Improving the Human Element: Usable Privacy and Security

Jason Hong, Lorrie Cranor, Norman Sadeh
jasonh@wombatsecurity.com
Wombat Security Technologies
1000 Heberton Street
Pittsburgh, PA 15206

1. Introduction

Information and communication technologies are pervasive in all aspects of every day life, including transportation, manufacturing, utilities, finance, and entertainment. It would not be an understatement to say that we depend on these systems being highly reliable for modern society to function. However, we have been witnessing an increasing number of privacy and security failures in these systems. What is interesting is that many of these failures happen not because of breakdowns in algorithms, software, or hardware, but because of failures in the user interface.

As we become more reliant on our computing infrastructures, the consequences of breaches in privacy and security are becoming more severe. These breaches might occur due to misconfigurations of firewalls or file servers[1], difficult to use software for encryption or ecommerce, lost laptops that contain sensitive corporate information, or social engineering attacks by malicious criminals intended to steal sensitive information.

There is growing recognition that privacy and security failures are often the results of cognitive and behavioral biases and human errors. Many of these failures can be attributed to poorly designed user interfaces or secure systems that have not been built around the needs and skills of their human operators: in other words, systems which have not made privacy and security *usable.*

To underscore this point, in 2003, the Computing Research Association (CRA) issued a grand challenge for computer security and privacy: "Give end-users security controls they can understand and privacy they can control for the dynamic, pervasive computing environments of the future." Similarly, the National Academy of Engineering (NAE) included "secure cyberspace" in their 2008 Grand Challenges for Engineering, arguing that more research is needed on the psychology of computer users, how people interact with their computers, and how "cultural and social influences can affect how people use computers and electronic information in ways that increase the risk of cybersecurity breaches."

The design and analysis of secure systems must consider humans as an integral part of the system under consideration, rather than a secondary constraint. Humans, however, have strengths and weaknesses considerably different from those of the rest of the system. Mismatches between what users can actually be expected to do and what the rest of the system assumes they will do is one of the main causes of security failures.

As such, we argue that usable privacy and security is critical to our nation and to society in general, so that we can sustain and continue to enjoy the benefits of information and communication technologies we have seen over the past decades. Success in this endeavor could lead to more reliable and dependable systems, including those designed for personal use, for corporate use, as well as those that help run our national infrastructure.

In this white paper, we outline some barriers to effective usable privacy and security, give a case study with our work in protecting people from phishing scams, describe the potential for

[1] See, for example, "GOP clerks nabbed Democratic Data, says probe", for an example of how a misconfigured file server led to everyone being able to see files, http://news.cnet.com/GOP-clerks-nabbed-Democratic-data%2C-says-probe/2100-1029_3-5170987.html

breakthrough research and development in this area, and close with how usable privacy and security matches the goals and functions of NIST.

2. Barriers to Effective Usable Privacy and Security

In this section, we provide an overview of some of the challenges facing usable privacy and security.

Usable privacy and security spans multiple disciplines

Usable privacy and security currently spans a number of disciplines, including such diverse fields as human-computer interaction, computer security, distributed systems, mobile computing, networking, machine learning, cognitive psychology, social psychology, decision sciences, learning sciences, and economics. This breadth of areas makes a challenging backdrop for this emerging discipline, as each of these fields has different values, tools, and methods. One consequence is that there is currently not yet widespread agreement on the best methods and techniques for developing and evaluating systems. Furthermore, no person is trained in all of these areas, making it hard for an expert in, say, computer security, to assess the value of a research paper (or product) having a novel user interface design for helping manage security.

Gaps in the Body of Knowledge

There are also many gaps in terms of basic facts, statistics, and models. Many of these gaps lie between the disciplines listed above. For example, what is the best way of designing warnings and alerts so that people can see them, understand them, and be motivated to act on them? What are better ways of motivating individuals in organizations to practice better security? How do people perceive and understand risks online, and what are better ways of shaping people’s perceptions? What is the best way of allocating the division of labor for security tasks, between automation with computers and monitoring and actions with people?

Privacy and Security are Secondary Tasks

Applying traditional human-computer interaction methods to privacy and security studies is challenging, because achieving security and privacy are typically *secondary tasks* for end users. Usable privacy and security requires the design of studies centered around realistic primary tasks that lead users to interact with security or privacy features. Because users often feel safe in a laboratory environment and may not believe their security or privacy is at risk, designing studies requires special care, and user studies must often be conducted outside the laboratory.

Related to this issue is the notion of threat models. That is, in some cases, the system needs to be evaluated to see if end-users can correctly configure things. In other cases, the system needs to be evaluated with active adversaries.

Growing Number and Complexity of Services and Devices

Finally, usable privacy and security is hard to achieve in practice due to the rapidly growing number of services and devices available to people. The challenge here is in managing the complexity of multiple operating systems, network services, data distributed across these different services and devices, and the interactions between these services and devices.

3. Case Study – Protecting People from Online Phishing Scams

In this section, we give an example area of research in usable privacy and security, taken from the authors' past work in protecting people from phishing scams. This case study illustrates some of the challenges, opportunities, and science behind usable privacy and security.

Semantic attacks are a kind of attack on computer systems that targets the users of a system rather than the hardware or software. The most common semantic attack today is phishing, where criminals impersonate legitimate people or organizations and trick people into giving up sensitive information (such as passwords, credit card numbers, or corporate secrets) or installing dangerous malware on their computers, such as viruses and worms. The most common form of phishing are those fake "please update your account" emails that direct people to fake sites that appear like legitimate sites. However, phishing also comes in the form of fraudulent instant messages and even Voice over Internet Protocol (VoIP) calls.

Phishing is a form of *social engineering* that takes advantage of people's general lack of understanding of how email, the Web and other technologies really work: what is legitimate and what can easily be spoofed. The past several years have seen a steady rise in attacks targeting banks, e-commerce sites, universities, and government organizations, especially the Department of Defense.

The Anti-Phishing Working Group, an international consortium of organizations committed to wiping out Internet scams and fraud, keeps track of phishing activity, including the number of unique phishing Web sites detected every month. In 2007 monthly totals ranged as high as 55,643. During each month in 2007, anywhere from 92 to 178 different company brands were "phished"—meaning their names or logos were used to fool victims into thinking they were dealing with a trusted institution. According to research and consulting firm Gartner, an estimated 3.6 million Americans fell victim to phishing last year, leading to losses of more than $3.2 billion. Note that this does not include indirect damage to an organization's reputation or loss of potential sales.

For corporations and government organizations, though, this is just the tip of the iceberg, as more targeted "spear phishing" attacks can lead to potentially devastating security breaches (e.g. loss of sensitive company or national security information), as reported in the cover story of Business Week's April 10, 2008 edition. Concurrently, regulations such as HIPAA, Sarbanes-Oxley, or Gramm-Leach-Bliley, are requiring corporations to secure the use of all electronic forms of communications, including Web-based communications. It is no surprise therefore that phishing has grown to be viewed as a particularly high priority threat by many organizations.

While in principle technologies such as Public Key Infrastructures (PKI), digital signatures, and multi-factor authentication could help minimize the impact of phishing, practical deployments of these technologies are prone to numerous vulnerabilities which industry has not been able to address.

With so much at stake, the computer security community has been scrambling to develop technologies to combat phishing. Researchers and commercial vendors have developed filters for e-mail and Web browsers that flag phishing attempts. Although such software has helped stop many attacks, phishers are constantly evolving their tactics to try to stay a step ahead of such technologies. Since phishing plays on human vulnerabilities—a successful attack requires a victim to succumb to the lure and take some action—it is also not strictly a technological problem.

Our research group at Carnegie Mellon University investigated several lines of research to understand the best way to protect people. These include studies to understand why people fell for phishing attacks[2],[3], as well as studies to understand why certain browser warnings were ineffective in protecting people[4]. We also developed several technologies to protect people, including automated filters that made use of machine learning and information retrieval techniques to detect fake phishing emails and web sites[5],[6], an embedded training system named PhishGuru that sends fake phishing email and trains people that falls for our messages[7], and a game called Anti-Phishing Phil that teaches people how to identify fake URLs and where to look in their browser for URLs[8].

Thus far, our work has generated a great deal of interest and collaboration from a number of partners. Our automated email filter is undergoing a field trial at Carnegie Mellon University's main email servers, where it will filter several million emails per day. Our research evaluating anti-phishing toolbars has been cited by several companies, with ongoing evaluations being presented to the Anti-Phishing Working Group, a consortium of companies "committed to wiping out Internet scams and fraud." Design suggestions from our studies to understand browser warnings have been incorporated into the latest version of Microsoft's Internet Explorer 8. PhishGuru's methodology of sending fake phishing emails to train individuals has undergone field trials at three different companies, and been cited by two different companies trying to commercialize the work. PhishGuru's training materials have also been adopted by APWG on their landing page, a page that ISPs and web sites can show after taking down a phishing web site. Anti-Phishing Phil has been played by over 80,000 people, licensed by two companies, demoed at many security days meant to teach people about good security practices, and translated into Portuguese with several more translations underway. Given all of this interest, we decided to commercialize all of this work through a startup we have founded, named Wombat Security Technologies.

[2] See J. Downs, M. Holbrook, and L. Cranor. Behavioral Response to Phishing Risk. Proceedings of the 2nd Annual eCrime Researchers Summit, October 4-5, 2007, Pittsburgh, PA, p. 37-44. http://www.ecrimeresearch.org/2007/proceedings/p37_downs.pdf

[3] J. Downs, M. Holbrook, and L. Cranor. Decision Strategies and Susceptibility to Phishing. In *Proceedings of the 2006 Symposium On Usable Privacy and Security*, 12-14 July 2006, Pittsburgh, PA. http://cups.cs.cmu.edu/soups/2006/proceedings/p79_downs.pdf

[4] S. Egelman, L. Cranor, and J. Hong. You've Been Warned: An Empirical Study of the Effectiveness of Web Browser Phishing Warnings. CHI 2008. http://doi.acm.org/10.1145/1357054.1357219

[5] I. Fette, N. Sadeh, and A. Tomasic. Learning to Detect Phishing Emails In *Proceedings of the 16th International conference on World Wide Web*, Banff, Alberta, Canada, May 8-12, 2007. http://doi.acm.org/10.1145/1242572.1242660

[6] Y. Zhang, J. Hong, and L. Cranor. CANTINA: A content-based approach to detecting phishing web sites. In *Proceedings of the 16th International conference on World Wide Web*, Banff, Alberta, Canada, May 8-12, 2007. http://doi.acm.org/10.1145/1242572.1242659

[7] P. Kumaraguru, Y. Rhee, A. Acquisti, L. Cranor, J. Hong, and E. Nunge. Protecting People from Phishing: The Design and Evaluation of an Embedded Training Email System. In CHI 2007: Conference on Human Factors in Computing Systems, San Jose, California, 28 April - May 3, 2007, 905-914. http://doi.acm.org/10.1145/1240624.1240760

[8] S. Sheng, B. Magnien, P. Kumaraguru, A. Acquisti, L. Cranor, J. Hong, and E. Nunge. Anti-Phishing Phil: The Design and Evaluation of a Game That Teaches People Not to Fall for Phish. In *Proceedings of the 2007 Symposium On Usable Privacy and Security*, Pittsburgh, PA, July 18-20, 2007. http://cups.cs.cmu.edu/soups/2007/proceedings/p88_sheng.pdf

4. Potential for High Risk Transformational Research

In this section, we outline how ideas in this white paper could lead to proposal submissions that represent transformational research that could greatly improve the state of computer security and usable privacy and security.

One area for proposals would be better security for **computer networks in the home**. The number of networked devices and services targeted for the mass market is increasing, including personal computers, portable devices (such as iPods, digital video recorders, digital cameras, and mobile phones), home entertainment systems (such as the Tivo, stereos, DVD players, televisions, and game consoles like the Nintendo Wii and Xbox), smart toys, web cameras, electronic photo frames, RFIDs, sensors (for detecting motion, carbon monoxide, etc), and a number of wireless networking systems (such as WiFi, ZigBee, and LonWorks). The increasing number of these devices and services is making it so that every home requires their own system administrator, a burdensome proposition. The complexity of making all of these systems work reliably is increasing potential vulnerabilities as well, as home system administrators need to make sure that all of these are configured properly, have been patched with the latest fixes, and do not have malware installed on them.[9] We have already seen fairly large failures in security on home personal computers, with numerous open web cameras that use default passwords, open WiFi access points because people cannot figure out how to configure them properly, and increasing numbers of botnets based on home computers. These problems will only get worse unless academia and industry can figure out better ways of protecting people. Proposals in this space could look at better logging and visualizations for helping people understand and diagnose problems, and better monitoring and visualization tools to help people understand what is being shared and with whom.

Another area for proposals would be **better usable security for mobile devices**. Mobile devices are becoming an intimate part of our lives, storing highly personal information such as photos, contact lists, instant messages, emails, and personal notes. Mobile devices are also being used by corporations to punch through their corporate firewalls. Future mobile devices will also be able to sense one's current location and activity, capturing even more sensitive information about our daily activities. It would not be surprising for future mobile devices to be able to do mobile commerce or even control aspects of one's home (for example, opening doors, turning ovens on and off, and controlling video recorders). As such, the damages that could result from losing a mobile device, or having a mobile device be compromised by malware, can be quite high. Here, there needs to be more support in terms of better authentication schemes for small devices, more usable and automated encryption, and better tracking and recovery for lost devices.

A third area for proposals would be **better tools and user interfaces for configuring policies**. Every system needs some form of configuration, including firewalls, web servers, Facebook profiles, digital video recorders, and Wifi access points. Failures in configuration can lead to accidental disclosures of sensitive information. One of the primary challenges here is complexity, in that even a small number of options can lead to an exponential number of states, all of which cannot be realistically tested by end-users. Here, there needs to be more support in terms of default policies, visualization, machine learning, and automated and semi-automated testing.

A fourth area for proposals would be **protecting people from online phishing scams and other social engineering attacks**. Confidence scams are nothing new, but the scale of the Internet and lack of cues as to an entity's identity is making it easier than ever to create sophisticated scams. As noted in an earlier section, phishing scams have targeted banks, ecommerce sites, social networking sites, corporations, and government employees. Here, a combination of techniques is

[9] See, for example, the Insignia photo frames that came installed with several computer viruses on them. http://www.sfgate.com/cgi-bin/article.cgi?f=/c/a/2008/02/15/BU47V0VOH.DTL

needed, including better training, better user interfaces for displaying salient information about entities, better systems for judging the legitimacy of a message and/or entity, and better filters for eliminating fake messages before an end-user ever sees them. Furthermore, these techniques need to be applied to a number of domains, including email and web, as well as VOIP, SMS, and instant messaging.

A fifth area for proposals is **usable authentication and biometrics**. The most common means of authenticating a user's identity is through a text password, but text passwords can be easy to guess or steal, and because there are a growing number of systems that each demand a unique and difficult-to-guess password. Research is needed to find ways to make user authentication more convenient without sacrificing security, as well as usable ways to allow users and systems to mutually authenticate. Several technical, sociological, legal, and practical challenges remain unsolved and have hampered the widespread deployment of biometric technology. On the technical side, the recognition ability of most biometric systems still needs a lot of improvement to be dependable. Moreover, better security protocols are needed for biometric template storage and transmission, and for securing biometric databases. On the non-technical side, more consideration of privacy issues as well as the sociological impact of biometric technology is needed. Privacy issues pose major obstacles to more widespread adoption of biometrics.

It is also important to emphasize that research and development in the above areas would also greatly contribute to our basic understanding of usable privacy and security, leading to generalizable results that could be applied to a number of applications. That is, the above research would lead to **good applications as well as good science**.

For example, research along the lines of this proposal could lead to a deeper understanding of human-in-the-loop systems, helping us to understand when certain security functions should be automated (for example, due to likelihood of error or lack of understanding by end-users), when functions should have better user interfaces, and when functions should require more training on the part of end-users. In many cases, we expect that all three approaches of automation, better interfaces, and training will be used, though the question for specific domains is what the proper division is.

Note that not all research for usable privacy and security need to directly include a human element. For example, work on machine learning to create better automated filters for protecting people would also fit under the primary goals as described by this white paper.

Research in usable privacy and security could also lead to better models to help guide practical application development. These might include, for example, models of how social engineering works and proven techniques for training and protecting people, better models for how alerts and warnings are seen and understood by people[10], better ways of running realistic usability studies to assess the effectiveness of tools

5. Usable Privacy and Security with Respect to NIST

The goals of this proposal mesh well with NIST's Information Technology Security and Networking (ITSN) division. From NIST's web site, the goal of ITSN includes "establishing, implementing, and testing information security policies, procedures, and technologies for NIST's administrative and scientific environments." This proposal also matches the goals of NIST's Information Technology Laboratory's Computer Security Division as well, which focuses on

[10] For example, in some of our past work, we studied why the anti-phishing warnings in Internet Explorer were not as effective as the warnings in Mozilla Firefox. See S. Egelman, L. Cranor, and J. Hong. You've Been Warned: An Empirical Study of the Effectiveness of Web Browser Phishing Warnings. CHI 2008.

“cryptographic standards and applications, security of emerging technologies, security management, and security testing.”

NIST also has a history of supporting usability. For example, NIST was the main developer of the Industry Usability Report (IUSR) for reporting the results of usability tests.[11] NIST also was instrumental in developing several web metrics and tools for assessing these metrics.[12] Finally, NIST has also held several workshops on improving usability, for the web and for computer applications in general.

Given this background of computer security and usability, we argue that this proposal fits well with NIST’s existing portfolio of areas. Furthermore, effective computer security is not a set of silver bullet technologies, but rather a process that must involve people as well as hardware and software. Thus, we view this proposal for usable privacy and security as a natural extension of NIST’s interests in computer security.

[11] http://zing.ncsl.nist.gov/iusr/

[12] http://zing.ncsl.nist.gov/WebTools/