain private. They should know whether data they received was protected. And above all else, they should have some indication of the trustworthiness of the Web site itself.

Human Interface

A critical component of that trustworthiness model is a reasonable and understandable human interface, and it is there we now turn our attention.

Does users really see what they have? Do they have what they see? And do they know the difference?

Research at several universities shows that there’s a dichotomy regarding whether people trust what they see. They trust a site displaying logos from normally trusted services[1], even when a toolbar, which they also see, tells them the site is bad! The latest browsers appear to be tackling this issue, but several other issues need to be addressed, including adoption and user training.

There is some concern about the whole logo-icon approach. E-commerce businesses have been telling consumers for 10 years to look for the “lock icon” at sites like Amazon, eBay and various banks. And consumers *still* enter personal information, passwords and user IDs at malicious web sites. A real human-accepted interface may require something like the browser putting a watermark over web page content, or the browser filtering out displays of trusted logos like the Verisign Logo, Trust-E Logo, Better Business Bureau and similar services from Web sites not really authorized to use those logos. So, for example, a Trust-E (http://www.truste.org/) logo wouldn’t be displayed by a browser unless the browser had determined that the Web site was officially partnered with Trust-E. A spoofed Web site showing obvious gaps, or even explicit warning graphics, where a trusted logo should appear would have a significant human interface impact.

Reputation Services

Interface impacts would be greatly muted, however, without comprehensive management of reputation guarantors. Trust-E, the Better Business Bureau and other companies represent an early type of reputation service, but, as we’ve noted, these models are fairly limited—anyone can copy the Trust-E logo and place it on their Web sites. Only an alert and rare user will go to Trust-E to verify that the Web site has earned the right to display the logo.

To augment the existing reputation services, there are alternatives:

Domain naming: Originally, the domain name (.com, .net, .org) formed a sort of early reputation service, indicating a company’s business model. There’s a longstanding but never implemented proposal to create a domain specifically for pornography, to help consumers avoid accidentally landing on porn sites. Perhaps there is potential for building reputation guarantees into the naming system itself.

Certificates have potential, but we need a better process review system before certificate authorities are automatically trusted in the distributed builds, and an improved user interface for adding certificate authorities.

Certificates were always a reputation service type of system, since they were issued by a “trusted third party.” Today Internet Explorer, Firefox and other browsers come preloaded to “trust” hundreds of certificate authorities. This implies trust of both the key management those vendors use and the certificate authority process itself. Bank of America has investigated the process for determining which authorities the browsers trust, and we found a minimum of due diligence. Essentially a Certificate Authority is either trusted or not. And a user can add even an untrusted authority, with little understanding of the process—or the vulnerability being created by the extended trust inherent in this acceptance.

As noted above, we need a consistent multi-value trust system for certificates. There could, for example, be “high,” “medium” and “low” trust built into the browser and certificates. Verisign has always had a four-level certificate model, which nicely aligns in quantity (but not substance) with the National Institute of Science and Technology’s (NIST) use of four authentication levels

[1] One study showed consumers trusted a site displaying the Verisign logo even when a toolbar noted that the site was malicious and not to be trusted.

in Special Publication 800-63. Whatever the levels, users would enjoy greater confidence and awareness if the browsers consistently represented a level of trust, not just the "on" or "off " ("trusted vs. "non-trusted") represented by the lock icon.

There are other reputation dimensions as well. Is this really Joe Cool? Can Joe repay a loan? Does Joe have the assets to buy a new Rolls Royce? eBay offers a reputation service (also proving that such services are not foolproof) where users rate each others sales and purchasing habits. Dun&Bradstreet (http://www.dnb.com/us/) has long offered information about the financial stability of companies doing business in the US and around the world. Banks have a "know-your-customer" responsibility that lets them manage customer reputation in identity and financial dimensions, and credit bureaus offer a credit reputation service.

Ultimately, we need methods for giving customers several types of reputation service with which they can evaluate web sites, e-mail, and other internet services, including payment options. And we need a reputation service that is equally accessible to our consumer and commercial customers. Ease of use is critical to reputation services. Cryptographic services like digital certificates are available and could be deployed to support trustworthy reputation services. However, certificate trust alone has proven so far to be burdensome and confusing to the home user population.

Summary

These three issues—trustworthiness, the human interface, and reputation services—need to be addressed as part of the W3C efforts to improve Web Authentication for consumers. These issues address whether the software does what it's supposed to, whether customers are really getting what they see, whether businesses representing data and applications on the Internet are reliable, and indeed whether the data and code are reliable.

As these issues are addressed, solutions need to support not only conventional computers, laptops, and servers, but must also work well on the growing number of mobile devices being used to access online services, from Sony PSPs to Blackberrys to Windows Phones and beyond. Trustworthiness, user interface, and reputation considerations are just as critical on mobile devices as they are on the larger systems. Companies are already encountering the issues caused by a lack of security in features like Short Messaging Service (SMS) and instant messaging that they have been working around in e-mail since long before phishing was a problem.

New services in particular must be designed to support a multi-value trust and reputation capability beyond the minimal security features in the first generation of Internet services. Those minimal features cost all of us in consumer and client trust as we attempted to encourage them online.

We can do better, and we at Bank of America believe we will do better. We are proud to support security leadership by offering these comments and ideas to the W3C effort to improve Web Authentication. We look forward to additional participation in the future.

Authors:

Todd Inskeep (Todd.Inskeep@BankofAmerica.com, +1 704-386-6441)

Brian Shea (Brian.A.Shea@bankofamerica.com, +1 206-358-6702

Don Provencher (Donald.P.Provencher@bankofamerica.com, +1 518-626-2609

Sam Phillips (Sam.Phillips@bankofamerica.com, +1 704-388-6375)

Craig Worstell (Craig.Worstell@bankofamerica.com, +1 253-265-8849)