To err is all too human

Frank Stajano discusses the human factor behind cyber attacks and explains what can be done to make a company's computer network more secure

Is it true that humans are the weakest link in computer security? To simplify the discussion, let us start by leaving them aside. How easy is it for a malicious attacker to get into a corporate network without exploiting its users? It depends. If the attacker's objective is to penetrate any random network in order to attack other targets (for example using it as a base for phishing or spamming), then they will certainly find one that is vulnerable. If, on the other hand, their objective is to penetrate a specific network because of the assets it contains (for example for financial fraud, industrial espionage, extortion or sabotage), then penetration might range from very easy to very difficult, depending on how competently that network is protected.

The first lesson here is that most bad guys do not care about you in particular: your first concern should simply be to avoid being an easy target. When the script kiddies fire off their scanning programs, performing the electronic equivalent of rattling the handles of all the doors in the neighbourhood, you just do not want your own door to open right away. Let them attack someone else. As a baseline, ask your security experts to configure your systems securely, to keep the security patches up to date and to monitor for intrusions. (You have competent security experts you trust, right? OK, we could all do with a few more.)

The second lesson is that it is essentially impossible to make a system totally invulnerable: with enough resources, any system can be penetrated. With government-class "advanced persistent threat" attackers, who have access to zero-day vulnerabilities ("zero-day" meaning that the vulnerability has never been flagged up, so there is no time to mitigate it), and who can inject a trojan into the firmware of the routers you just bought before they are shipped to your premises, all bets are off. There are countermeasures to such threats but are they worth it? Do not ask for invulnerable security: it requires military-style operational security practices that would make it impossible to get any productive commercial work done, and it comes with infinite cost, such as designing and fabricating your own microprocessors, chipsets and motherboards (witness the major security vulnerability discovered in May and affecting essentially every Intel platform built between 2008 and 2017).

So, the third and most important basic lesson is that information security is not cryptography but risk management. The only mature approach to security is to start with an assessment of what your assets are, how valuable they are to you, how valuable they are to potential attackers (not the same thing), how easy they are to compromise,

> *"It is essentially impossible to make a system totally invulnerable – any system can be penetrated*

how expensive they are to defend, and so forth. And then to decide, as a strategic board-level decision, how risk-seeking or risk-averse you are, and how much you are willing to invest in preventive measures (a potentially unnecessary but certain expense) to avert the possibility of being successfully attacked – a loss that might or might not happen.

But let us go back to our original question. It is true that the prevalent network attacks today, phishing and ransomware, are computer-aided frauds that target humans rather than machines. This often prompts the comment that it is those pesky users (your employees) who make the system insecure. I believe this is a wrong-headed attitude that will not make your company safer.

A few years ago, I partnered with Paul Wilson,[1] co-author and star of the popular TV series *The Real Hustle*, to investigate the psychological foundation for frauds and

1. Stajano F and Wilson P (2011), 'Understanding Scam Victims: seven principles for systems security'. *Communications of the ACM*, 70-75. Available at: http://dl.acm.org/citation.cfm?id=1897872.

scams. Many of the hundreds of scams we considered worked by exploiting a handful of psychological traits that are part of human nature. As you might expect, most of today's computer-based frauds exploit the same psychological traits, which were around long before computers. For example, according to what we called "the social compliance principle", society trains us not to question authority, and fraudsters exploit this ("I am your bank and, if you do not verify your login now, we will close your account").

Another trait, the "herd principle", suggests that, when a situation looks dodgy, we feel reassured if many others are engaging in it. Fraudsters, therefore, surround us with accomplices who engage confidently in the suspicious behaviour and implicitly reassure us that it is all legitimate. Think fake reviews on restaurant-rating websites. Subtle manipulation of swing voters is a more ominous example that is topical nowadays, with the involvement of Cambridge Analytica in the Brexit and Trump campaigns.

Other psychological buttons that the fraudsters might push include greed, distraction, time pressure and, cynically, our kindness. The exact set of exploitable psychological vulnerabilities is not so important. What does matter is that the hundreds of frauds we examined all exploited combinations of the same few psychological vulnerabilities. The crucial insight is that these vulnerabilities are part of human nature, and that they are there for a reason.

Why did medieval suits of armour have joints? Were they not the weak point that enemy swordsmen always targeted? Yes, but without joints you could not move. With psychological vulnerabilities,

> *Many of the scams worked by exploiting a handful of psychological traits that are part of human nature*

it is similar: each exists for some good reason. We react to time pressure by switching to quick heuristics rather than full logical reasoning because that is what saved our ancestors from becoming lunch when they heard the roar of the sabre-toothed tiger.

The most important lesson for the security engineer is that you cannot hope to remove the vulnerability just by telling users what they should do, especially against an adaptive adversary who will invent a different scenario to trigger the same reaction. An explanation, or blaming the user's gullibility, will not change human nature and will not get the desired results. A better solution is to design the system with the expectation that users will continue to be human: the technical defence must protect the system even if users react to well-crafted malicious stimuli according to those predictable failure modes.

Is it possible to manage internet security without restricting the sites that can be accessed? Sure. This may not apply to basic and repetitive jobs but, where appropriate, productivity and morale are increased by empowering team members to take the right decisions: the warning ("we rate this site as 80 per cent likely to be fraudulent; are you sure you wish to proceed?") might be coupled with a logged and explicit assumption of responsibility ("please explain why you need to access the site despite the warning, and click here to accept responsibility for the consequences").

A similar approach, without the logging, is used by the Firefox web browser: when a user attempts to open a website that others have reported as fraudulent, the browser instead puts up a red page with a conspicuous warning. It is still possible to proceed to the website if desired, but it cannot happen inadvertently. Similarly, Firefox makes it difficult to visit a website whose Transport Layer Security certificate (a cryptographic protocol that provides communication security over a computer network) fails to verify, except if users confirm they know what they are doing by clicking the correct sequence of options in a purposefully technical dialog box.

A useful guideline to protect your system, while acknowledging the existence of these universal psychological vulnerabilities, is to be parsimonious in imposing security policies on your employees. Security measures are "a tax on the honest": something annoying that gets in the way of employees doing their work.[2]

Analysis of user response to security policies suggests that every person had a "compliance budget": a finite amount of goodwill that is gradually depleted every time they comply with some annoying corporate policy. Once it runs out, the person becomes fed up and stops cooperating. Spend that budget wisely and do not annoy your employees with trivia, such as gratuitous password requests that interrupt the workflow, if you want them to have some goodwill left to comply with the policy items that really matter.

> *Many of the scams worked by exploiting a handful of psychological traits that are part of human nature*

To conclude, here are three parting thoughts for security, three important goals to which I have wholeheartedly devoted my academic and entrepreneurial efforts.

First, we need more security experts. In addition to my university teaching, with support from academia, government and industry, I started two hacking competitions, Inter-ACE and Cambridge 2 Cambridge, to raise a new generation of skilled cyber defenders.

Second, computer people have a moral duty to build a digital society that is secure and fair for its citizens. I embarked on a crusade to eliminate passwords because we cannot give people an impossible task and blame them for not completing it. We are now turning this project into an open-source start-up, Pico – a hardware token that relieves the user from having to remember passwords and PINs .

Last but not least, users are a crucial component of the system. My message to system architects is that, rather than blaming users, understanding and accepting human nature is a necessary step towards making systems truly secure. ■

Frank Stajano *is a tenured academic at the Faculty of Computer Science and Technology of the University of Cambridge, where he is the head of the Academic Centre of Excellence in Cyber Security Research. He is a founding director of two cyber security start-ups, Cambridge Cyber and Pico Authentication, and a founder of the Inter-ACE and Cambridge 2 Cambridge cyber security competitions*

2. Beautement A, Sasse M and Wonham M (2008), 'The Compliance Budget: managing security behaviour in organisations'. *New Security Paradigms Workshop*, 47-58. Available at: http://dl.acm.org/citation.cfm?id=1595684.