Social Engineering 2.0: What's Next

By Markus Jakobsson

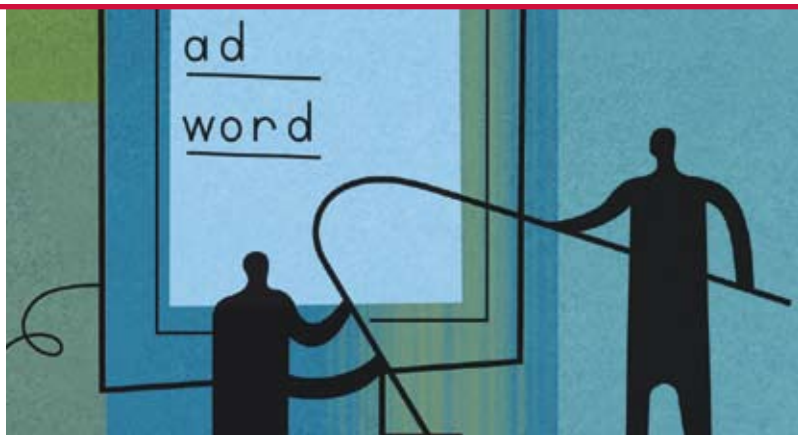

Although social engineering has probably been around since the dawn of human civilization, many are concerned that it is currently transforming and wreaking havoc on the Internet. In this article, we'll offer some predictions about what may come next.

Few would disagree that the current crimeware wave is fed by economic incentives. The current state of affairs stands in stark contrast with the past. Early viruses were simply an expression of intellectual curiosity, competitiveness, and maybe a bit of ennui. The case is even clearer as we turn to click fraud and phishing. What other possible motivation is there other than to make a shady buck or two? (Or often a whole lot more.) The same holds for spam in its various forms. If spammers couldn't make money from it, there would be no spam. It is, therefore, rational to consider the ways that criminals can monetize abuses of existing Internet features so that we can predict trends in fraud.

Internet Fraud: A Socio-Technical Crime

An increasing number of experts recognize that fraud is no longer only a technical matter, but that to an increasing extent there is also a social engineering component. Phishing is a prime example of this, but not the only one. It is more and more common these days to see crimeware attacks that hinge on social engineering for installation. A recent example of this is the so-called Better Business Bureau scam, shown in Figure 1. In this phishing attack, a potential victim receives an email appearing to come from the Better Business Bureau and relating to a case against the organization of the recipient. The attachment, which supposedly contains the details of the complaint, in reality contains a Trojan downloader. To make matters worse, these emails are often sent to people high up in the targeted organization—often to individuals who deal with customer complaints on a daily basis.

Defenses Shape Attacks

From the point of view of criminals, Internet fraud is a relatively safe and comfortable crime. Apart from being a crook's telecommuting dream, Internet fraud offers scalability, high profits, and very low traceability—and thus very limited risk. It is no wonder that Internet fraud has taken off. Now to understand the attacks, we must also understand the defenses. It is clear that the crimes are being fought on three separate planes today: technical features (such as anti-virus software, spam filters, and anti-phishing browser plug-ins); educational campaigns (such as those run by FTC, eBay, SecurityCartoon.com, banks, and the Carnegie Mellon University Usable Privacy and Security Laboratory (CUPS) group); and finally, by legal means. The legal efforts typically involve tracking origination, raiding drop boxes, and finally, prosecuting offenders.

Whereas the technical and educational efforts—if successful—result in a lower yield to criminals, the legal efforts result in a higher risk. These risks are a big deal, especially given how well Internet fraud scales. It is, therefore, fair to assume that the next frontier in Internet crime will involve a component that makes it less traceable. We will make that assumption here, and investigate what that could mean for the future. We will do this by considering two types of highly untraceable attacks, neither of which has occurred to date, but both of which are waiting to happen. But first, to truly understand the importance of the legal aspect, we will take a slight tangent and review why "ransomware" never became the calamity people thought it would be.

Ransomware Fails

In the late 1990s, researchers at Columbia University posited that the next wave of malware might attempt to hold the files on the victim's computer hostage by encrypting them using a public key carried in the malware body and demand a ransom to get the secret key—to regain access to the encrypted files. Years later, the Archiveus Trojan carried out an attack just like that, although with a small difference: it used symmetric-key cryptography instead of a public key. The attack was foiled when the Trojan was reverse engineered and the encryption/decryption key was extracted and distributed to anybody who was attacked. But maybe the Archiveus attack would not have succeeded even if it had used public-key cryptography (which, by its nature, would have prevented anyone's reverse-engineering the decryption key from the code, since it would never be contained there in the first place). The reason Archiveus might have failed is not technical, but lies in the monetization aspect: there was no way the criminals could have safely collected the ransom without being traced.

Vandalware Strikes

With the ransomware example in mind, let us now consider a new type of attack, which we can call vandalware. This attack does not carry out vandalism for fun or defiance, but rather for profit. Here is what the criminal would do: first, he or she would select a company to target, and use data-mining techniques to get as much detailed information as possible about vulnerable employees. By vulnerable employee, we mean an employee with access to sensitive data or access to the web page façade of the company. From the vulnerable employee, a vandal might learn about the internal structure of the company, the names of key employees, and the format of email addresses. Second, the criminal would buy put options for that company. (We are assuming that it is a publicly traded company.) A put option is a financial instrument that increases in value if the corresponding stock falls in price; investors and speculators use put options to turn a profit from an insight that a given stock is soon to lose value. Most likely, other investors, not just the criminal, would also buy put options, especially if the stock of the targeted company has a reasonable trading volume. Third, the criminal would unleash an attack against the company, perhaps by sending selected employees spoofed email appearing to come from another employee, such as their boss: "Hi Jim. Please take a look at the attached PowerPoint slides and let me know what you think. If possible, I'd like a quick assessment by tomorrow morning. Hope you can make it." Or perhaps from a system administrator: "There is a dangerous new computer virus, and our systems are not properly patched yet to defend against it. Please install the attached program on your computer right away to help us stay secure. Do this as soon as you can."

And what would happen if someone were to open or execute the attached file? Assuming that the email would not end up in the spam folder in the first place and that the anti-virus system would not catch it, we would have an infection—on a computer with access to sensitive data or to the corporate web site. What if some of that sensitive data were to make its way onto the Internet, maybe even onto the web site of the company itself? There would be a public uproar, and the stock price would suffer. Then the criminal would exercise his or her put options, cashing in on the previous bet that the stock of the company would go down in value. Doing so does nothing to make the attacker traceable, as every investor with put options would be in the same situation. Who is the criminal? Nobody would be able to tell.

Faking the Clicks

Click fraud is another common type of online fraud. It takes advantage of the fact that when a consumer clicks on an advertisement, the advertiser pays a commission both to the web site displaying the ad and to the portal that provided the web site with the ad. Related types of fraud take advantage of advertising in which money is transferred when the consumer views a banner ad (whether or not he takes action), and other approaches in which a sale or other action is generated as a result of someone viewing an ad. The objective could be to profit from these transfers (criminals benefit when their web sites display the advertisements) or to drain the advertising budgets of competitors

BBB complaint case

BBB CASE #569822971

Complaint filed by:	Michael Taylor
Complaint filed against:	Business Name: Contact: BBB Member:
Complaint status:	-
Category:	Contract Issues
Case opened date:	2/28/2008
Case closed date:	-

*** Attached you will find a copy of the complaint. Please download and keep this copy so you can print it for your records.***

On February 26 2008, the consumer provided the following information: (The consumer indicated he/she DID NOT received any response from the business.)

The form you used to register this complaint is designed to improve public access to the Better Business Bureau of Consumer Protection Consumer Response Center, and is voluntary. Through this form, consumers may electronically register a complaint with the BBB. Under the Paperwork Reduction Act, as amended, an agency may conduct or sponsor, and a person is not required to respond to, a collection of information unless it displays a currently valid OMB control number. That number is 502-793.

<Complaint_569822971.doc>

Figure 1: The Better Business Bureau scam. The email contains an infected attachment, which the attacker hopes will be opened by the recipient.

(when the competitors are the advertisers from whom money is transferred). Often, criminals generate traffic in an automated manner, making it appear as though real people viewed the ads. Automation can include some form of malware, such as a botnet. Another common approach is for criminals to hire people to click on selected ads; this is referred to as a "click farm."

We will now describe how social engineering can be used in a new kind of click fraud attack. First, we'll begin by explaining a common scenario that is not click fraud:

- **Scenario 1** Standard web site. Consider a legitimate web site that provides some service, and that displays advertisements which relate to this service. The contents of the ads are typically determined in an automated manner by the ad portals (for example, Google and Yahoo) by automatically reviewing the contents of the web site and selecting ads on topics related to the contents. If the web site is devoted to cooking, for example, then the ads may relate to pots, pans, and coffee machines. These sites also commonly place ads that bring in traffic. Thus we would expect to see ads that use keywords such as "knife," "Calphalon," "Teflon," and similar terms. There is nothing unusual about this type of site.
- **Scenario 2** Using arbitrage. Consider now a second web site that has content which selects ads corresponding to the keywords "find a attorney." (We do mean "a," not "an." We'll explain why soon.) The site can do this by having lots of text (whether visible or not) that repeats this phrase. At the time of writing this article, the cost for this type of strategy falls in the range of $1.07 to $7.05 per keyword. The exact price depends on the venue, the time of the day, and, of course, the competing bids for the keywords because all keyword prices are established by auctions. Thus, if a user clicks on an ad on this site, the owner of the corresponding ad would pay that amount to the portal, which in turn would transfer the amount—minus commission—to the web site that displayed the ad.

  Next, imagine that the site in question places an advertisement using the keyword "find an attorney." The only difference here is the article—"a" versus "an." The price range for this keyword is $0.87 to $3.82. We will assume that the web site pays $2.00 for each visitor it brings in, and receives $4.00 for each visitor who clicks an ad on the site. As long as 50 percent of the visitors who arrive via the $2.00 ad click on a $4.00 ad, then the site makes a profit, without providing any service. This is referred to as keyword arbitrage. It is not quite click fraud, but it's close, as we shall see.
- **Scenario 3** An attack using social engineering. Now we'll see how a criminal might use social engineering and extend the arbitrage technique to make a spectacular profit. Let's assume that the criminal produces a web site that generates the keyword "mesothelioma" (a rare form of cancer caused by asbestos exposure). As we write, this is a Google keyword worth $63.42. The criminal buys traffic for the keyword "asthma" ($0.10) to bring visitors to his site. If at least one 634 people coming to his site clicks on the mesothelioma ad, he makes a profit. But why would someone do that? Assume that the web site content is an article, apparently written by a medical doctor, asking "Did you know that 10 percent of asthma sufferers are at risk to contract mesothelioma?" Although this is not a truthful statement, it will make many people who are concerned with asthma and who are unaware of what mesothelioma is to do exactly what the criminal wants—to click. Will half of all visitors fall for it? With a thousand visitors per day, that means a daily profit of more than $30,000. Even with less conspicuous keywords, the criminal can still make a pretty decent profit.

What makes the three scenarios differ is the intent—and the use of social engineering. From the ad providers' perspective, these three scenarios are very similar in structure. A visitor comes in, reads content, and clicks an ad. Although it is possible to match keywords coming in and going out to find anomalies, it is also possible for criminals to use one service provider to bring traffic in, and another one to carry traffic out. This strategy makes it hard to detect and stop this kind of attack, especially if it is carried out at small scale using a large number of sites.

Conclusion

Social engineering on the Internet is here to stay. We have already witnessed its effects through phishing scams, and we are starting to see how criminals use social engineering to improve the efficiency of spam and crimeware. Even more skilled applications than we currently see are just around the corner, we fear, as is the use of social engineering for other types of fraud—such as click fraud. We can design technical countermeasures with this in mind, and understanding the ways attacks are likely to occur will help improve the defenses. But we must also understand that our strategy requires better user interfaces, better procedures, stronger legislation, and improved education. The good guys still have a lot of work to do.

Dr. Markus Jakobsson is a Principal Scientist at Palo Alto Research Center. He researches phishing and countermeasures, click fraud, the human factor in security, cryptography, network security, and protocol design. He is an editor of *Phishing and Countermeasures* (Wiley, 2006) and co-author of *Crimeware: Understanding New Attacks and Defenses* (Symantec Press, 2008).

Image courtesy of PARC, photographer Brian Tramontana