

DOI:10.1145/2063176.2063197

**Looking past the systems people use,
they target the people using the systems.**

BY JASON HONG

The State of Phishing Attacks

PHISHING IS A kind of social-engineering attack in which criminals use spoofed email messages to trick people into sharing sensitive information or installing malware on their computers. Victims perceive these messages as being associated with a trusted brand, while in reality they are only the work of con artists. Rather than directly target the systems people use, phishing attacks target the people using the systems. Phishing cleverly circumvents the vast majority of

an organization's or individual's security measures. It doesn't matter how many firewalls, encryption software, certificates, or two-factor authentication mechanisms an organization has if the person behind the keyboard falls for a phish.

On the surface, phishing attacks may seem to be a variant of spam. However, such attacks can lead to damaging losses in terms of identity theft,^{14,25} sensitive intellectual property and customer information, and national-security secrets.

Phishing attacks are also increasingly pervasive and sophisticated. Phishing has spread beyond email to include VOIP, SMS, instant messaging, social networking sites, and even massively multiplayer games.^{4,6,35} Criminals have also shifted from sending mass-email messages, hoping to trick anyone, to more selective "spear-phishing" attacks that use relevant contextual information to trick specific victims.

Academic research and commercial work in phishing is a dynamic area combining social psychology, economics, distributed systems, machine learning, human-computer interaction, and public policy. In 2006, Jakobsson and Myers²⁰ published an overview of how phishing works and what countermeasures were available at the time. This article serves as an introduction, as well as overview, of the

» key insights

- **Phishing attacks initially targeted general consumers, aiming to steal identity and credit-card information, but evolved to also include high-profile targets, aiming to steal intellectual property, corporate secrets, and sensitive information concerning national security.**
- **Developers must go beyond blaming users if they expect to deploy effective countermeasures against phishing attacks.**
- **The three general strategies for protecting end users from phishing scams: make things invisible, develop better user interfaces, and provide effective training.**



current state of phishing. It starts with how phishing attacks work, why people fall for them, the debate over the actual damage they cause, and finally a survey of countermeasures against phishing.

Anatomy of an Attack

Phishing attacks involve three major phases: The first is potential victims receiving a phish; the second is the victim taking the suggested action in the message, usually to go to a fake Web site but can also include installing malware or replying with sensitive information; and the third is the criminal monetizing stolen information.

Fake phishing email. Most phishing email messages use social techniques rather than technical tricks to fool end users. Conveying urgency is a well-known method used by criminals to misdirect people's attention;³⁴ an example is pretending to be a system administrator warning people about a new attack, urging them to install the attached patch. Another is notifying people there have been multiple failed logins for their account and they must verify their account now or risk dire consequences.

Appealing to people's sense of greed is an ancient technique now adapted to

the digital world. One phish the author of this article almost fell for about five years ago was filling out a survey for a bank in return for a small amount of money. The survey seemed innocuous until it asked for a bank-account number for depositing funds. So-called Nigerian 419 scams, offering "free" money in exchange for helping the sender move large amounts of money, also fall into this category. However, such obvious get-rich-quick scams are morphing to appeal to other emotions. Phishers today might pose as a relief agency asking for help with a recent natural disaster or as a random person appealing to

prurient interests, as in, say, “see Britney Spears naked.”


More sophisticated spear-phishing attacks use specific knowledge of individuals and their organizations; for example, an attack on military personnel might contain an invitation to a general’s retirement party, asking recipients to click on a link to confirm they will attend. People who wouldn’t normally fall for phish might in this case, due to the context. Jagatic et al.¹⁹ experimented in 2007 with how to exploit social-network information, showing that people were 4.5 times more likely to fall for phish sent from an existing contact over standard phishing attacks. Criminals indeed heavily target online social-networking sites partly for this reason.

Spear-phishing is also being used against high-level targets, in a type of attack called “whaling”; for example, in 2008, several CEOs in the U.S. were sent a fake subpoena along with an attachment that would install malware when viewed.²⁶ A *Communications* blog entry¹⁶ outlined several successful spear-phishing attacks in late 2010 and early 2011, with victims including the Australian Prime Minister’s office, the Canadian government, the Epsilon mailing list service, HBGary Federal, Oak Ridge National Laboratory, and RSA SecurID.


Setting up fake Web sites. Most phishing attacks try to convince people to go to a fake site where personal information is collected. To host a fake site, scammers use free Web space and a compromised machine or register a new domain.²⁷

When registering new domains, criminals look for names similar to the site they want to impersonate; for example, impersonating eBay, scammers might register ebay-login.com. Criminals also commonly use homograph attacks that exploit the visual similarity of characters; for example, bankofthevest.com⁸ uses two v’s to look like a w. Internationalized domain names facilitate this kind of attack, since characters in different language sets may appear identical.

However, in practice, criminals have opted for even simpler approaches. One is to put the domain name in plain sight, as in, say, paypal.com.phishsite.com. Surprisingly, many attacks make no attempt to disguise the destination site,



It doesn’t matter how many firewalls, encryption software, certificates, or two-factor authentication mechanisms an organization has if the person behind the keyboard falls for a phish.



relying on people’s lack of understanding of URLs. Unfortunately, even these simple tactics still fool many people.

When phishing attacks were just starting, scammers would create Web pages by hand, so they tended to be of poor quality, often including misspellings and hotlinks to images on the original site. The majority of phishing sites today are created with toolkits that might let a phisher specify what legitimate page to copy and where to direct stolen data, then generate all needed content. In 2008, Cova et al.⁷ identified more than 500 working kits. One surprising finding was that over one-third of these toolkits would send phished information to a location different from the one specified by the phisher, targeting inexperienced criminals who would do the work (and bear the legal risk) of breaking into sites.

When phishing attacks began, law enforcement, industry, and academic researchers were not organized in preventing and responding to attacks. However, as countermeasures (such as blacklisting and takedowns) were deployed (discussed later in the section on invisible attacks), criminals began introducing new techniques, thus starting an arms race that continues today. The most innovative approach so far is called “fast flux,” using a large pool of proxies and domain names to hide the true location of a phish. Fast flux makes it more difficult to blacklist sites since many URLs must be checked manually. Finding and taking down offending sites is also difficult since more work is needed to find the actual server. While an average phishing site lasts an average of 62 hours before being taken down, sites using fast flux tend to last an average of 196 hours.²⁷

Monetizing stolen information. The final phase of phishing is the monetization of stolen information. In some cases, the path is direct (such as when stealing banking credentials). In other cases, the path is convoluted (such as when stealing credentials for online games and social networking sites). Criminals have shown ingenuity here; for online games, they might transfer all of a victim’s virtual gold to an accomplice, then sell it to other players for real money. These attacks are common enough that Blizzard Entertain-

ment, creator of the popular online game *World of Warcraft*, sells special authenticators and offers in-game gifts for using them.⁵

Phishing on social networks is also somewhat indirect in terms of monetization. One attack involves notifying the victim's friends that the person is in trouble and needs money fast. Another involves using compromised accounts to spread malware; for example, the Koobface worm sends messages to a victim's friends urging them to go to a site that contains malware. Another is to steal victims' passwords and break into their email and bank accounts, working all too well since many people reuse passwords and existing password-reset mechanisms send responses directly to an individual's email address.

Marketplaces have also developed for acquiring and trading legitimate credentials. In the early days of phishing, phishers might use stolen credentials directly; today, many sell such credentials through underground networks to other criminals. These purchasers in turn might recruit unsuspecting people as "mules" to launder money and goods, reduce the risk the criminals face, and circumvent existing countermeasures; for example, some so-called "work at home" jobs involve receiving money transfers into the mule's bank account, with the funds actually coming from a hacked bank account. The mule then wires the money to a different account in another country, keeping a small commission. Such activities are illegal, and many perpetrators have been indicted around the world.²²

This evolution in how stolen credentials are monetized is due to specialization and perceived risk. A person good at creating phishing sites might not necessarily be good at stealing money from the accounts, especially given increasing vigilance by banks and law enforcement. Thus, rather than risk being traced, a phisher could opt to sell stolen information to others who are less risk averse.

Many researchers have examined how criminals trade stolen information on open Internet Relay Chat (IRC) channels. Herley and Florencio¹⁵ found that criminals often sell credentials for pennies on the dollar, explaining the

situation as a classic case of a marketplace for lemons. Given the anonymity of IRC, sellers find it easy to swindle purchasers by offering fake credentials or selling the same ones multiple times. Likewise, it is also easy for law enforcement and banks to offer honeypot credentials; as such, it is difficult for buyers to assess the quality of stolen data before buying. This asymmetric information about sellers and their goods leads buyers to dramatically lower what they are willing to pay.

Why We Fall for Attacks

An unfortunate response by the technically savvy is to dismiss end users as stupid and gullible, but it overlooks the fact that phishers deliberately exploit the poor usability of many interfaces that provide few cues for assessing the legitimacy of email messages and Web sites. Moreover, a deeper understanding of end-user motivations, beliefs, and mental models is essential for the security community to build effective countermeasures.

Dhamija et al.⁸ conducted one of the earliest studies (2006) investigating why people fall for phishing scams, asking participants to identify various Web sites as legitimate or fake. They found that good phishing sites fooled 90% of their participants and that most browser cues were opaque. Many par-

ticipants incorrectly judged sites based on their content and how professional they appeared, not realizing that Web pages are easily copied. Dhamija et al. also found that even experienced participants in the study had trouble with picture-in-picture attacks showing screenshots of a Web browser at a given site (see Figure 1). Picture-in-picture attacks point to an even greater challenge—that many people cannot differentiate between the browser "chrome," or buttons and URL area, that can mostly be trusted, and the browser content area, where attackers can show whatever they want.

Also in 2006, Downs et al.⁹ conducted a complementary study examining phishing email messages. As in Dhamija et al.,⁸ Downs et al. found their participants used basic, often incorrect heuristics in deciding how to respond to email messages; for example, some participants reasoned that since the business already had their information, it would be safe to give it again.

Sheng et al.³⁰ conducted a follow-up study involving a large-scale survey examining demographics and phishing susceptibility. Surprisingly, they found women were more vulnerable to phishing than men, primarily due to women having less exposure to technical knowledge. They also found that younger participants (ages 18 to 25) per-

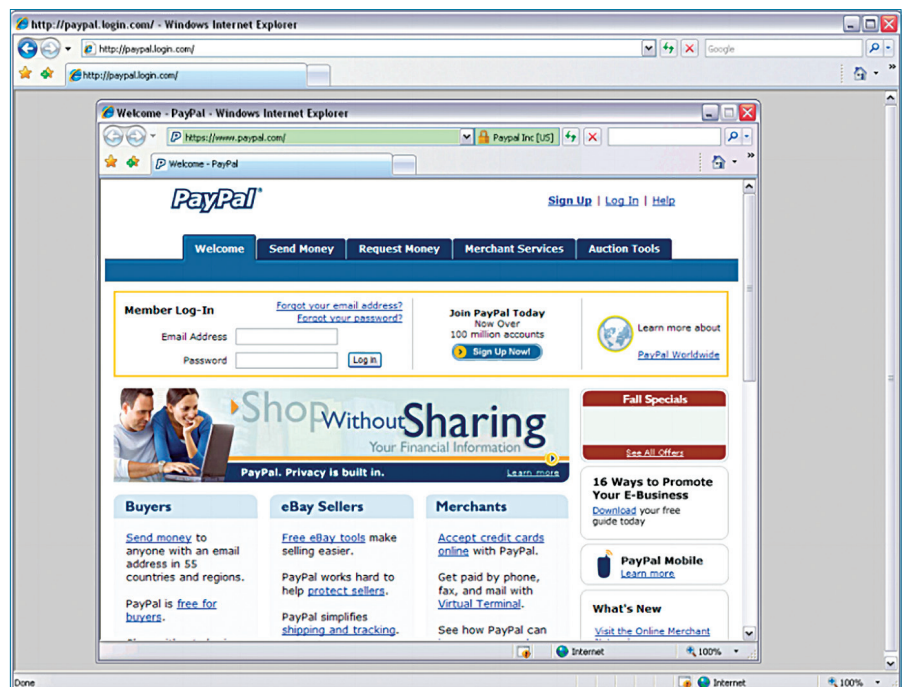


Figure 1. Picture-in-picture phishing attacks show an image of a Web browser in the content area of a Web browser and fool even experienced users; adapted from Jackson et al.¹⁸

formed worse than all other age groups, possibly due to fewer years of experience online, less exposure to training materials, and less aversion to risk.

How Bad?

The Anti-Phishing Working Group (<http://www.antiphishing.org>) is an international consortium of law enforcement, industry, and academic researchers devoted to combating Internet scams and online fraud. For phishing attacks against consumers, the peak was in the third quarter of 2010 when APWG identified more than 115,000 unique phishing email messages worldwide, along with more than 150,000 unique phishing sites worldwide.³ Subsequent APWG reports have shown a slow downward trend.

However, estimates of damage caused by phishing varies widely, ranging from \$61 million per year¹⁴ to \$3 billion per year²⁵ of direct losses to victims in the U.S. The main problem is a lack of data from banks and other institutions that suffer losses; as such, these estimates are heavily dependent on the methods used and assumptions made by the organizations compiling and reporting the statistics.

While there is not yet full agreement among security analysts regarding how to calculate direct damages, there is increasing agreement that the indirect costs of phishing are substantial. One bank the author has spoken to said it cost it about \$1 million per attack in terms of call-center costs, recovery costs, and actual money (relatively small) that could not be recovered. A more difficult metric to measure is the damage to personal and corporate reputations. In presentations on the economics of computer security, Cormac Herley of Microsoft Research captured the problem succinctly, asking: "What is the first thing you think of when you hear the words 'Nigerian businessman'?"

Estimates of direct costs to the public also fail to capture the damage from specialized spear-phishing attacks. A number of successful high-profile attacks were reported in 2011, victimizing agencies in the Australian and Canadian governments, the Epsilon mailing list service, Gmail, Lockheed-Martin, Oak Ridge National Laboratory, and RSA.¹⁶

In 2009, the Operation Aurora attacks used spear-phishing and malware to target a number of organizations, most notably Adobe, Google, Symantec, and Yahoo!. In many cases, the attackers stole source code and other intellectual property. However, there are no good estimates as to the damage caused by spear-phishing, due to victims' unwillingness to share information and the basic difficulty in assessing damages.

Countermeasures

Given the risks of phishing, what can individuals and organizations do to protect themselves? Taking an end-user perspective, three general strategies are: make things invisible, so users do not have to do anything different; provide better user interfaces that either make the situation more obvious to users or offer additional protection; and train end users to proactively recognize and avoid phishing attacks. All three are needed to provide the strongest possible protection against phishing attacks.

Make it invisible. The first line of defense is to prevent phishing attacks from reaching end users in the first place. Solutions include filtering phishing email messages, blocking fake sites, and taking down fake sites.

Filtering phishing email. A large body of research covers spam detection. However, research on detecting phishing email messages is sparse, in part because phishing is a relatively new phenomenon and because phishing email messages sent by technically savvy criminals are so convincing. Fette et al.¹¹ developed the first email phishing filter (in 2007) identifying several features that are highly indicative of phishing, including, for example, URLs that use different domain names. Researchers and developers alike have since then explored additional features and machine-learning techniques.

An alternative to heuristics is authentication and verification technologies. For example, the Sender Policy Framework (SPF) uses the Simple Mail Transfer Protocol (SMTP) to reject forged email addresses. Another is DomainKeys Identified Mail (DKIM), verifying the DNS domain of a sender and message integrity. However, all have proved difficult to deploy on a large scale and do not provide protection

against several kinds of phishing attacks;¹³ for example, while focusing on preventing email spoofing, attackers easily create alternative fake addresses.

Blocking phishing sites. There are two general ways of detecting phishing Web sites: heuristics that examine the URL, HTML, and server characteristics to classify sites; and manually verified blacklists.

For heuristics, researchers have investigated a number of ideas using machine learning; examples include looking for patterns in URLs,¹² words in Web pages,¹ and search engines.³⁹ Researchers have also looked at linguistic characteristics of Web pages, identifying the brand name a Web page claims to be.³⁷ The effectiveness of all such techniques is reasonable, with true positive rates (correctly identifying a phishing site) of 90% or better and false positive rates (incorrectly labeling a legitimate site as phish) approaching 1% or less.

The best-known anti-phishing blacklists are operated by Google, Microsoft, and PhishTank, each containing URLs manually verified as phish. Google's is integrated with Firefox and Chrome, so no special action is required of end users to protect themselves. Microsoft's is integrated with Internet Explorer. And PhishTank's uses a wisdom-of-crowds approach to identify phish, letting people submit potential phish. Once enough other people vote that a submission is indeed a phish, it is added to their blacklist. Since October 2006, PhishTank has received close to four million votes from volunteers, labeling more than half-a-million phishing sites.²⁸

Several commercial browser add-ons are designed to block phish. Since they can be installed in Web browsers, their effectiveness can be evaluated empirically. For example, in 2009, Sheng et al.³³ examined major blacklists and browser tools, showing that zero-hour protection offered by blacklists had a false positive rate of 0% but a true positive of less than 20%. Even after 12 hours, the best blacklist identified only 83% of phish. They also found that deployed heuristics were somewhat effective in identifying phish but were used only to warn people in the Web browser rather than block likely phishing sites.

Sheng et al.'s research identified a gap between research and industry in terms of true positives. Academic research has generally focused on heuristics and machine-learning techniques with very good true positives though somewhat high false positives. These heuristics are good at identifying phishing sites not seen before. On the other hand, industry relies primarily on blacklists, which have middling true positives but no false positives. However, blacklists do not generalize well to future unseen cases, can be slow to respond to zero-hour attacks, and are easily overwhelmed by automatically generated URLs, a tactic phishers have already adopted.

In follow-up work, Sheng et al.³¹ probed the issue of heuristics vs. blacklists by interviewing people in industry, law enforcement, and academia, finding that concern over liability for false positives is the major barrier to deploying more aggressive heuristics. However, the first few hours of an attack are critical for blocking it, as a substantial fraction of users will have read their email by the time blacklists are updated. Jagatic et al.¹⁹ found that during regular work hours, most users who fell for a phishing attack did so in the eight hours following the start of the attack.

Sheng et al.³¹ identified several ways to ameliorate the situation; foremost is to clarify the legal issues surrounding false positives. Another is to have a central clearinghouse for phish, rather than piecemeal efforts that take longer to identify phish due to duplicated effort. A third is for researchers to develop better heuristics that minimize false positives. An early example of such heuristics was developed by Xiang et al.³⁸ observing that many phish are near or exact duplicates because they are generated by toolkits. Once a phish is on a blacklist, other copies of it can be identified quickly and blocked with virtually no risk of false positives. Using probabilistic-matching methods, the obvious countermeasure of adding noise can also be mitigated.

Taking down phishing sites. Several organizations identify and take down phishing sites, and private mailing lists help share information about fake sites, as well as find contact information for specific ISPs and Web sites.

When phishing sites are taken down, end users who click on a phish are typically shown a "page not found" error. One innovation developed by APWG and Carnegie Mellon University is to have ISPs and take-down providers replace the phishing page with a training message, teaching people who click on phishing email messages about such attacks. The APWG landing page,² in use since September 2008, is available in several languages. As of April 2010, it has been displayed in place of 1,285 phishing pages and viewed almost 200,000 times.¹⁷ While measuring the effect of the landing page is difficult, it is a step in the right direction, offering multiple ways of protecting people worldwide.

Better interfaces. The second major strategy for protecting people is to provide better interfaces. The following paragraphs cover innovations in warnings, support for properly identifying Web sites, and authentication.

A general problem with security warnings is that users often close them the instant they appear, a perfectly rational behavior, as many warnings are so obtuse people don't understand what the problem is or what they should do. Other warnings annoyingly

interrupt what people are trying to accomplish. Warning notifications can also be too subtle, with people not even seeing them.

A "passive indicator" warns of potential dangers without interrupting the user's task. In contrast, "active indicators" force users to notice the warnings by interrupting them. Studies by Wu et al.³⁶ and Egelman et al.¹⁰ found passive warnings ineffective in protecting people from phishing scams, as they are easily missed.

Egelman et al. also examined the effectiveness of active anti-phishing warnings in Firefox and Internet Explorer 7; see Figure 2 for an example of Firefox's active warning. Using simulated phishing attacks, they found no participants fell for phishing attacks when seeing Firefox's warning, but, surprisingly, half of the participants using IE did. Egelman et al. analyzed the results using a model from the warning sciences describing how people see, understand, believe, and act on warnings in the regular physical world. Using this framework, they found that most people simply did not "see" the warning in IE, since it looked like a standard "page not found" warning. Several participants also did not believe the warn-

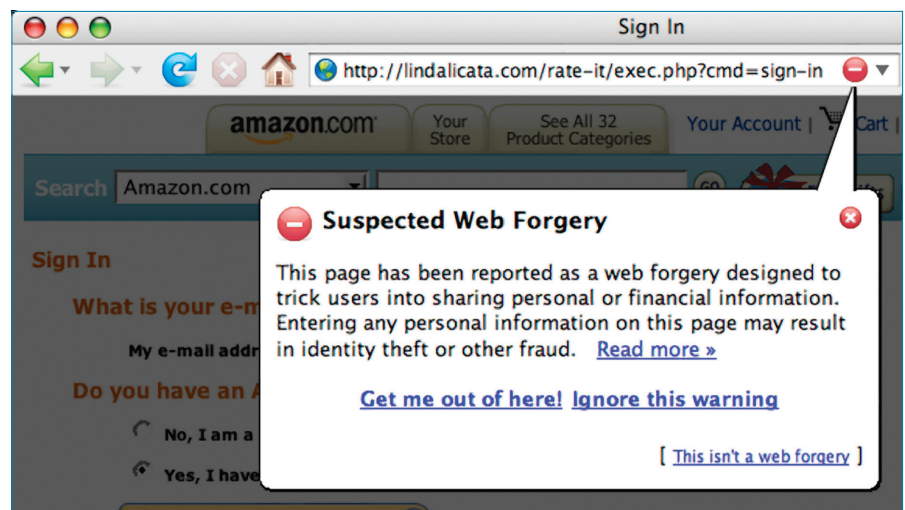


Figure 2. The active warnings used by Mozilla Firefox when blocking phishing pages are more effective than passive warnings.




Figure 3. Extended Validation certificates in Microsoft's Internet Explorer and Mozilla Firefox.

ing, thinking Microsoft would not put them at risk, and went on to give sensitive personal information. In response to this work, Microsoft redesigned its anti-phishing warnings in IE8.

Techniques have been developed to help people identify the sites they are on. However, it is unclear how much they help in practice; for example, Extended Validation (EV) certificates are a special kind of certificate with specific guidelines for verifying that an organization purchasing the certificate is legitimate. When a site with an EV certificate is loaded, the browser's URL bar changes automatically to show the site's brand name (see Figure 3). However, a 2007 study by Jackson et al.¹⁸ found EV certificates ineffective in protecting people from phishing attacks.

Using SiteKey, a technique used by many financial organizations worldwide, users first select a secret picture. When logging in, they see if the picture is displayed to verify they are on the right site. However, Schecter et al.²⁹ found SiteKey suffers from the same problem as passive indicators, in that the absence (or presence) of an indicator is easily missed or even rationalized away by the typical user.

An alternative to indicators is to improve the way users sign into sites. Two-factor authentication (2FA) strength-

ens authentication by requiring two separate ways to prove one's identity. One of the most common forms is key fobs with a periodically changing number synchronized with a remote server. Users login by using both their password and this number. While 2FA increases the cost of conducting phishing attacks, phishers have also developed workarounds, including switching to real-time man-in-the-middle attacks using malware (such as the Zeus Trojan horse).

Train users. Training is the third way of protecting people from phishing scams. Though an essential aspect of computer security, it is also arguably the least-popular approach, given the inherent challenges of motivating people to be secure, as well as the fact that training does not guarantee complete protection (though neither does any other solution today).

Many Web sites offer advice on how to identify phishing sites. A 2010 study by Kumaraguru et al.²⁴ found this kind of information useful in helping people identify fake Web sites but only if they actually read the material. In an earlier study, Kumaraguru et al.²³ found that simply emailing anti-phishing material was ineffective, because people were habituated to receiving such warnings and thought they already knew how to protect themselves.

Two lines of research have sought to address these problems: The first is micro games designed to teach people about phish. Micro games are a popular format for games played for short periods of time. Sheng et al.³² developed a micro game for computer security called Anti-Phishing Phil (see Figure 4) that teaches about browser address bars, domain names, and phishing pages, then tests users on what they learned. Phil incorporates many ideas from learning science, a body of empirical research that seeks to understand the best methods for learning and retaining knowledge. An example principle is "conceptual-procedural," holding that high-level concepts should be interleaved with concrete procedures on how to achieve given tasks. An evaluation of Phil with more than 4,500 people demonstrated it improved novices' ability to identify phish by 61% while also dramatically lowering false positives.

The second approach is "embedded training," teaching people in the specific context of use in which they would normally be attacked. Embedded training is in contrast to other forms of security training that might take place in a classroom and give people few opportunities to test what they've learned. Kumaraguru et al.²⁴ developed an embedded training system called PhishGuru that sends simulated phishing email messages to people. If they fall for one, they see an intervention that teaches them about phishing and how to protect themselves. In a study with more than 500 participants, Kumaraguru et al. found this approach led to a 45% reduction in falling for phish even a month after being trained, helping lead to creation of the APWG landing page² described earlier.

Conclusion

This article has emphasized criminals' tenacity and creativity, a trend that will only continue. We will also likely see increased spear-phishing and whaling attacks, as phishers look for vulnerable targets with valuable information.

Blurring traditional security perimeters, phishing also causes new problems for organizations. Even their lawyers and accountants can be attacked to surreptitiously gain access to

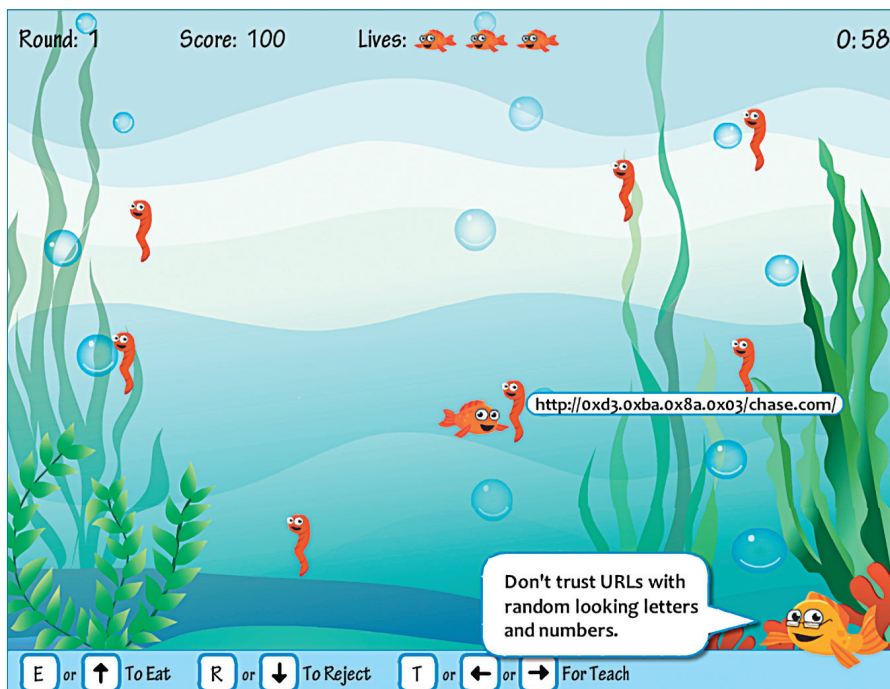


Figure 4. Anti-Phishing Phil is a micro game that teaches people how to identify phishing scams.

documents. Facebook and other social media provide more contextual details that can be used for spear-phishing attacks. Individual employees falling for phish in one context can cause headaches for their organizations over reused passwords. Finally, instant messaging, VOIP, SMS, and other relatively new ways of communicating offer criminals new vector options for delivering attacks.

On the positive side, law enforcement, industry, and academic researchers are getting better organized in terms of reporting phishing attacks, sharing information, analyzing data to identify trends, and focusing resources. More organizations are devoted to combating online fraud, including APWG, the Internet Crime Complaint Center (IC3) (<http://www.ic3.gov/>), and National Cyber-Forensics and Training Alliance (NCFITA) (<http://www.ncfta.net/>). There are also initiatives dedicated to educating people about phishing scams, including StaySafeOnline.com. Law enforcement has stepped up efforts in gathering evidence and cooperating with international partners in shutting phishing sites and phishing gangs. Legislators in the U.S. have also passed laws to explicitly spell out what phishing is and the related penalties, including California's Anti-Phishing Act of 2005,²¹ though such laws face many of the same challenges as anti-spam laws in terms of attackers being outside a particular jurisdiction, the sheer number of attacks, and limited resources available to law enforcement.

Phishing will continue to be an arms race. Since any communication medium can be used for phishing, it is also a problem that can never be solved. The best we can hope for is to blunt the worst aspects of phishing and continue to work on better ways to prevent, detect, and respond to this new form of a very old crime.

Acknowledgments

Special thanks to the Supporting Trust Decisions group (<http://cups.cs.cmu.edu/trust/>) at Carnegie Mellon University and Wombat Security Technologies (<http://www.wombatsecurity.com/>) for their contributions and comments. C

References

1. Abu-Nimeh, S., Nappa, D., Wang, X., and Nair, S. A comparison of machine learning techniques for phishing detection. In *Proceedings of The Anti-Phishing Working Group's Second Annual eCrime Researchers Summit* (Pittsburgh, PA, Oct. 4–5, 2007), 60–69.
2. Anti-Phishing Working Group. APWG & Carnegie Mellon University's phishing education landing page; <http://education.apwg.org/r/en/>
3. Anti-Phishing Working Group. *Phishing Activity Trends Report: Third Quarter Report*, Jan. 2010; http://apwg.org/reports/apwg_report_Q3_2009.pdf
4. Arthur, C. Facebook hit by phishing attack. *The Guardian* (Apr. 30, 2009); <http://www.guardian.co.uk/technology/2009/apr/30/facebook-phishing-scam>
5. Blizzard Entertainment. *Battle.net Authenticator FAQ*; http://us.blizzard.com/support/article.xml?locale=en_US&articleId=24660
6. Cavalli, E. World of Warcraft phishing attempts on the rise. *Wired* (Apr. 29, 2009); <http://www.wired.com/gamelife/2009/04/world-of-warcraft-phishing-attempts-on-the-rise/>
7. Cova, M., Kruegel, C., and Vigna, G. There is no free phish: An analysis of 'free' and live phishing kits. In *Proceedings of the Second USENIX Workshop on Offensive Technologies* (San Jose, CA, July 28, 2008). Usenix; <http://portal.acm.org/citation.cfm?id=1496706>
8. Dhamija, R., Tygar, J.D., and Hearst, M.A. Why phishing works. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Quebec, Apr. 24–27). ACM Press, New York, 2006, 581–590; <http://portal.acm.org/citation.cfm?id=1124861>
9. Downs, J.S., Holbrook, M.B., and Cranor, L.F. Decision strategies and susceptibility to phishing. In *Proceedings of the SOUPS Symposium on Usable Privacy and Security* (Pittsburgh, July 12–14). ACM Press, New York, 2006.
10. Egelman, S., Cranor, L.F., and Hong, J.I. You've been warned: An empirical study of the effectiveness of Web browser phishing warnings. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Florence, Italy, Apr. 5–10). ACM Press, New York, 2008, 1065–1074.
11. Fette, I., Sadeh, N., and Tomicic, A. Learning to detect phishing emails. In *Proceedings of the 16th International World Wide Web Conference* (Banff, Canada, May 8–12, 2007), 649–656.
12. Garera, S., Provos, N., Chew, M., and Rubin, A.D. A framework for detection and measurement of phishing attacks. In *Proceedings of the WORM Workshop on Rapid Malcode* (Alexandria, VA, Nov. 2). ACM Press, New York, 2007; <http://portal.acm.org/citation.cfm?id=1314391>
13. Görling, S. An overview of the Sender Policy Framework as an anti-phishing mechanism. *Internet Research* 17, 2 (2007), 169–179.
14. Herley, C. and Florencio, D. A Profitless endeavor: Phishing as a tragedy of the commons. In *Proceedings of the New Security Paradigms Workshop* (Lake Tahoe, CA, Sept. 22–25, 2008).
15. Herley, C. and Florencio, D. Nobody sells gold for the price of silver: Dishonesty, uncertainty, and the underground economy. In *Proceedings of Workshop on the Economics of Information Security* (London, June 24–25, 2009).
16. Hong, J. Why have there been so many security breaches recently? *Blog@CACM* (Apr. 27, 2011); <http://cacm.acm.org/blogs/blog-cacm/107800-why-have-there-been-so-many-security-breaches-recently/fulltext>
17. Hong, J.I. Statistical analysis of phished email users intercepted by the APWG/CMU phishing education landing page. In *Proceedings of the Anti-Phishing Working Group Counter eCrime Operations Summit IV* (Sao Paulo, Brazil, May 11–13, 2010); http://www.antiphishing.org/events/2010_opSummit.html
18. Jackson, C., Simon, D.R., Tan, D.S., and Barth, A. An evaluation of extended validation and picture-in-picture phishing attacks. In *Proceedings of the 11th International Conference on Financial Cryptography* (Trinidad/Tobago, Feb. 12–15, 2007), 281–293.
19. Jagatic, T.N., Johnson, N.A., Jakobsson, M., and Menczer, F. Social phishing. *Commun. ACM* 50, 10 (Oct. 2007), 94–100.
20. Jakobsson, M. and Myers, S. *Phishing and Countermeasures: Understanding the Increasing Problem of Electronic Identity Theft*. Wiley-Interscience, 2006.
21. Keizer, G. California enacts tough anti-phishing law. *InformationWeek* (Oct. 3, 2005); <http://informationweek.com/news/171202672>
22. Krastev, N. U.S. indicts dozens from Eastern Europe in Internet theft scheme. *Radio Free Europe* (Oct. 1, 2010); http://www.rferl.org/content/US_Indicts_Dozens_From_Eastern_Europe_In_Internet_Theft_Scheme/2173545.html
23. Kumaraguru, P., Rhee, Y., Sheng, S. et al. Getting users to pay attention to anti-phishing education: Evaluation of retention and transfer. In *Proceedings of the Anti-Phishing Working Group's Second Annual eCrime Researchers Summit* (Pittsburgh, Oct. 3–5, 2007), 70–81.
24. Kumaraguru, P., Sheng, S., Acquisti, A., Cranor, L.F., and Hong, J.I. Teaching Johnny not to fall for phish. *ACM Transactions on Internet Technology* 10, 2 (2010), 1–31.
25. Litan, A. Phishing attack victims likely targets for identity theft. Gartner Group, May 2004.
26. Markoff, J. Larger prey are targets of phishing. *New York Times* (Apr. 16, 2008); <http://www.nytimes.com/2008/04/16/technology/16whale.html>
27. Moore, T. and Clayton, R. Examining the impact of Website take-down on phishing. In *Proceedings of the Anti-Phishing Working Group's Second Annual eCrime Researchers Summit* (Pittsburgh, Oct. 3–5, 2007), 1–13.
28. PhishTank. *PhishTank Stats, 2011*; <http://www.phishtank.com/stats.php>
29. Schechter, S.E., Dhamija, R., Ozment, A., and Fischer, I. The emperor's new security indicators: An evaluation of Website authentication and the effect of role playing on usability studies. In *Proceedings of the IEEE Symposium on Security and Privacy* (Washington, D.C., 2007), 51–65.
30. Sheng, S., Holbrook, M.B., Kumaraguru, P., Cranor, L.F., and Downs, J.S. Who falls for phish? A demographic analysis of phishing susceptibility and effectiveness of interventions. In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Atlanta, Apr. 10–15). ACM Press, New York, 2010, 373–382.
31. Sheng, S., Kumaraguru, P., Acquisti, A., Cranor, L.F., and Hong, J.I. Improving phishing countermeasures: An analysis of expert interviews. In *Proceedings of the Fourth Anti-Phishing Working Group eCrime Researchers Summit* (Tacoma, WA, Oct. 20–21, 2009).
32. Sheng, S., Magnien, B., Kumaraguru, P. et al. Anti-Phishing Phil: The design and evaluation of a game that teaches people not to fall for phish. In *Proceedings of the Third Symposium on Usable Privacy and Security* (Pittsburgh, July 18–20). ACM Press, New York, 2007, 88–99.
33. Sheng, S., Wardman, B., Warner, G., Cranor, L.F., Hong, J.I., and Zhang, C. An empirical analysis of phishing blacklists. In *Proceedings of the Sixth Conference on Email and Anti-Spam* (Mountain View, CA, July 16–17, 2009).
34. Stajano, F. and Wilson, P. Understanding scam victims: Seven principles for systems security. *Commun. ACM* 54, 3 (Mar. 2011), 70–75.
35. Verisign. *Fraud Alert: Phishing: The Latest Tactics and Potential Business Impact*. White Paper, 2009; <http://www.verisign.com/static/phishing-tactics.pdf>
36. Wu, M., Miller, R.C., and Garfinkel, S. Do security toolbars actually prevent phishing attacks? In *Proceedings of the CHI Conference on Human Factors in Computing Systems* (Montréal, Apr. 24–27). ACM Press, New York, 2006, 601–610.
37. Xiang, G. and Hong, J.I. A hybrid phish detection approach by identity discovery and keywords retrieval. In *Proceedings of the International World Wide Web Conference* (Madrid, Apr. 20–24, 2009), 571–580.
38. Xiang, G., Rose, C., Hong, J.I., and Pendleton, B. A hierarchical adaptive probabilistic approach for zero-hour phish detection. In *Proceedings of the ESORICS 15th European Symposium on Research in Computer Security* (Athens, 2010), 571–589.
39. Zhang, Y., Hong, J.I., and Cranor, L.F. Cantina: A content-based approach to detecting phishing Web sites. In *Proceedings of the 16th International World Wide Web Conference* (Banff, Canada, May 8–12, 2007), 639–648.

Jason I. Hong (jasonh@cs.cmu.edu) is an associate professor in the School of Computer Science and Human Computer Interaction Institute at Carnegie Mellon University, Pittsburgh, PA.

© 2012 ACM 0001-0782/12/01 \$10.00