



Technical Perspective

The Effectiveness of Security Measures

By Nicolas Christin

IN THE LATE 1990S, we came to the realization that users were central to computer and information security. Ross Anderson famously argued that “the threat model was completely wrong” when referring to our historical focus on securing technical components while ignoring possible human mistakes. A large and growing body of research has subsequently attempted to study how people face computer security challenges. Studies in the adjacent field of information privacy revealed that user behavior is complex. People may profess caring about their privacy, but frequently end up making decisions that prove costly, for example, due to limited information or to behavioral biases that lead them to miscalculate long-term risks.

Measuring security behavior turns out to be even more difficult than measuring privacy preferences and actions but imagine for a second that we had the ability to do so. For instance, we could examine the practical relevance of the following well-known, but rarely evaluated, security advice: updating software frequently, browsing reputable websites, using encryption whenever possible, and trying to avoid operating systems that are too common and targeted by villains.

How should we go about it? Carefully controlled experiments, such as asking users to come to a lab and run through a set of predetermined activities, are unlikely to display the potentially risky behaviors in which people engage in the comfort of their home. On the other hand, field measurements, in which users are left to their own devices and are passively observed, without direct interaction, are cumbersome to set up for a range of reasons. Meaningful data collection must be done at scale to exhibit statistical patterns; such studies are highly sensitive, and researchers must take the greatest care in preserving

their participants’ privacy; and, perhaps most importantly, getting access to such user data is nearly impossible without access to a large infrastructure provider.

In the following paper, DeKoven et al. manage to overcome these challenges. The authors leverage their campus network to get access to 15,000 computers in university dorms. They build a remarkable infrastructure to acquire this network traffic, anonymize it to safeguard user privacy, and whittle it down to levels suitable for analysis. The scale of the experiment is particularly daunting. The authors collect approximately 4Gbps to 6Gbps of traffic for six months, which corresponds to roughly 100 petabytes of data in aggregate. A clever idea in the paper is to partially repurpose intrusion detection systems such as Bro/Zeek and Suricata to extract relevant information from these large volumes. In particular, the authors ingeniously create a set of fingerprints to automatically detect operating systems and installed software, and to track individual user behavior across many different, disjoint sessions, with potentially different IP addresses.

They analyze this data to provide a unique perspective into how users choose to implement common security advice in practice. The authors notably discover that there is no difference in system update frequency between users, regardless of whether their computer ends up being infected by malware. However, users tend to more frequently update their Web browsers and Flash software after having been compromised. The main and perhaps most interesting outcome of this part of the study is the absence of very strong, obvious correlations between the likelihood of compromise and lack of adherence to generally accepted security practices.

However, there is one important ex-

ception: Potentially risky Web browsing behavior tends to lead to riskier security outcomes. This result is particularly meaningful because it echoes what other studies run with different users, under different circumstances, had also observed. Going beyond this paper, Web browsing behavior appears to be an important determinant of the potential risk of security compromise, which in turn justifies the large amount of effort we should continue to put in securing the Web.

But, perhaps, the main contribution of this paper lies more in asking questions than providing definitive answers. This research clearly shows that while passive measurements are powerful in helping us establish correlations, they are much more limited when it comes to exhibiting causal relationships. For instance, this paper shows that users of the anonymous Tor network are more likely to get in trouble than others—but is it because users erroneously believe that Tor provides increased protection against many security compromises, or because the malware itself installs Tor to communicate anonymously?

Moving forward, to improve security practices and distinguish between folk remedies and advice rooted in empirical evidence, we should focus on understanding causal relationships between user actions, exposure to vulnerabilities, and security compromises. Combining large-scale passive measurements as described by the authors with finer-grained timeline reconstructions and user interviews could help reach these objectives. As this paper clearly demonstrates, it is a highly ambitious, technically challenging, but overall worthy goal. 

Nicolas Christin is a professor in the School of Computer Science and in the Department of Engineering and Public Policy at Carnegie Mellon University, Pittsburgh, PA, USA.

Copyright held by author.