Can Phishing Be Foiled?

Understanding the human factors that make people vulnerable to online criminals can improve both security training and technology

By Lorrie Faith Cranor

KEY CONCEPTS

- A form of online crime that lures people into giving up personal or corporate information, phishing is a growing security threat that already costs victims billions of dollars every year.

- Because phishing exploits human vulnerabilities, studying the factors that make people fall for phishing scams can improve antiphishing training and technology.

- The combined efforts of law enforcement, computer security experts and computer users are needed to reduce the success of phishing.

—*The Editors*

Over just a few weeks, I received e-mail messages from several banks warning me that my online banking services were in danger of being deactivated, from eBay telling me that I needed to change my password, from Apple complaining that I had unpaid bills for music downloads, from an airline offering me the opportunity to earn a quick $50 for filling out a survey and from the Red Cross asking me to contribute money to help earthquake victims in China. These messages were all very convincing and looked authentic. Except for the eBay message, however, they were all fraudulent e-mails known as "phish."

Phish e-mails are constructed by con artists to look like legitimate communications, often from familiar and reputable companies, and usually ask victims to take urgent action to avoid a consequence or receive a reward. The desired response typically involves logging in to a Web site or calling a phone number to provide personal information. Sometimes victims need only click on links or open e-mail attachments for their computers to become infected by malicious software—known as malware—that allows phishers to retrieve the data they want or take control of the victim's computer to launch future attacks. Although the details of phishing scams can vary, the result is usually the same: thousands of unsuspecting victims give information to criminals who then use it to break in to their accounts and steal their money or identities, or both.

The Anti-Phishing Working Group, an international consortium of organizations committed to wiping out Internet scams and fraud, keeps track of phishing activity, including the number of unique phishing Web sites detected every month. In 2007 monthly totals ranged as high as 55,643. During each month in 2007, anywhere from 92 to 178 different company brands were "phished"—meaning their names or logos were used to fool victims into thinking they were dealing with a trusted institution. According to research and consulting firm Gartner, an estimated 3.6 million Americans fell victim to phishing last year, leading to losses of more than $3.2 billion.

With so much at stake, the computer security community has been scrambling to develop technologies to combat phishing, such as filters for e-mail and Web browsers that flag phishing attempts. Although such software has helped stop many attacks, phishers are constantly evolving their tactics to try to stay a step ahead of such technologies. Since phishing plays on human vul-

nerabilities—a successful attack requires a victim to succumb to the lure and take some action—it is also not strictly a technological problem. For that reason, my research group at Carnegie Mellon University is studying the best ways to teach people to recognize and avoid phishing scams. This research, in turn, is informing our design of antiphishing software so people are more likely to use it correctly. Because human factors are a critical element in the success of phishing attacks, we have found that they can be essential weapons to foil phishers as well.

Teachable Moments

When we began trying to understand why people fall for phishing attacks in 2004, my coworkers Mandy Holbrook and Julie Downs recruited people on the streets of Pittsburgh to interview. Most were unaware of phishing and assumed the term had "something to do with the band Phish." Others knew about e-mail scams that used the names of financial institutions, but they did not realize that messages seemingly from retailers might also be fraudulent. Most

people had little sense of how to identify a phishing e-mail and tended to rely on superficial features, such as a logo or a professional look, to determine whether it was legitimate. They also did not understand the security messages displayed by their Web browsers and did not know how to use cues in Web addresses and within e-mail messages to judge their authenticity.

After confirming that a great need exists to educate Internet users about phishing, our next step was to review existing antiphishing training efforts to try to understand why they apparently do not work. We found a wide range of Web sites devoted to antiphishing training provided by companies, government agencies and industry associations. Some of these included a lot of technical jargon and more information than a nontechnical computer user was likely to digest. A few sites provided good background to raise awareness of the phishing threat but little in the way of actionable advice about how people could protect themselves. In fact, we found in a laboratory study that some of the best antiphishing materials in terms of raising awareness left peo-

[THE AUTHOR]

Lorrie Faith Cranor is an associate professor of computer science and of engineering and public policy at Carnegie Mellon University, where she directs the Usable Privacy and Security Laboratory and leads an antiphishing research project. She also recently cofounded Wombat Security Technologies, Inc., to commercialize products developed by her group. Cranor has published four books and scores of research papers about online privacy, phishing, spam, electronic voting, and other topics related to computer security and usability. She hopes one day people will no longer consider "usable security" to be an oxymoron.

ple overly suspicious of legitimate Web sites.

Worse still, messages that companies send to their employees or customers to warn them about phishing attacks are largely ignored. We did learn, however, that it was much easier to get research volunteers to read e-mail that looks like a phishing message than to get them to read a security-related e-mail. Our studies seemed to show, therefore, that awareness of phishing in the abstract does not translate into protection but that firsthand experience with phishing could provide a powerful teachable moment.

With some of these insights in mind, members of my team, Ponnurangam Kumaraguru, Alessandro Acquisti and others, developed a training system called PhishGuru, which delivers antiphishing information *after* users have fallen for simulated phishing messages. The program incorporates a set of succinct and actionable messages about phishing into short cartoons, wherein a character named PhishGuru teaches wouldbe victims how to protect themselves. In a series

of studies, we demonstrated that when people read the cartoons after falling for the simulated phishing e-mails that we sent to them, they were much less likely to fall for subsequent attacks. Even a week later our test subjects retained what they had learned. In contrast, those who read the PhishGuru cartoons sent to them by e-mail, without experiencing a simulated attack, were very likely to fall for subsequent attacks.

Extending this principle, Steve Sheng, one of my graduate students, also developed an online training game called Anti-Phishing Phil that teaches people how to identify suspicious Web site addresses while providing an experience of getting "caught" by a phisher. Players take on the role of Phil, a young fish that must examine the Web addresses associated with the worms he encounters and determine which are safe to eat. When Phil tries to bite a worm with a fraudulent address, he gets caught on a fishing hook and hauled out of the water. An older and wiser fish then appears on the scene and explains where

[THE BASICS]

HOW PHISHING WORKS

Phishing can take several forms, but the goal of phishers is always to lure people into giving up information by making them think they are interacting with a known and trusted company or person. Phishers are criminals seeking to profit from the information they acquire. In some cases, they also implant malicious software that controls a computer so that it can participate in future phishing scams.

THE LURE

1 Phisher sends e-mail that appears to be from a source the user trusts and urges quick action, such as clicking on a link in the e-mail or opening an attachment.

Phishing e-mail
Phisher
Variations:
• Social-networking messages
• Instant messages
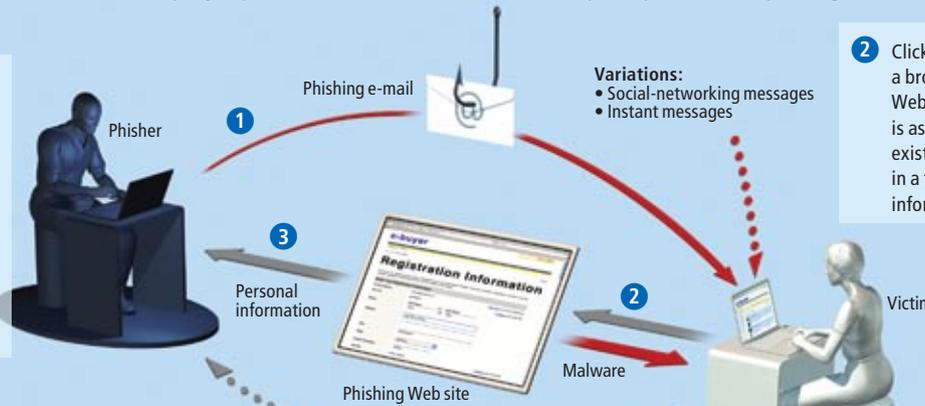
2 Clicking the link opens a browser window on a Web page where the user is asked to log on to an existing account or to fill in a form with personal information.

3 Personal information

Phishing Web site
Malware
Victim

3 Information is relayed to the phisher.

Personal or employer information

THE PAYOFF

6 Personal data allow phishers to steal identities and money or government and corporate secrets.

6

Botnet

5

MALWARE

4 Merely opening the Web page or an attachment to the e-mail might also download malicious software onto the user's computer. This malware may spy on the user's activities or mine the computer's data and relay them to the phisher.

5 Malware can also cause a user's computer to automatically send out more phishing e-mails or turn it into part of a "botnet"—a network of compromised computers—that hosts a phishing Web site without users' knowledge.

GEORGE RETSECK

SPOTTING PHISHY E-MAIL

Phishers' preferred way to lure victims is through a mass e-mail, constructed to look like an authentic message from a well-known company. Computer users often trust such e-mails based on the presence of a familiar brand name or logo. These phishing messages do contain clues that can help identify them as fraudulent, however. Many are visible to the attentive user; others are detectable by software filters.

VISIBLE CLUES

Professional-looking design. Familiar corporate logo.

Urgent message requiring action.

Account status threat.

When cursor is held over the link, the visible address does not match underlying link shown in mail program status bar.

Subject: Wombank Urgent E-mail Verification
From: "Wombank" <creditcards@Wombank.com>
Date: Mon, November 24th, 2008 3:12 pm
To: janywhere@sciam.com
Priority: Normal
Options View Full Header | View Printable Version

❖ WOMBANK

Dear Wombank Member,

This email was sent by the Wombank server to verify your e-mail address. You must complete this process by clicking here or on the link below and entering in the small window your Wombank User ID and Password. This is done for your protection — because some of our members no longer have access to their email address and we must verify it. For security reasons, if your account information is not verified within next 72 hours we are required by law to limit access to your account.

To verify your e-mail address and access your bank account, click on the link below. If nothing happens when you click on the link, copy and paste the link into the address bar of your web browser.

http://www.wombank.com/verifyEmail

Thank you
Accounts Management

http://www.wombank-accountonline.com/accountonline/AccountSummary.htm?verify=email

FILTER FODDER

HTML or JavaScript: Both forms of code appear in many legitimate e-mails, but phishing e-mails would be difficult to construct without them because they allow information, such as a linked address, to be hidden.

"Here" link: Phishers often use legitimate corporate links within an e-mail to lend an authentic feel, but the primary link they intend the victim to click will have a different domain address.

Domain age: A filter can search domain registries to see if the linked Web site is newly created.

Phil went wrong [*see illustration on next page*]. Through both laboratory and field studies, we have shown that the game makes a significant difference in users' ability to identify phishing sites. Comparing their performance before and after the training, we saw a drop in the number of false negatives, phishing sites mistakenly deemed to be legitimate, and false positives, legitimate sites judged to be phishing sites. The game players also outperformed participants who trained with a tutorial or with materials from other sources.

Although we have shown that we can teach people to protect themselves from phishers, even those educated users must remain vigilant and may require periodic retraining to keep up with phishers' evolving tactics. The Anti-Phishing Working Group reported that the number of programs and Web sites devoted to infecting computers with password-stealing code jumped dramatically this year, for instance. "Spearphishing" attacks, which are tailored to their victims, are another growing trend. These can take the form of e-mails sent to the employees of a company that appear to have come from a manager in that company, leading the employees to trust the message and open its attachment. Information available on corporate Web sites and through social-networking sites can help attackers to craft these targeted messages.

Because phishers are such determined criminals, individual computer users cannot be expected to defend themselves alone. Our group also develops automatic filters that can identify likely phishing attacks. But in this work, too, we have found that human responses can be critical to a filter's success.

A Multipronged Defense

Many browser programs already include built-in security filters or can work with add-on programs for detecting suspicious Web sites. Yet even when antiphishing software tools are able to correctly identify phishing Web sites, they may still be ineffective if users choose to ignore their warnings. To understand why some people do not heed such security messages, another of my graduate students, Serge Egelman, sent simulated phishing e-mails to the volunteers participating in our research. When the recipients fell for the phishing messages and clicked on the links, warnings were triggered in their Web browsers. Egelman then found that all the participants who used the Mozilla Firefox 2 browser heeded the warnings, whereas those who used Internet Explorer 7 (IE7) often ignored them.

A high rate of false positives can undermine a filter's credibility and cause users to ignore it after a while.

We determined that the dramatic difference in the responses of the two groups was largely attributable to the fact that the IE7 users either did not notice the warning messages or confused them with less severe warnings. Microsoft appears to have learned this lesson too, and the next generation of the Internet Explorer browser, IE8, now has clearer warning messages that are similar to those shown by Firefox.

In addition to clarity, we have found that accuracy is another critical factor affecting whether users respect the warnings of automatic filters. A high rate of false positives can undermine a filter's credibility and cause users to ignore it after a while. The antiphishing filters we tested employ a mixture of approaches to identify phishing messages and Web sites. Most commercially available tools use a blacklist of known phishing sites, for example. As new sites are reported, they are quickly added to the lists. Some tools also use a white list of known legitimate sites.

Most filters do not rely solely on such lists, however. Some analyze each Web site a user visits and apply a combination of heuristics to determine whether the site is likely to be fraudulent. A few of these are the same kinds of signals we train people to look out for, such as Web addresses beginning with all numbers or addresses that look similar to those of well-known brands. Other features the filters scrutinize include things people could not readily see; for example, the tool may take into account the age of the Web site because phishing sites are typically extremely short-lived, remaining active for as little as a few hours to days or weeks.

AVOID FALLING FOR PHISH

- Never click on a link in a suspicious e-mail or instant message, particularly one asking for personal information. If you do business with the purported sender, open your browser and type the company's usual Web address yourself.

- Look carefully at Web addresses for subtle errors, such as "Annazon.com." Learn to parse Web addresses for other clues to the site's legitimacy [*see examples below*].

- If unsure of a Web site, perform a Google search for the company. The address of the suspicious site is unlikely to appear in the top results, whereas the real company Web site will.

- See consumer tips and resources from the Anti-Phishing Working Group at: **http://apwg.org/advice** Play Anti-Phishing Phil at: **http://cups.cs.cmu.edu/antiphishing_phil**

The time element can make a difference in the performance of the filters that rely heavily on blacklists. Our group recently tested eight consumer antiphishing programs, for instance, by feeding them fresh phishing URLs. We discovered that most of the blacklist programs caught fewer than 20 percent of the phishing sites when we tested them within minutes of receiving the URLs. After five hours, most could detect about 60 percent of the active phishing sites. The programs that used a combination of blacklists and heuristics fared much better, with one detecting almost 90 percent of phishing attacks from the beginning of our test.

Our group has been working on programs that employ machine-learning techniques to detect phishing e-mail. This is a common approach used to detect spam e-mails, but spam detectors are not very accurate when it comes to phishing messages, which generally look legitimate. A member of our team, Norman Sadeh, has been leading an effort to develop a tool, which we originally called PILFER, that analyzes e-mails for a variety of features that may be indicative of phishing. For example, phishing e-mails often contain hyperlinked text that looks like the address of a well-known Web site, but the actual embedded computer code directs users to the attacker's site. In addition, the Web addresses in phishing e-mails often contain five or more dots and point to recently registered domain names. Not all phishing e-mails contain these features, however, and sometimes legitimate e-mails contain them as well. Researchers therefore train the program—which we have renamed Phish-

ANTI-PHISHING PHIL, **an online game, teaches users to identify the addresses (URLs) of phishing Web sites by having a player take the role of a young fish named Phil that must choose to eat or reject worms associated with URLs (left). During and after each round,** a player gets feedback on the choices and new tips (*right*). **In laboratory studies, the ability of subjects who had played the game to distinguish legitimate from fraudulent URLs improved nearly twice as much as that of those trained with standard materials.**

© 2008 SCIENTIFIC AMERICAN, INC.

FILTERING PHISH

To be effective, phishing filters need to apply criteria flexible enough to work in the face of evolving phishing tactics. A filter the author's group created to recognize phishing Web sites works with 95 percent accuracy in laboratory tests. In addition to applying some commonly used heuristics, the filter extracts a "lexical signature" of important words on the page, then performs a Google search to find a legitimate site containing those words.

HEURISTIC	SUSPECTED PHISHING CRITERIA
Age of domain	Less than or equal to 12 months
Known images	Page contains known logos but is not a domain owned by the logo owner
Suspicious URL	URL contains @ sign, hyphen, an IP address or more than five dots
Suspicious links	Link on page contains @ or hyphen
Forms	Page contains a text entry field
Lexical signature search result	URL does not match address of Google-ranked legitimate page

➡ MORE TO EXPLORE

Phishing Exposed. Lance James. Syngress, 2005.

Phishing and Countermeasures. Edited by Markus Jakobsson and Steven Myers. Wiley, 2007.

Anti-Phishing Phil: The Design and Evaluation of a Game That Teaches People Not to Fall for Phish. Steve Sheng et al. in *Proceedings of the 2007 Symposium on Usable Privacy and Security;* July 18–20, 2007.

Behavioral Response to Phishing Risk. Julie S. Downs, Mandy Holbrook and Lorrie Faith Cranor in *Proceedings of the 2nd Annual eCrime Researchers Summit,* pages 37–44; October 4–5, 2007.

The "Supporting Trust Decisions" Web site has information and links related to Lorrie Faith Cranor's laboratory research:
http://cups.cs.cmu.edu/trust

Patrol—by providing it with a large collection of legitimate and phishing e-mails so it can analyze these messages and learn which combinations of features are most likely to appear in phishing e-mails. In our most recent experiments, Phish-Patrol was able to detect more than 95 percent of phishing messages while only triggering false positives for around 0.1 percent of legitimate messages.

We have also combined some of the features used in PhishPatrol with other approaches to detect phishing Web sites. Jason Hong has been leading our group's development of a tool called CANTINA, which analyzes the content of a Web page in combination with other heuristics to determine whether or not the page is part of a phishing site. CANTINA first employs a well-known information-retrieval algorithm to identify five terms that are important on a given Web page but relatively uncommon on the Internet overall. For example, on an eBay log-in page, this "lexical signature" might be, "eBay, user, sign, help, forgot." If you were to search for these five terms using Google, the legitimate eBay log-in page would appear among the top search results. Phishing Web sites that have replicated the eBay log-in page are unlikely to appear because one of the criteria Google's proprietary algorithm uses in ranking a Web page is the number of links to it from other pages on the Internet, so legitimate pages are more likely to be in the top results. This approach is not foolproof, however, especially if a legitimate site was recently created; thus, it is only one of several features that CANTINA considers in assessing a Web site.

The Evolving Threat

We in the computer security community are not the only ones continually seeking to improve our performance. As antiphishing technologies get better, attackers adapt their tactics. Phishing messages are now being sent via instant messenger and mobile phone text messaging. Phishers are using online games such as World of Warcraft and messaging features of social-networking sites such as MySpace and Facebook to lure their victims. Another type of phishing attack involves setting up Wi-Fi access points in public places and spoofing (imitating) the log-in pages of legitimate Wi-Fi vendors. These attacks are used to steal victims' passwords as well as to infect their computers with malware.

Organized gangs of phishers leverage thousands of compromised computers as launch points for their attacks. For instance, a group believed to be based in eastern Europe and known as the "Rock Phish gang" uses compromised computers to relay messages to phishing sites. It can thus send phishing messages that appear to originate from those computers, masking the Web address of the actual phishing site and making it difficult for law enforcement to find the real source of the attack.

Another evasive tactic this gang uses is a system that security experts have dubbed "fast-flux," in which the phishers manipulate Internet domain name servers to continuously change the numerical addresses corresponding with phishing domain names.

Phishing is only lucrative, of course, if phishers have a way of converting stolen credit-card numbers and other credentials into cash. Thus, phishers often recruit "mules" by advertising for people to fill work-from-home jobs or by befriending Internet users and convincing them that the phishers need their help. Mules are often unsuspecting victims themselves, who may believe they have been employed to perform a legitimate job. Yet a mule's real job is to transfer stolen money and to be the person who gets caught if law enforcement catches on.

By constantly improving phishing detection software and educating users about new types of phishing attacks as they are discovered, the number of phishing victims can be reduced. Coordinating international law-enforcement efforts and finding ways to make phishing less lucrative will also help. Still, phishing remains an arms race that will be hard to eliminate completely without stopping it at the source, so consumers need every form of protection they can get. ∎

© 2008 SCIENTIFIC AMERICAN, INC.