THE MYTH OF USER EDUCATION

Stefan Görling
Royal Institute of Technology, Sweden

Email stefan.gorling@indek.kth.se

ABSTRACT

Discussions in the security community often tend to end in agreement that the only way to really address many of our current problems is ‘user education’. User education has, in many respects, become the default way to address the fact that our security environment is becoming too complex for us to secure with applications.

But is user education the way forward or is it merely a term used to avoid admitting our failure to create a secure environment for our users/customers? Is there any reason to expect that the user would be interested in educating themselves? Is there any research indicating that user education actually helps?

This paper aims to discuss two questions.

First, should we expect our users to be interested in education? After all, they pay us for taking care of security, so that they can get on with their real work.

Second, do we have any evidence that user education leads to a higher level of security? Do the users actually change their behaviour in a way that mitigates risks? Are the risks we are seeing today addressable by increasing awareness?

INTRODUCTION

For quite a while now, many of us have regarded the largest problem of computer security to be the users. Many discussions have ended with the conclusion that if we could only remove the user from the system, we would be able to make it secure.

And if we cannot completely get rid of the naïve user, we conclude that the best approach to computer security is to educate the user (*cf.* [1, 2]). Although this might sound like the best possible way to increase security in complex IT systems, it is based upon a lack of recognition of the needs and wishes of the users – those who we are supposed to please.

Instead of trying to help the users achieve their goals, we look upon them as the weakest link in security, and thus blame them for the current state of affairs. As Brostof & Sasse write [3]:

> ‘... labelling users as the “weakest link” implies that they are to blame for the current state of affairs. We argue that this is an unfortunate repeat of the “human error” perspective, which blighted the development of safety-critical systems in the mid-eighties: pilots and operators were blamed for accidents whenever they took a wrong action when dealing with a critical incident.’ (p.41)

One source of this problem could be the frame of reference, as suggested by Cooper [4]:

> ‘Programmers aren’t evil. They work hard to make their software easy to use. Unfortunately, their frame of reference is themselves, so they only make it easy to use for other software engineers, not for normal human beings.’ (p.17)

This problem is not limited to the field of computer security. Many specialized departments fight similar issues. Some of us can probably admit that we are reluctant to fill out forms properly; we might fail to provide the proper signatures, thus not respecting the proper procedures imposed by our accountants.

The field of computer security has historically been closely related to the field of computer science. Much of the focus has been on what could be described as more theoretical security aspects; access control and identification. Computer security has therefore come to be associated with barriers for the users, added to which is a layer of annoyance such as hard-to-remember passwords, administrative routines without a proper background explanation and a large number of rules on how not to use your computer equipment.

The increasingly common view of the user as an active, and weak, link in the chain of security has led to expanded security policies. These range from password selection and management routines to lengthy codes of conduct, many of them including rules that are contradictory to common procedures in place (instant messaging, mailing sensitive documents, etc.).

Instead of using this knowledge as a means to mitigate the conflicts, this problem has instead increased the clash between the IT department and its users by sending out contradictory messages, such as:

1. IT department: ‘*Do not perform the actions necessary to carry out your work, it would violate policy X.*’
2. Management: ‘*Do perform the actions necessary to carry out your work, that’s what we are paying you for.*’

It might be unnecessary to state that few things have affected this field as much as the move towards Internet connectivity. Connecting computers to the Internet increases the risk of attacks immensely. Further, unleashing a large number of uneducated users onto this network creates a whole new world, where many old principles of security, which are based upon sealed systems with security-aware operators, are put to the test.

One of the understandings of this shift in power was presented in [5].

> ‘Engineers attempt to solve problems by designing mechanism with predictable consequences. Successful engineering yields bridges that predictably don’t fall down, planes that predictably don’t fall out of the sky, and calculators that give the “right” answer. The essence of engineering is the development and codification of models, techniques and tools that deliver predictable, desirable behaviour.
>
> […]
>
> ‘The **technical** development of the Internet has followed this path. As a community, we focus on design principles that deliver such virtues as robustness, scalability and manageability in the face of complexity, component failures, growth, and other challenges.’

The authors further note that the Internet today is no longer a playground for techies, but has evolved into something larger, something more complex and we are now unable to describe its behaviour and development using simple, predictable arguments of cause and effect.

> ‘The operation of societies follows a different model. Historically the essence of successful societies is the dynamic management of evolving and conflicting interests. Such societies are structured around “controlled tussle” – regulated by mechanisms such as laws, judges, societal opinion, shared values, and the like. Today, this is the way the Internet is defined – by a series of ongoing tussles.’

Analogous, we could argue, to the fact that computer engineering in general once was a strict engineering endeavour where both systems and tasks could be constructed in a predictable way. The computer has evolved into a general utility constructed to perform a large number of tasks, in a large number of settings, cultures and contexts. This has to be taken into consideration when designing secure systems. Old paradigms may fail to address these new types of threat in an effective way.

PREVIOUS RESEARCH ON USER EDUCATION AND SECURITY

Research in the area of computer security has been on-going since the birth of the computer itself (in some cases even long before that). In this section a few studies relevant for the discussion in this paper are presented. This is by no means an attempt to give a full, correct, representation of the field as a whole.

Several studies have been conducted in order to discuss the problems caused by users selecting weak passwords. (*cf.* [6]). Other studies have discussed methods of user education and campaigns to raise awareness (e.g. [7, 8]), but most fail to discuss whether the campaigns have any more than a short-term effect.

A study of user education carried out at Harvard University indicated that people who had anti-virus protection actually imposed a greater threat towards the network than those who did not, as they had a false sense of security, or lack of knowledge about how the software worked, leading to fewer precautions when downloading files or clicking on attachments [8].

A recent user study on phishing attacks found that when users were put to the test of judging whether a website was valid or not, the incorrect conclusion was reached in 40% of the cases [9]. This not only tells us that users have a hard time judging web pages for authoritative content, but also that carrying out that validation for each important web page they visit is not only time-consuming, but would also make it impossible for them to carry out their tasks due to the high number of false positives.

Furthermore, the study concluded that a well-crafted phishing site fooled 90% of the participants, that participants rarely took notice of browser cues such as the address bar, and that a popup warning for fraudulent security certificates was ineffective. There were no significant correlations between vulnerability and factors such as age, sex, education, computer experience, etc.

One attempt to handle phishing attacks has been to add a ‘security toolbar’ to Internet browsers – a piece of software which gives more indicators as to whether a web page is authentic. However, studies have shown that these are ineffective, as users judge web pages by the content that makes up most part of the screen, rather than by the security indicators that are located in the corners of the screen [10]. Even with popup boxes that warned users: ‘*Internet security experts believe that this page is part of a fraudulent site*’, the spoof rate was as high as 10%.

The now somewhat legendary study in which random commuters passing by Liverpool Street station in London were asked to write down their username and password in exchange for a chocolate bar showed that 70% of people would reveal their password under these circumstances [11]. (The passwords were never validated.)

In order to discuss user education as a way to reduce threats it is necessary to split the problem into two parts. The first part is whether users can identify security problems when they are asked to. The second part is whether it is possible to educate the users to exercise their knowledge when carrying out every single task during the day.

Brostoff & Sasse [13] suggest that many of the principles learned from safety research could also be used in security research as they share a number of similarities, most importantly that they are *secondary goals* which must be upheld while performing the primary tasks.

Dahimja, Tygar and Hearst [9] suggest that a traditional cryptography-based security approach is ineffective and that user design must be taken into account. Min Wu and Garfinkel [10] conclude that if the users must make security-critical decisions, it should be integrated into the critical path so that users *have* to deal with it.

SECURITY – THE SECONDARY GOAL

Before discussing further when and how user education might be applicable we should assert whether there is a need for security at all. From an economic perspective, there is no such thing as security for the sake of security. For us to believe that money is worth investing in such things we must see that it affects the bottom line in a positive way. (This might seem cynical and naïve, but so is economics in the most simplistic representations.)

From a corporate perspective we are interested in having our employees carry out their tasks in an efficient way. We want them to serve our customers and deliver results. There is a certain risk that a lack of security will lead to a data leakage or downtime of computer resources, but these are somewhat abstract/theoretical effects which will often affect us in the future, rather than while performing the actual task. We must never forget that both humans in general, and market economy more specifically, are inherently short-sighted. We are burdened by the tyranny of small decisions.

> ‘A market economy makes its large allocations and reallocation of resources on the basis of a summing up of the “votes” recorded by customers in a host of small, individual market transactions. A critical task in appraising the efficiency of such an economy, then, is to determine whether and under what conditions this adding up process produces optimal results. The “smallness” of the decisive, individual transaction – their limited size, scope and time perspective – can, it is argued, be a source of misallocations, in the sense that consumers might disapprove of the larger result thereby produced, if they were ever given the opportunity explicitly to vote for or against it.’ [12] (p. 45)

As such, most people realize that security is a good thing – that we should protect our assets and keep viruses and parasites out of our systems. However, we must understand that this is always a secondary task. No matter how important we believe security to be, the primary task we are trying to carry out will always take precedence over it. If we have to choose between sending a file unencrypted, or failing to deliver on time, we will deliver on time. Our reasoning is very different when studied in perspective rather than by small incremental decisions:

> 'Suppose, 75 years ago, some being from outer space had made us this proposition: "I know how to make a means of transportation that could in effect put 200 horses at the disposal of each of you. It would permit you to travel about, alone or in small groups, at 60 to 80 miles an hour. I offer you this knowledge; the price is 40,000 lives per year." Would we have accepted?' [12] (p.30)

Even though many social engineering attacks could, in theory, be avoided by teaching employees never to give out information and simply to say 'no' to suspicious requests, this advice is often going to be task conflicting with the primary goals of the employee. It is even likely that such a policy would cost a company more in customer losses as levels of customer satisfaction may to decrease. Further, if employees are evaluated on customer satisfaction, they are not only fighting the conflict between primary and secondary goals, but are also given economic incentives not to follow policy.

This same argument could be utilized outside the limits of the corporation. No matter how provocative to the security researcher. Even though the sources of the *Wired* article [13] have been put in doubt, the argument within it could still be taken as representative of another perspective. There exist plenty of pragmatic incentives for one user simply to click on links with little hesitation – even if it means that parasites will infest their machines.

THE LIMITS OF USER EDUCATION

If we are to look to user education as a way of increasing security, and thus strengthening the weakest link, we first need to be aware of the scenarios in which an educated user will make different decisions from the naïve user.

A wise man once claimed: '*If market economy worked, computers would come with pre-installed porn*.'

Many users browse porn regularly, and research indicates that 5% to 10% of the workforce even do it from their workplace [14]. Many more access 'insecure' web pages from their homes. We know this for a fact. Still, we fail to take it into consideration.

Obviously, advising users not to click on such links is no solution to any problem, as we are then trying to achieve the secondary goal by not performing the primary. When did you last consider the design of your software from the use-case of a porn-browsing user?

User education can never protect us to a larger degree than raising security consciousness up to second place on in the attention span. As previous studies have shown, an educated user (such as university students and staff commonly used in these studies) is often unable to judge whether a site is fraudulent, even when asked to actively make this decision, and with the aid of several visual aids.

This means that in the current state of affairs, the technologies utilized during social engineering attacks are so advanced that user education is often unable to avoid it. In some cases (for example sophisticated DNS-poisoned attacks or man-in-the-middle attacks) it is not possible to detect this no matter how much effort you spend evaluating the level of trust ability.

Even more interesting is the level of false positives in these studies. The fact that many users using these techniques take real websites as fakes severely limits them to perform their primary task. Because of the structure we have created, where security features and visual queues are added on top of the actual process, the transaction cost of behaving as an educated user is simply larger than the cost of problems created by the irresponsible acts.

Relating back to the earlier discussion. Even if we can train users to be more accurate when they are actively asked to question the authenticity, we cannot expect them to perform this evaluation continuously as the (economic) incentives steer against such behaviour.

CONCLUSIONS

Perhaps the argument put forward in this paper should be summarised as follows: computer security experts must cease to consider themselves as a theoretical sub-field of computer science, but rather expand and borrow knowledge from various disciplines, including behavioural fields such as Human Computer Interaction (HCI), and a broad range of other disciplines which may help to put security back into context, such as the fields of organization theory.

The field of HCI-SEC is emerging and there is hope for a greater understanding in the future. However, in order for applications to become secure, results from this area must be utilized to a larger degree. Security must not be 'added', it must be integrated in an early stage in order to be integrated with the users' anticipated behaviour.

The user will always circumvent a security model where the security features clash with the tasks the user is trying to carry out. The fact that users are interested in browsing pornographic web pages must be embraced in the design process when designing browsers.

If we are to protect our systems by educating users, if we are to put the question of security in the hands of the users, it is not enough to educate them. The relevant question is whether we can get them to ***act*** as educated users.

This paper has tried to highlight the fact that this is not only a question of knowledge, but of utilizing this knowledge to regulate behaviour. And that the regulation of behaviour is dependent on many more aspects other than simply the amount of education we have given to the user.

Realizing that user education is no general solution to computer security problems is not an argument not to educate and inform the users. It is vital that we help users, spread our knowledge and inspire better behaviour. However, we must not limit ourselves to our narrow frame of reference, expecting them to become security experts – that is, after all, what they are paying us for.

One might claim that nobody is interested in security. And they should not have to be. Security can never be added on later. It can never be the primary goal. It must be built into the systems and processes we utilize. It must be supportive to the user rather than restrictive.

REFERENCES

[1] Mitnick K. The art of deception – controlling the Human Elements of Security. Hungry Minds Inc.; 2002.

[2] Schneier B. Secrets & Lies. John Wiley & Sons; 2000.

[3] Brostoff S, Sasse MA. Safe and Sound: A Safety-Critical Approach to Security. NSPW'01; 2002.

[4] Cooper, A. The Inmates Are Running the Asylum. SAMS Publishing; 1999.

[5] Clark, Wroclawski, Braden, Sollins, Tussle. In Cyberspace: Defining Tomorrow's Internet. SIGCOM'02; 2002.

[6] Adams A, Sasse MA. Users are not the enemy: Why users compromise computer security mechanisms and how to take remedial measures. Communications of the ACM. 1999; 42; 41–46.

[7] McCoy C, Thurmond Fowler R. 'You are the key to security': Establishing a Successful Security Awareness Program. SIGUCCS'04; 2004

[8] Davis K. Saving Users From Themselves: Creating an Effective Student-Oriented Anti-Virus Intervention. SIGUCCS'01; 2001.

[9] Dahimja R, Tygar JD, Hearst M. Why Phishing Works. CHI 2006.

[10] Min Wu R, Garfinkel S. Do Security Toolbars Actually Prevent Phising Attacks?. CHI 2006.

[11] BBC News. Passwords revealed by sweet deal. 2004. http://news.bbc.co.uk/1/hi/technology/3639679.stm.

[12] Kahn AE. The tyranny of small decisions: Market failures, imperfections and the limits of economics. Kyklos. 1996; 19(1); 23–47.

[13] Wired. Spyware on My Machine? So What?. 2004. http://www.wired.com/news/technology/0,65906-0.html.

[14] Websense. Gender differences in employee computing exposed in Websense's seventh annual Web@Work survey. 2006. http://www.websense.com/global/en/PressRoom/PressReleases/PressReleaseDetail/index.php?Release=0605161213

[15] Whitten A, Tygar JD. Why Johnny Can't Encrypt – A Usability Evaluation of PGP 5.0, Proceedings of the 8th USENIX Security Symposium. 1999 ; 169–181.