The TIPPI Point: Toward Trustworthy Interfaces

Sara Sinclair and S.W. Smith
Dartmouth College

Although technologists might hate to admit it, computing systems exist not for their own sakes, but rather to serve human processes. Although such processes often involve actual humans, the research community has only recently started considering humans when designing secure systems (for example, see the Secure Systems department, "Humans in the Loop: Human–Computer Interaction and Security" from May/June 2003 [pp. 75–78] or the usability and security special issue from September/October 2004). Unfortunately, this neglect has created an infrastructure that makes humans more likely to divulge their authentication information to spoofed interfaces (as the Attack Trends department, "Interface Illusions," in the November/December 2004 issue discusses).

This problem's urgency led Burt Kaliski of RSA Security and Dan Boneh of Stanford to organize the first Workshop on Trustworthy Interfaces for Passwords and Personal Information (TIPPI; http://crypto.stanford.edu/TIPPI), held 13 June 2005 at Stanford.

As a yellow legal pad went around the room for invited participants to note their contact information, many chuckled when they saw that next to the traditional "Name" and "Email" columns, was a "SS#" column. Identity theft has steadily increased with Internet popularity, but the introduction and rapid proliferation of phishing attacks has shown that elements of current authentication systems—be they the authenticators themselves (usually passwords) or the authentication interfaces—are growing as insufficient as the US's Roosevelt-era social security numbers.

The TIPPI workshop brought security and user interface professionals together to explore this problem and determine ways to improve our authentication methods so that users won't be tricked into giving away personal information. Here, we consider some of the themes discussed at TIPPI, including the nature of the authentication problem, systems that might help solve it, and other observations on necessary components of secure systems designed for human users.

Problems

Todd Inskeep from Bank of America offered insight into the state of PC security among his corporation's user base during his presentation, "Roots of Trusted Interfaces and the User Experience." He noted in particular the challenge of providing online services to users with ancient (by today's standards) machines. He said that most hardware and software platforms from the past 10 to 15 years are still alive and trying to connect to the bank's site—it even sees the occasional connection from machines running DOS! Another former bank employee related a similar anecdote regarding the difficulty of supporting users on WebTV; however, when his organization tried to lock out those machines, they faced a lawsuit from one of the roughly 70 bank clients still using that technology.

Inskeep also cited information regarding the client-side human factor. A Gartner survey (www.pcworld.com/news/article/0,aid,118841,00.asp) found that users want the option of using a more secure authentication method than passwords, but they don't want to be forced to do so. They prefer lower-tech options, such as challenge-response or user-selected image schemes. Less popular were the common smartcard/USB token and secure software download solutions. Inskeep said that providers must be able to maintain trustworthiness in users' eyes—especially in a world in which 2 to 25 percent of phishing emails result in users actually visiting fraudulent Web sites. Bank of America is considering implementing two-factor authentication, one-time passwords, and digital certificate schemes as possible solutions in the next five years (at least for its high-value customers, who have more financial incentive for using strong authentication). The bank is also rolling out a simple "site-key" mechanism that lets users choose a customized picture that will provide assurance that the bank actually runs the Web site they're visiting (pictures are a popular tool for authentication among users, and other presenters at the conference had similar proposals, as we discuss later).

In "Evolution of the Threat and Its Impact on Requirements," Dave Jevans of the Anti-Phishing Working Group provided a sobering look at the phishing community's structure and tools. In particular, he showed that identification information has become the commodity of choice, citing organizations such as the International Association for the Promotion of Cybercrime, and Web sites where visitors can buy hundreds, even thousands, of credit-card numbers or other types of personal information at a time (often paying with more hijacked personal information). The phishing sites this community produces are often short-lived—the average duration was 5.8 days in April 2005 (http://antiphishing.org/APWG_Phishing_Activity_Report_April_2005.pdf)—but plentiful, with 2,854 unique attacks in that same month.

Computers in the US host roughly 30 percent of identified phishing sites, although Jevans stipulates that many of these are compromised personal machines being used as parts of botnets. Sending phishing emails from such networks also makes it much harder for real-time blackhole spam filters to identify and block attacks. Asian countries such as China and Korea are quickly rising through the ranks as hosts of phishing servers as well, and Jevans suggests that a non-US location is desirable because the time, legal, and lingual differences makes it harder for US victims to convince ISPs in other countries to shut down the offending connections.

Additionally, these tools that attackers are using are getting more sophisticated: botnet utilities let the attackers who control such networks of zombies download and execute files, store and upload regular screenshots on compromised machines, and block or redirect Web browsers when users try to access specified sites. Compromised machines that become part of a botnet no longer really belong to their physical owners, although owners are likely to continue to use them for day-to-day tasks.

Countermeasures

Luckily, the TIPPI workshop focused not only on current vulnerabilities, but also on ways to build more resilient systems. In "Trustworthy User Interface Design: Dynamic Security Skins," Rachna Dhamija of the University of California, Berkeley, codified five properties that attackers exploit and that solutions must address:

- Humans have limited skills and abilities (for examining certificates, remembering long, secure passwords, and so on).
- Humans are unmotivated and will choose the path of least resistance, even if it's less secure. (For example, Dartmouth students prefer to use passwords on public machines rather than freely available USB tokens, even though they know the security risks.)
- Given that most program interfaces aren't customized on a per-user basis, a malicious program can easily spoof general-purpose graphics.
- Users often fall victim to the "Golden Arches" property: users trust names and brands (such as that of the internationally ubiquitous restaurant chain), and attackers can easily use a logo or a similar domain name to gain that trust.
- As soon as users give away their secret credentials (such as passwords), this information can never be secret again (the "Barn Door" property).

An old tool from the crypto arsenal can help. In "Solutions for Secure and Trustworthy Authentication," Ramesh Kesanupalli of Phoenix Technologies presented the Simple Password-authenticated Exponential Key Exchange (SPEKE) protocol, which removes the need for users to provide a relying party with a password. Instead, SPEKE constructs a proof of knowledge based on the password that provides mutual authentication. This scheme also lets users maintain their current views of the authentication process.

Behind bars

Many of the other schemes presented focused exclusively on the client software running on users' computers. In trying to combat users' inability to recognize untrustworthy sites, numerous browser toolbars now present information about a server's name, location, and certification status. In "Net Trust," Allan Friedman of the Kennedy School and Alla Genkina of Indiana University proposed a solution that takes into account not only users' strengths but also their online interests. Particularly, they suggested using a social networking toolbar—instead of just providing information about the site being viewed, their tool also lets users know whether friends and associates with whom they're linked have visited and trust the site.

Similarly, Trustbar, a Web browser toolbar by Amir Herzberg of Bar Ilan University, looks toward users' strengths in presenting information about companies through their logos. Instead of asking users to compare URL strings and domain names, Trustbar presents a site's logo as well as the logo of the certification authority that vouches for the site's identity (and has signed the cert and logo). This approach encourages users to use the Golden Arches property to actually improve security.

In addition to her five exploitable properties, Dhamija also presented a secure interface for Web browsers. In her scheme, a remote trusted party, such as a bank, lets users choose a

Humans are unmotivated and will choose the path of least resistance, even if it's less secure.

customized picture, which the bank then uses as the background of the browser space into which users type their usernames and passwords. Phishers would have difficulty accurately guessing users' customized pictures, so users can be fairly certain that they're visiting the sites they think they are. Furthermore, the textboxes in the secure browser space are transparent, such that users can see the pictures' details; if an attacker tries to spoof just the textboxes, users should notice that they can't see pictures through them.

Users must be forced to confirm the security for them to do so at all.

Such browser customization occurs automatically and without user interaction; when an authentication session is initialized, the browser generates a random number and from that, a visual hash. The result is a colorful yet simple image that the browser tiles as a border around secured browser windows or regions (including the login region with the user-specific picture background).

Steve Myers of Indiana University presented a similar image-based approach to mutual authentication in "Delayed Password Disclosure." He proposed giving users a sequence of pictures that they receive in order as they input their passwords in the browser, after either each letter or every few letters. In this manner, the server authenticates itself to the user continuously; it would be very difficult for a phishing site to correctly guess all the pictures to be displayed.

Workshop attendees gained further insight into the relationship between users and their interfaces through the work of MIT's Min Wu, Simson Garfinkel, and Robert Miller: "Users are Not Dependable—How to Make Security Indicators to Better Protect Them." Wu's group tried to evaluate existing antiphishing tools from a user perspective. They divided browser toolbars into three categories:

- *neutral information*, which states simply the name or location of the server to which the user is connected,
- *system decision*, which gives the user an informed opinion on a site's trustworthiness, and
- *positive information*, which informs the user of the site's trustworthy characteristics, such as whether it's been signed and which authority signed it.

The study's participants were divided into three groups, one for each type of toolbar. They visited several Web sites to accomplish certain tasks, but didn't provide personal information if they suspected that a site was fraudulent. Halfway through the experiment, the users received a brief email tutorial on how to identify phishing attacks.

Wu's work found that users in the system-decision toolbar group fared the best against various phishing attacks, with a 35 percent spoof rate before the tutorial, and a 13 percent spoof rate afterward. Users in the neutral information toolbar group were the most susceptible to attack before the tutorial, with a 54 percent spoof rate, but this dropped to 28 percent afterward, whereas the positive-information toolbar was susceptible 39 percent of the time before and 33 percent afterward. All in all, the tutorial clearly improved users' ability to recognize an attack, but no matter which toolbar they had, they were still extremely susceptible: 20 out of 30 users were fooled by at least one attack.

This susceptibility seems due in part to users' reliance on content to provide them with security information—if a site looks trustworthy and is well designed, users will almost go out of their way to justify trusting it: "Yahoo might have just opened a branch in Brazil. That must be why this site is registered there." This reliance on content is another expression of the Golden Arches property.

Trusted paths

Wu also presented work that his group has done using cell phones as proxy devices for authentication to a remote party from a potentially untrusted public terminal.[1] As part of the authentication process, users must verify that the session name displayed on the workstation is the same session name the proxy server gives their cell phone. When the cell phone program asked users simply to confirm that the session names on the two displays matched, users fell prey to various attacks 30 percent of the time. In general, users were inclined to explain the odd behavior the system exhibited during the attack as a momentary glitch that prevented the correct session name from being displayed in both places, and not a compromise.

To improve the authentication system, Wu's group had the cell phone interface ask users to choose the session name from five displayed on the computer monitor (one option was "none of them"), instead of just confirming that the session names match. This minor interface change had dramatic results: in the second round of testing, not a single attack was successful. According to Wu, the moral is that users will not use security indicators well unless their correct usage is part of the critical action sequence—users must be forced to confirm the security for them to do so at all.

Aaron Emigh of Radix Labs and Dan Boneh of Stanford University considered several lower-level solutions to authentication problems in their separate presentations. In "Trusted Path in Heterogeneous Environments," Emigh suggested considering several models of trust interaction, which invest trust in alternating elements of secure systems, including the user, the user's computer (its hardware, OS, and applications), the network, and the remote relying party. When users trust some component of their computers, for example, such as an application or its OS, they can use a "secure attention sequence" to tell the trusted element to protect the data about to be input from the computer's potentially malicious components. The security community has long considered such a signal to be a useful mechanism for trusted communication between user and computer. (Browser and OS designers frequently assume that a "yes" click to a dialog box constitutes a secure attention sequence; in work presented in July 2005 at EuroPKI, Adil Alsaid of Royal Holloway showed how user-level malware can spoof this response—and silently wreak havoc with things such as the browser's trust roots.[2])

In some cases, the remote party to which the user is connecting isn't trustworthy, as with phishing attacks. Boneh's work on hashed passwords, "Trusted Interfaces for Sensitive Data," uses a browser plug-in to take the user's password input, combine it with the name of the site to which the user is connecting, and use this new string as the password during authentication. This way, if the user always tries to authenticate to "hotmail.com," his or her password will always be the same; however, if the user authenticates to "fake-hotmail.com," the hash the phisher receives as a password won't be the same as the one given to the real Hotmail site. To prevent JavaScript attacks in which an attacker steals a password before the plug-in can hash it, Boneh suggests using the secure attention sequence to input the password directly into the computer's trusted element. In some ways, this treats even users as untrusted entities—that is, a relying party doesn't have to trust users to know when it's safe to give out passwords.

Emigh also considered situations in which users don't trust their computers' software or I/O devices. The recent Sumitomo Bank exploit, in which attackers installed hardware keyloggers throughout the bank's London branch—most likely an inside job—shows that this vision of the trust landscape is becoming increasingly prevalent in the real world. To establish a trusted path for secret information in this case, the password and credit-card hashing could occur on a remote device, such as a Palm Pilot. Along these lines, Sara Sinclair presented our PorKI project, which uses PDAs to mitigate the risks of untrustworthy interfaces and machines. Microsoft's Dave Stevens also suggested using a trusted device for authentication in "Securing Online Transactions with a Trusted Digital Identity." In his scheme, a smart-card reader could be enhanced with a display that lets users verify and approve any actions before they're performed with the credentials stored on the smart card.

To prevent attackers from using spyware for keylogging, Boneh also suggested employing VMware to separate the open virtual machine used for normal activities from a closed virtual machine that's used only for secure actions (such as password hashing). Securing the border between the two machines would be easier than securing an OS from a malicious software attack because the line is more clearly drawn and thus easier to defend. Kesanupalli also suggested implementing a firmware-level encryption engine, arguing that spyware might take the path carved by rootkits and burrow so far down into a machine that even full OS reinstallation can't extricate it.

As security practitioners, we still have a lot to learn about creating usable, trustworthy user interfaces for authentication. In the future, we expect that many of the schemes discussed at TIPPI will come to widespread deployment. Those that exploit humans' strengths, such as the ability to recognize pictures quickly, are particularly promising. Portable hardware devices will also likely become more prevalent with time, especially if we can integrate them into already prolific PDAs and mobile phones. Solutions that require more infrastructure development, such as those that involve secure attention sequences or firmware protection of credentials, will likely become necessary as attackers' methods become more sophisticated. Through continued collaboration between researchers and those in industry at venues such as TIPPI, we will hopefully learn how to build such authentication solutions to use our strengths without allowing our weaknesses to provide means of compromise. □

References

1. M. Wu, S. Garfinkel, and R. Miller, "Secure Web Authentication with Mobile Phones," *Proc. MIT Project Oxygen: Student Oxygen Workshop 2003*, MIT Project Oxygen, 2003; http://sow.csail.mit.edu/2003/proceedings.html.
2. A. Alsaid and C. Mitchell, "Installing Fake Root Keys in a PC," *Proc. 2nd European PKI Workshop*, to appear in LNCS, Springer-Verlag, 2005.

S.W. Smith *is an assistant professor of computer science at Dartmouth College. Previously, he was a research staff member at IBM Watson, working on secure coprocessor design and validation, and a staff member at Los Alamos National Laboratory, doing security reviews and designs for public-sector clients. Smith received a BA in mathematics from Princeton University and an MSc and a PhD in computer science from Carnegie Mellon University. Contact him at sws@cs.dartmouth.edu; www.cs.dartmouth.edu/~sws/.*

Sara "Scout" Sinclair *is a PhD student in Dartmouth College's PKI/Trust Lab. Her current research interests include security-enabling portable devices and human factors in security design. She received a BA in computer science and French from Wellesley College. She also serves as president of the Dartmouth College Graduate Student Council. Contact her at sinclair@cs.dartmouth.edu.*