Security User Studies: Methodologies and Best Practices

Serge Egelman
Carnegie Mellon University
egelman@cs.cmu.edu

Jen King
Yahoo!, Inc
jenking@alumni.
sims.berkeley.edu

Robert C. Miller
MIT CS & AI Laboratory
rcm@mit.edu

Nick Ragouzis
Enosis Group LLC
nickr@enosis.com

Erika Shehan
Georgia Institute of Technology
erika@cc.gatech.edu

CHI 2007, April 28–May 3, 2007, San Jose, California, USA.
ACM 978-1-59593-642-4/07/0004.

Abstract

Interest in usable security -- the research, development, and study of systems that are both usable and secure -- has been growing both in the CHI and information security communities in the past several years. Despite this interest, however, the process of designing and conducting security-related user studies remains extremely difficult. Users deal with security infrequently and irregularly, and most do not notice or care about security until it is missing or broken. Security is rarely a primary goal or task of users, making many traditional HCI evaluation techniques difficult or even impossible to use. This workshop will bring together researchers and practitioners from the HCI and information security communities to explore methodological challenges and best practices for conducting security-related user studies.

Keywords

Information security, usable security, user studies

ACM Classification Keywords

H5.m. Information interfaces and presentation (e.g., HCI): Miscellaneous.; H.5.2 User Interfaces: Evaluation/methodology

Introduction

As networked computing weaves itself into many aspects of daily life, ensuring the security of networked systems is becoming vitally important. Much of the existing body of information security research focuses on making cryptographic algorithms that are harder to break, protocols that are more robust, and computer systems that are resistant to attack. Although this work provides a foundation necessary for creating secure environments, lack of attention to usability in security has resulted in a world where security-related functionality is often complex and counterintuitive. Interest in usable security -- the research, development, and study of systems that are both usable and secure -- has been growing both in the CHI and information security communities in the past several years, as indicated by a number of research articles on usable security, a CHI workshop focused on HCI and security systems [9], and the establishment of a conference on exclusively focused on usable security and privacy [1]. Despite this increased interest in usable security, the process of conducting effective, ethical security-related user studies remains daunting, even to experienced HCI practitioners and researchers.

Challenges of Security-Related User Studies

Conducting security-related user studies can be extremely difficult. Users deal with security infrequently and irregularly, and most do not notice or care about security until it is missing or broken. Security is rarely a primary goal or task of users, making many traditional HCI evaluation techniques difficult or even impossible to use.

Security-related user studies employing observational methods are extremely difficult to design. Gaw et al. [5] and Dourish et al. [2] argue that user practices related to security need to be understood, however obtaining observational data about user practice is challenging, as users only deal with security on rare occasions. How could researchers modify existing observational techniques used in HCI to collect sporadically-occurring data about security practices?

Qualitative techniques such as interviews and surveys have been successfully used in a number of studies [4, 3, 6, 13], but have significant limitations. For example, it is well known that participants may claim to take particular actions with respect to their security (or privacy), when in fact they actually take completely different ones in reality [7]. How should researchers and practitioners deal with this inconsistency between what people say and what they do with respect to security? Given the infrequency and irregularity of security-related user activity, as well as unreliability of self-reported data, how can researchers design studies to better understand user practices related to security?

Designing security-related lab experiments can also be difficult. Whalen and Inkpen note that in their study of web browser security, people didn’t act to protect data treat as if it was their own [10]. Wu et al. [12] and Whitten [11] also acknowledge that user motivation in security lab experiments is a significant problem. Moreover, there is a chance that when an evaluator tries to motivate a study participant to complete security related tasks, he or she may be introducing bias into the study by priming the participant to focus more on security than he or she would outside of the experimental setting [12]. How can evaluators design laboratory experiments that are faithful to the fact that in the real world, security is almost never a primary goal? How can evaluators motivate study participants to explicitly act on security related tasks without overemphasizing security?

To overcome the limitations of lab experiments, some usable security researchers have attempted experiments that have involved launching attacks or deceiving subjects in real-world situations. These studies, however, can raise significant ethical concerns. In fact, the usable security community has already faced its first major ethical dilemma associated with this sort of user study. In 2005, researchers at Indiana University conducted a study examining factors influencing response rates to phishing attacks, which had been approved by the IU institutional review board. The study generated results extremely interesting to the usable security community, as over 70% of people involved in the experiment fell victim to the phishing technique used. This study, however, also left many members of the IU community embarrassed, violated, and outraged; a few students even threatened legal action against the researchers [8]. Given the potential for causing emotional harm to participants, when is it appropriate for usable security researchers to conduct attack studies?

Finally, to conflate the difficulties of conducting security-related user studies, there are no comprehensive discussions of best practices available for researchers or practitioners, nor are there any resources such as processes, checklists, or criteria that could assist people without deep knowledge of both HCI and security in conducting user studies [13].

Objective

This one-day workshop will bring together researchers and practitioners from the HCI and information security communities to discuss methodological challenges and best practices for security-related user studies. This workshop will focus specifically on the following issues:

- *Study Design:* How can evaluators design studies that are faithful to the fact that in the real world, security is almost never a primary goal? How can evaluators motivate study participants to complete security-related tasks without overemphasizing security? How should evaluators even decide what to test in a security user study? How can researchers handle the problem that users may claim to take particular steps to protect their security, but in reality do something else?
- *Ethical Issues:* How can evaluators conduct realistic studies involving attacks on users, yet at the same time protect study participants from harm or embarrassment? When is it appropriate to launch security attacks or employ deception in studies?
- *Lessons Learned & Best Practices:* Why have previous security user studies succeeded or failed? What are best practices for security user studies? What would security user study processes, checklists, and criteria look like?

Workshop Structure

The workshop will be divided into two sessions. In the first session, six authors of accepted papers will present their work in 10-15 minute talks, followed by approximately 15-20 minutes of discussion.

The second session will be devoted primarily to small group breakout discussions of methodological issues associated with security user studies. Participants will be split into groups of approximately four people. If possible, each group should have members from industry, academia, HCI, and security. Each group will receive a proposal for a possible user study (these proposals will consist of participant-submitted study proposals as well as proposals created by the workshop organizers in order to stimulate discussion). For one hour, each group will identify and discuss methodological challenges, alternatives to the proposed study design, and ethical concerns. After the hour has passed, each group will present their findings to the entire workshop, followed by a short discussion session. After each group has presented their findings, the workshop will conclude with a brainstorming session that will culminate in the creation of a list of best practices and pitfalls for conducting security user studies.

A poster discussing the outcomes of the workshop will be displayed at the conference poster session. If appropriate, the results of this workshop will result in the publication of one or more papers discussing techniques, best practices, and pitfalls of security user studies in a special issue of a journal or in a trade publication. We also hope that this workshop facilitates future collaborative endeavors among participants.

References

[1] Cranor, L.F. Proceedings of the Symposium on Usable Privacy and Security (SOUPS 2005). Pittsburgh, PA, USA. July 6-8, 2005.

[2] Dourish, P., Grinter, R.E., Delgado de La Flor, J., and Joseph, M. Security in the Wild: User Strategies for Managing Security as an Everyday, Practical Problem. Personal and Ubiquitous Computing, 8, (2004), 391-401.

[3] Downs, J.S., Holbrook, M.B., and Cranor, L.F. Decision Strategies and Susceptability to Phishing. Proc. SOUPS 2006, ACM Press (2006), 79-90.

[4] Friedman, B., Hurley, D., Howe, D., Felten, E., and Nissenbaum, E. Users' Conceptions of Web Security: A Comparative Study. Ext. Abstracts CHI 2002, ACM Press (2002), 746-747.

[5] Gaw, S. and Felten, E.W. Password Management Strategies for Online Accounts. Proc. SOUPS 2006, ACM Press (2006), 44-55.

[6] Gaw, S., Felten, E.W., and Fernandez-Kelly, P. Secrecy, Flagging, and Paranoia: Adoption Criteria in Encrypted Email. Proc. CHI 2006, ACM Press (2006), 591-600.

[7] Gideon, J., Egelman, S., Cranor, L., and Acquisti, A. Power Strips, Prophylactics, and Privacy, Oh My!. Proc. SOUPS 2006, ACM Press (2006), 133-144.

[8] Jagatic, T., Johnson, N., Jakobsson, M., and Menczer, F. Social Phishing. http://informatics.indiana.edu/fil/Net/social_phishing.pdf

[9] Patrick, A.S., Long, A.C., and Flinn, S. HCI and Security Systems. Proc. CHI 2003, ACM Press (2003), 1056-1057.

[10] Whalen, T. and Inkpen, K.M. Gathering Evidence: Use of Visual Security Cues in Web Browsers. Proc. Graphics Interface 2005, ACM Press (2005), 137-144.

[11] Whitten, A. and Tygar, J.D. Why Johnny Can't Encrypt: A Usability Evaluation of PGP 5.0. Proc. USENIX 1999, USENIX Press (1999), 169-184.

[12] Wu, M., Miller, R.C., and Garfinkel, S.L. Do Security Toolbars Actually Prevent Phishing Attacks? Proc. CHI 2006, ACM Press (2006), 601-610.

[13] Wu, M., Miller, R.C., and Little, G. Web Wallet: Preventing Phishing Attacks by Revealing User Intentions. Proc. SOUPS 2006, ACM Press (2006), 102-113.