Computer Networks 56 (2012) 3891–3895

Contents lists available at SciVerse ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Speculations on the science of web user security

Ravi Sandhu *

Institute for Cyber Security, University of Texas at San Antonio, One UTSA Circle, San Antonio, TX 78249, USA

ARTICLE INFO

Article history:
Available online 26 October 2012

Keywords:
Web user security
Security science
Cyber security

ABSTRACT

There appears to be consensus among seasoned cyber security researchers that there is substantial disconnect between the research community's priorities and the real world—notwithstanding numerous intellectual advances in the theory and practice of cyber security over the past four decades. This is in part manifested by recent recurring calls for dramatic shifts in cyber security research paradigms, including so called game-changing approaches that go beyond the typical computer science and engineering perspectives. This article focusses on a specially important piece of cyber security called web user security where the prime concern is security for the ordinary consumer of web application services. The proliferation of web services and their enthusiastic reception by the ordinary citizen attests to the tremendous practical success of these technologies. As such it is prima facie evident that the current web is "secure enough" for mass adoption. Now, one certain prediction about the web is that it will continue to evolve rapidly. This article gives the author's personal perspective on what web user security science might be developed to address the need to be "secure enough" in light of continued evolution. To this end the article begins by considering what happened in evolution of the web in the past and how much of it, if any, was guided by "science." The article identifies some security principles that can be abstracted from this short but eventful history. The article then speculates on what directions the science of web user security should take.

1. Introduction

The current state of web user security is rather paradoxical. On one hand the general public continues to enthusiastically embrace web-based applications with increased productivity and efficiency for all. At the same time stories of cyber breaches and risks of being online proliferate in the media. The basic premise of this article is that this is the steady state situation which will persist indefinitely into the future, and that this is a good situation. The web will never get to the point where crime and mischief is completely eliminated. That could happen only if criminals and miscreants disappeared from humanity, which is not going to occur. Entrepreneurs and innovators will continue to find new ways of utilizing and evolving the web. Some of these will become winners and gain mass adoption. Others will be relegated to niche uses or fall by the wayside, perhaps to come back another day or disappear forever. The conflict between evolving the web through innovation while "securing" it for the public, will ensure a steady state of relative safety for the vast majority of users with a small fraction of unfortunate victims who suffer losses ranging from the merely annoying to substantial.

Given this premise, what should a science of web user security seek to do? This article explores this question from the author's personal perspective. The discussion is impressionistic and speculative rather than definitive. Implicit in the question is that the current science, such as it is, is inadequate. Putting it more strongly, the current science has not been terribly helpful so far. In the past four decades there have been numerous intellectual advances in the theory and practice of cyber security. Nonetheless, many researchers and practitioners agree that the priorities of

* Tel.: +1 210 458 6081.
E-mail address: ravi.sandhu@utsa.edu

1389-1286/$ - see front matter © 2012 Published by Elsevier B.V.
http://dx.doi.org/10.1016/j.comnet.2012.10.010

the real world are substantially disconnected from those of the research community. Research funding agencies in the USA have expressed this concern in recent reports calling for "game-changing" shifts in cyber security research paradigms [1,2]. These influential reports will certainly have impact on future research efforts.

This article seeks to complement these insights by considering this fundamental question from a novel perspective. Our core assumption is that the state of web user security is about as good as it can ever be. It is never going to be much better than where we are today. At first thought this may sound like a pessimistic, and almost defeatist, assessment with lack of faith in technical advancement. But it is actually an optimistic yet pragmatic statement. Consider that people have overwhelmingly voted to go online and become increasingly dependent on the web as a routine part of their daily lives. If web user security was really that bad how could this have happened? In light of mass adoption it is untenable to argue that the current web is not secure enough for the services it provides to the typical user. The success of the web speaks for itself. Prima facie we must accept that the web is secure enough for mass adoption. This observation is further reinforced by the fact there has been no mass departure from the web. Once people are on it they seem to stay on it and even increase their activity. Perhaps we should be teaching web user security as a true success story in our cyber security classes, which to our knowledge is rarely if ever done.

To summarize our position so far: web user security is about as good as we are going to get it with relative safety for the vast majority of users and a tiny fraction of unfortunate victims with losses ranging from mere nuisance to substantial. Given this premise, the fundamental question is what should a science of web user security seek to do.

The rest of the article is organized as follows. Section 2 defines web user security for our purpose. Section 3 considers the short but eventful history of the web from the perspective of its security, so that we can learn from the past. Section 4 proposes a set of security principles that are grounded in the current web and that should guide web user security in the future. Section 5 addresses the fundamental question identified above. Section 6 concludes the article.

2. Web user security

For purpose of this article, web user security is primarily concerned with security of the ordinary consumer of web application services. Web provider security is also required, otherwise compromise of the application provider can cause compromise of the web application consumer. Operating system security is required to protect the user's client platform from where the web is being accessed. Network security is required to protect the user from network-based attacks. Likewise for browser security. Is web user security a meaningful term or does it expand to cover a much larger chunk of cyber security so as to become amorphous?

Considered from a technical perspective web user security does indeed reach out to permeate a large portion of cyber space. Attacks can come from a variety of layers in the software–hardware stack. Countermeasures must correspondingly deploy across these layers. In parallel to the technical complexities of web user security, there are accountability complexities. Who is responsible for securing the operating system on the user's client machine? Or securing the browser? Or preventing the user from phishing attacks? These are the essential complexities of cyber security. We can modularize, layer and compartment while designing the system and determining accountability, but the attacker knows no such boundaries. Nevertheless, from a requirements perspective web user security is a useful concept that provides critical focus on the end user whose interests are crucial to the evolution and growth of the web.

Clearly, web user security is a complex undertaking. It behooves us to scope the problem so that it does not equate to the entire problem of cyber security. For this article we will demarcate the concern of web user security as being entirely in cyberspace. In other words it does not encompass cyber physical systems [3], although in future it might do so [4]. Our focus is on the ordinary user so we are not concerned with user-related enterprise issues such as BYOD [5]. Our concern is with web security at the micro level of individual users and households, so we are not directly concerned with macro level aggregated threats such as botnets [6] and worms [7]. The issue of digital rights management to protect copyrighted entertainment content is also out of scope.

On the attack side we are primarily concerned with "ordinary" attacks conducted by ordinary criminals and miscreants, much as we are concerned with ordinary users. This excludes, for instance, attacks by nation states or their agents. It also excludes highly targeted attacks such as aimed at specific senior executives and officials. We exclude attacks on critical infrastructure, both cyber physical and cyber only. We specifically do include attacks by organized crime, be it directly on an end user, say by phishing [8], or indirectly, say by a data breach at the web application provider.

We emphasize that our claim about web user security being as good as it will ever get is applicable only under the caveats of the scope discussed above. The situation with critical infrastructure, including cyber physical, and enterprise security targeted by nation state adversaries is certainly not so sanguine. The adversary is much more sophisticated and the payoff worth magnitudes beyond what an ordinary user can deliver. Moreover, the tolerance for a "small fraction" of unfortunate victims amongst web users does not carry over to domains such as critical infrastructure and national and business sensitive information. Web user security as we have defined it here is a very special and very important domain. However, we cannot extrapolate our observations and insights beyond this domain.

In general the term security is understood to include privacy, although the term security and privacy is also commonly employed with the implication that there is

some difference between security and privacy. It is beyond the scope of this article to discuss this issue at length. For our purpose we will understand web user security to encompass elements of user privacy without exhaustively including all its aspects, much as web user security does not fully include user security. In particular protecting against deliberately malicious web application servers is considered out of scope for web user security. Thus a web application that misuses a user's credit card information to make fraudulent charges is not part of what we consider ordinary attacks. Likewise for misuse of a user's personal information by the web application such as address and date of birth, say to perpetrate identity theft. Typically such scams by the web application server tend not to be long-lived due to inability to attract repeat customers and loss of reputation. Nonetheless, we acknowledge that high assurance security and privacy is extremely difficult, perhaps impossible, to achieve.

3. The past

The web is not quite 20 years old. There has been concern about web user security from the early days, if only to foster growth of e-commerce. Among the earliest technologies to be deployed was that of internet firewalls [9] which constrain the flow of internet packets beyond the normal routing rules. Practitioners quickly came to appreciate the limitations of firewalls and turned to cryptographic techniques to establish authenticated secure communication. There was great hope that public-key infrastructure (PKI) [10] would be widely deployed and replace the much-hated and security-challenged passwords as the common authentication mechanism for users. PKI was further expected to enable a number of new services such as secure email. The credit card industry agreed on a common standard PKI for secure credit card transactions on the web, to be built around the newly developed SET (secure electronic transactions) protocol [11]. The dominant browser vendor in the early days (Netscape) brought forth a generic security protocol, called SSL (secure sockets layer) [12] to secure TCP sessions on the web. There was a feeling of confidence amongst security researchers and practitioners that a secure web was technically within reach.

The promise of PKI never did materialize. The SET protocol fell by the wayside. SSL continued to flourish and was renamed in later incarnations as TLS (transport layer security). Its design called for public-key certificates on at least the server side but preferably on both server and client sides. The latter and much more secure mode of SSL never got mass deployment, while the serious vulnerabilities of the former widely deployed mode to man-in-the-middle attacks were identified very early [13]. We continue to live with passwords till today. There have recently been calls to recognize the persistence of passwords and establish a research agenda to understand how to do passwords better [14]. The failure of PKI also effectively eliminated the potential for widespread use of secure email.

In the meantime the internet and the web itself became a conduit for rapid spread of computer viruses [15]. Denial of service attacks emerged as a threat for providers of web services [16] which impacted their users. Users were also swamped by spam emails [17] and phishing attacks [8]. The recent appearance of drive-by downloads [18] makes it possible for attackers to install malware on a user's machine just by visiting legitimate web sites that have been compromised. Vendors and users have become comfortable with frequent security updates of user-deployed software, including operating systems, browsers and applications. The earlier practice of leaving users' computers off-line till such time as the user engaged in internet activity, has been supplanted by an always online presence so as to be current on updates.

The past can be summarized as follows. While the high expectations for web user security based on PKI have not been realized, the spread of web application services and their enthusiastic embrace continues unabated leading us to conclude that currently deployed security technologies are adequate for purpose of good enough security [19]. Many of these technologies were not considered a priori but developed in response to real world attacks such as denial of service, spam, phishing and frequent security updates, which were not anticipated to the degree that they occurred.

4. Security principles

What security principles can be gleaned from this recent eventful history? We propose a number of them here. We reiterate that these are articulated in context of web user security and should not be simply extrapolated outside this setting. While we believe that all of these are more generally applicable, consideration of the extent of their generality is beyond the scope of this article.

Our first principle recognizes that the only thing certain about the future is that it is uncertain.

Principle 1. Anticipation of future web user security technologies, services and attacks cannot be perfect.

The consequence of this principle is that security must be largely reactive. While it would be grossly negligent to deploy web application services without any security, overdesign and expensive provisioning of their security is likely to be premature. The design of the service, the attacks it will attract and the security appropriate to its protection are intertwined issues and must be dealt with concurrently as the service evolves.

The second principle comes from the users' perspective.

Principle 2. Users do not expect perfect security. They will tolerate a small fraction of unfortunate victims in preference to higher cost and inconvenience.

The real world is not perfect in security and safety [20], so it is only natural to expect the same in cyberspace. It has been the hubris of security researchers that technology can make cyberspace more secure than the real world. Perhaps in theory it can, but in practice users will not tolerate the additional cost and inconvenience. This is especially true with respect to web user security where service providers operate in a fiercely competitive environment which drives security and its burdens to a lowest common denominator.

The next two principles come from the attackers' perspective.

Principle 3. Attacks are hard to discover but easy to replicate.

The ingenuity and persistence of attack inventors is truly remarkable. However, once invented, most attacks are easy to automate and replicate. Attacks do not need to be reinvented by every attacker. This is especially true in the realm of web user security. Highly targeted attacks on high value individuals and organizations can require considerable customization and reconnaissance. Attacks targeted at the masses are effective so long as they yield sufficient return to the attacker, which in percentage terms could be minuscule.

Principle 4. There is no upper bound to the sophistication of attacks that will be deployed in the wild (i.e., outside of the laboratory). Attackers always have the next low hanging fruit.

Given our knowledge of the theoretical possibilities of attacks, it is clear that attackers have barely scratched the surface of the iceberg. Historically, more sophisticated attacks have been manifested in the wild as older attacks become less effective due to defenses that have reduced their yield.

Our next two principles are respectively duals of the preceding two principles from the defenders' perspective.

Principle 5. Defenses are hard to discover but easy to replicate.

This would seem to offer a good business opportunity for security vendors.

Principle 6. While there are no perfect defense, defenders always have the next low hanging fruit.

Every new attack can be met with an appropriate defense at a reasonable cost. Defenders must strive to find defenses that will be adopted by the early-adopter segment of the market and then can be grown into mass adoption over time.

Our final principle recognizes the limitations of formal methods and proofs in guaranteeing security properties.

Principle 7. Mathematics and formal methods are useful but can only provide symbol security, not real security.

The term symbol security [21] denotes the fallacy of relying on mathematical arguments for proof of security without due consideration of real world context. Proofs of security of SSL and PKI notwithstanding, there are many insecure ways of deploying these provably secure protocols. The security proofs are of idealized abstractions and not of the real world implementations.

5. Science of web user security

What are the consequences of these security principles? Can we develop a science of web user security which will be conformant with these principles? What are the open questions that such a science should seek to address? We formulate some sample questions below.

Our first set of questions deals with the concurrent nature of designing security, services and attacks due to Principle 1.

Question 1. What is the appropriate level of security that should be provided in a new web application service?

Our current security theories provide no guidance in this regard. How do we characterize "appropriate"? We believe that attempting to quantify this characterization by numerical values is probably premature at this point. Are there qualitative measures we could start with? A related question is as follows.

Question 2. How much and how quickly should we adjust the level of security in light of changes in features of the web application service and the attacks on it?

For example, in reaction to a newly discovered attack we may quickly fix security bugs that enabled the specific attack and issue patches. Should we do more than just fix the one bug when it is possible that several related bugs may be present? Should we try to close a larger class of related attacks?

Our next questions turn to the cat and mouse game between attack and defense.

Question 3. Can we anticipate the next level of attack so we can be proactive about the next level of defense?

Question 4. Does the next enhancement to our defenses address the likely next level of attack?

As a defender one would like to be ahead of the attackers but not too far ahead. There is no point strengthening the front door while the back door is wide open. The current practice more or less boils down to fixing only those vulnerabilities that have been attacked. This results in a completely reactive posture. What is the value in being proactive? How do we demonstrate this value?

Next we deal with the issue of formal methods.

Question 5. What is the real world value of formal methods when they apply only to idealizations of the actual system?

Question 6. When is it appropriate to use formal methods?

Formal methods no doubt provide insights and additional assurance, but only about idealized abstractions of the real system. They cannot be directly applied to a real system due to well-known scalability limitations. How should they be used to provide meaningful value.

Finally, the question about the most important party in the system, the end user.

Question 7. How do we measure the exposure of the user relative to cost and convenience?

Here again numerical quantification may not be the best first step. Are there useful qualitative measures we could start with?

6. Conclusion

This article has defined the concept of web user security as a very special but very important piece of the overall cyber security challenge. Its fundamental claim is that web user security is about as good as it will ever be. In considerable part this is because users do not expect the web to be any more secure than the real world and will settle for much less than perfect security. They will rather accept lower security than bear higher cost and inconvenience, so long as the vast majority of users are relatively safe. Only those web application service providers who conform to this expectation will succeed in the market thereby driving security to the lowest common denominator of "good enough security." Overall this is a good thing since it is encouraging of innovation and evolution, without the impossible attempt at anticipating the future of the web. By looking at the past history of the web we have offered a number of security principles which should be the basis for a science of web security. We have speculated on some of the issues that such a science must address.

References

[1] Cybersecurity game-change research and development recommendations, 2010. <http://www.nitrd.gov/pubs>.

[2] Trustworthy cyberspace: strategic plan for the federal cybersecurity research and development program, 2011. <http://www.nitrd.gov/pubs>.

[3] R. Rajkumar, I. Lee, L. Sha, J. Stankovic, Cyber-physical systems: the next computing revolution, in: Proceedings of the 47th Design Automation Conference, DAC '10, 2010, pp. 731–736.

[4] L. Atzori, A. Iera, G. Morabito, M. Nitti, The social internet of things (SIoT)—When social networks meet the internet of things: concept, architecture and network characterization, Computer Networks 56 (16) (2012) 3594–3608.

[5] G. Thomson, BYOD: enabling the chaos, Network Security 2012 (2) (2012) 5–8.

[6] P. Barford, V. Yegneswaran, An inside look at botnets, in: M. Christodorescu, S. Jha, D. Maughan, D. Song, C. Wang (Eds.), Malware Detection, Advances in Information Security, vol. 27, Springer, US, 2007, pp. 171–191.

[7] N. Weaver, V. Paxson, S. Staniford, R. Cunningham, A taxonomy of computer worms, in: Proceedings of the 2003 ACM Workshop on Rapid Malcode, WORM '03, 2003, pp. 11–18.

[8] R. Dhamija, J.D. Tygar, M. Hearst, Why phishing works, in: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '06, 2006, pp. 581–590.

[9] W.R. Cheswick, S.M. Bellovin, Firewalls and Internet Security: Repelling the Wily Hacker, Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1994.

[10] C. Adams, S. Lloyd, Understanding PKI: Concepts, Standards, and Deployment Considerations, second ed., Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2002.

[11] C. Meadows, P. Syverson, A formal specification of requirements for payment transactions in the SET protocol, in: R. Hirchfeld (Ed.), Financial Cryptography, Lecture Notes in Computer Science, vol. 1465, Springer, Berlin/Heidelberg, 1998, pp. 122–140.

[12] E. Rescorla, SSL and TLS: Designing and Building Secure Systems, Addison-Wesley, 2001.

[13] J. Hayes, The problem with multiple roots in web browsers-certificate masquerading, in: Proceedings Seventh IEEE International Workshops on Enabling Technologies: Infrastructure for Collaborative Enterprises, 1998 (WET ICE '98), 1998, pp. 306–311.

[14] C. Herley, P. Van Oorschot, A research agenda acknowledging the persistence of passwords, IEEE Security and Privacy 10 (1) (2012) 28–36.

[15] F. Cohen, Computer viruses: theory and experiments, Computers & Security 6 (1) (1987) 22–35.

[16] A. Hussain, J. Heidemann, C. Papadopoulos, A framework for classifying denial of service attacks, in: Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, SIGCOMM '03, 2003, pp. 99–110.

[17] L.F. Cranor, B.A. LaMacchia, Spam!, Communications of the ACM 41 (8) (1998) 74–83

[18] M. Egele, E. Kirda, C. Kruegel, Mitigating drive-by download attacks: challenges and open problems, in: J. Camenisch, D. Kesdogan (Eds.), Open Research Problems in Network Security (iNetSection 2009), IFIP Advances in Information and Communication Technology, vol. 309, Springer, Boston, 2009, pp. 52–62.

[19] R. Sandhu, Good-enough security, IEEE Internet Computing 7 (1) (2003) 66–68.

[20] B. Lampson, Computer security in the real world, IEEE Computer 37 (6) (2004) 37–46.

[21] M. Schaefer, Symbol security condition considered harmful, in: Proceedings 1989 IEEE Symposium on Security and Privacy, 1989, pp. 20–46.

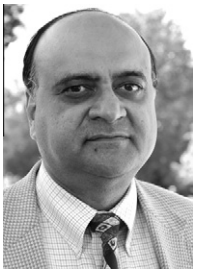

Ravi Sandhu is Executive Director of the Institute for Cyber Security at the University of Texas at San Antonio, where he holds the Lutcher Brown Endowed Chair in Cyber Security. Previously he was on the faculty at George Mason University (1989–2007) and Ohio State University (1982–1989). He holds BTech and MTech degrees from IIT Bombay and Delhi, and MS and PhD degrees from Rutgers University. He is a Fellow of IEEE, ACM and AAAS, and has received awards from IEEE, ACM, NSA and NIST. A prolific and highly cited author, his research has been funded by NSF, NSA, NIST, DARPA, AFOSR, ONR, AFRL and private industry. His seminal papers on role-based access control established it as the dominant form of access control in practical systems. His numerous other models and mechanisms have also had considerable real-world impact. He is Editor-in-Chief of the IEEE Transactions on Dependable and Secure Computing, and founding General Chair of the ACM Conference on Data and Application Security and Privacy. He previously served as founding Editor-in-Chief of ACM Transactions on Information and System Security and on the editorial board for IEEE Internet Computing. He was Chairman of ACM SIGSAC, and founded the ACM Conference on Computer and Communications Security and the ACM Symposium on Access Control Models and Technologies and chaired their Steering Committees for many years. He has served as General Chair, Program Chair and Committee Member for numerous security conferences. He has consulted for leading industry and government organizations, and has lectured all over the world. He is an inventor on 26 security technology patents. At the Institute for Cyber Security he leads multiple teams conducting research on many aspects of cyber security including secure information sharing, social computing security, cloud computing security, secure data provenance and botnet analysis and detection, in collaboration with researchers all across the world. His web site is at www.profsandhu.com.