# PHISH PHODDER: IS USER EDUCATION HELPING OR HINDERING?

*David Harley*
Small Blue-Green World, 8 Clay Hill House
Haslemere, Surrey GU27 1DA, UK

Tel +44 1428 749186
Email dharley@smallblue-greenworld.co.uk

*Andrew Lee*
ESET, Broad Quay House, Prince Street, Bristol
BS1 4DJ, UK

Email aj@eset.com

## ABSTRACT

Mostly, security professionals can spot a phish a mile off. If they do err, it's usually on the side of caution, for instance when real organizations fail to observe best practice and generate phish-like marketing messages. Many sites are now addressing the problem with phishing quizzes, intended to teach the everyday user to distinguish phish from phowl (sorry). Academic papers on why people fall for phishing mails and sites are something of a growth industry. Yet phishing attacks continue to increase, and while accurate and up-to-date figures for financial loss are hard to come by, indications are that losses from phishing and other forms of identity theft continue to climb.

This paper:

1. Evaluates current research on how end users are susceptible to phishing attacks and ID theft.
2. Evaluates a range of web-based educational and informational resources in general and summarizes the pros and cons of the quiz approach in particular.
3. Reviews the shared responsibility of phished institutions and phishing mail targets for reducing the impact of phishing scams. What constitutes best practice for finance-related mail-outs and e-commerce transactions? How far can we rely on detection technology?

## INTRODUCTION

We should define what we *don't* mean by phishing for the purposes of this paper, though we will allude to other forms of scam where appropriate. But we won't be focusing on the following:

- 419 (advance fee fraud) scams [1], though such scams usually involve masquerading. However, any ID theft involved is generally incidental.
- Pump and dump scams, which are more impersonal in nature, and don't involve direct access to the mail recipient's funds or identity [2].
- Mule recruitment scams (419- *or* phishing-related). However, mule recruitment is another side of the Black Hat Economy [3], of which phishing is one of the essential components.
- Deceptive messages that aren't generally or primarily profit motivated (hoaxes, urban legends and so on [4]), though these sometimes involve some element of identity theft ('This virus warning was issued this morning by *AOL*/*Microsoft*/*McAfee*/Loamshire police/whoever else might convince you to mistake this out-and-out fiction for the truth...') [5].

We have, however, detailed the distinctions between these and other scams in another phishing-related paper [2].

So what *do* we mean by phishing? Certainly the practice of posting a deceptive message (often, but not necessarily via email) [6] as part of an attempt at fraud and/or identity theft, and especially one manipulated to make it look as if it comes from a legitimate business or agency, when in reality it is from a criminal source [2]. We assume an intent to acquire sensitive data by malicious social engineering. The underlying assumption is of intent to plunder the victim's financial resources, to steal their identity for criminal purposes, to obtain information about them for sale to others, or a combination thereof. However, the posting of a deceptive message is only part of the phishing process: equally important is the dishonest acquisition of data from fake websites or other data capture methods, including fake forms, keyloggers, backdoor trojans and so on. So it's still important that the user education process includes, somewhere, awareness of the dangers of message attachments and downloads (including drive-by downloads) from unverified sites, however convincing they may look.

It's often assumed that phishing is about finance-related institutions (banks, credit unions, *PayPal*, auction sites etc.) We don't assume, however, that target data is always related to the victim's personal finances. In principle this kind of attack can be intended to access quite different forms of data (industrial espionage, ISP account information, information relating to access to restricted systems, and so on.) Non-commercial entities may also need to allow clients to volunteer financial information to pay for services electronically. In consequence, potential victims are conditioned to share sensitive data with groups masquerading as taxation departments, healthcare and social security agencies, law enforcement agencies and so on – even retail outfits in some circumstances, where ordering via telephone or web [7]. Quizzes assuming financial data could be considered over-specific: after all, we consider it more useful to teach generic scepticism than the recognition of highly specific scams.

Phishing activity is not necessarily restricted to short-term exploitation of financial data, but may be extended to full-scale identity theft. Unsurprisingly, quiz sites that cover this range of scams earn extra brownie points, in our view, though in this instance points don't mean prizes.

## WHY AND HOW DOES PHISHING WORK?

Mustaca [8] suggests that a number of factors are at work:

- The verisimilitude of the phish email and/or of the faked site to which it links.
- How fast and how far the mail is distributed before the site is shut down.
- Speed and timing of activation of the fake site.
- Susceptibility of phish mails and (possibly) fake sites to automated detection.

In fact, the technical skills of the criminal (in terms of presenting convincing counterfeit emails and websites) are far from being the only relevant factor.

While the general level of phishing presentation has risen dramatically, the continuing success of stereotyped 419 scams suggests that poor presentation doesn't always mitigate gullibility [9]. Skilful social engineering (offering rewards for information, or scare tactics like 'your account has been compromised – to re-authenticate, click here, or we'll cancel your account') is at least as relevant, though perhaps the most effective weapon in the scammer's armoury is simply that victims are confused about the nature of the problem. Of course, education, including phish quizzes, should dispel some of that confusion.

### Phishing attack components

A phishing attack can be regarded as having three parts, as described by Mustaca [8] in a model subsequently adapted by Harley & Lee [2].

- Bait distribution through email, instant messaging, or, increasingly, other channels. Vishing, for example, uses VoIP (Voice over IP) technology to extend phish-like scams to telephone services.
- Data collection through misdirection, most commonly through a fake website. In principle, data may also be collected through a direct response to email, or an intermediate form of misauthenticated response, or through the planting of spyware.
- The use of the misappropriated information for purposes of fraud and identity theft.

These definitions are purely functional; a phishing crew may comprise a far wider range of roles (bot herder, mule driver, programmer and so on). These issues are further addressed in several papers [2, 3] and other resources.

### Bait distribution

Phish emails range from crude, badly spelled plain text to sophisticated, well-conceived, graphic-rich messages distinguishable only by their content and provenance. They share characteristics with other forms of scam (hashbusters, use of images, and so on) but there are techniques particularly associated with phish-type mails. Since most of them are intended to misdirect victims to a site masquerading as a legitimate web page, they often feature some kind of concealment or obfuscation of the real target URL, using URL encoding, misused <map> tags, and so on [2].

Most phish quizzes are focused on the recognition of phishing emails, but don't always explore these details, even though some of them are not hard to spot in a modern mail client.

Bait emails are not the only way of decoying a potential victim away from a legitimate website to a spoofed site or an interpolated page or script. Other possibilities include:

- Cousin domains
- Typosquatting
- Pharming/DNS spoofing
- Forms that pop up over a legitimate site
- Some form of cross-site scripting

However, these techniques are not usually detailed in phishing quizzes, and not particularly easy to represent using static graphics.

### Data collection

- Websites constructed to resemble the phished organization's site, often incorporating elements of the genuine site.
- Dynamic insertion of code into the legitimate site through compromise of the browser or machine.
- Pop-up or pop-down forms designed to appear when the real site was accessed via a link in the email.

Bait and data collection are not dependent on email/messaging or web pages and forms. The installation of some form of spyware for data collection can be achieved by a number of alternative approaches: for example, over unsafe network shares, the use of unpatched vulnerabilities, drive-by downloads, and so on. We won't consider the role of bots and botnets in the phishing problem in this paper, but some of these attacks are part of the bot controller's standard armoury [10]. Again, we don't see these issues referenced directly in phishing quizzes; nor do we see much consideration of the third aspect of the model (misuse of the data for criminal/fraudulent purposes) except in text-oriented multiple choice tests.

## LITERATURE REVIEW

We can't cover the whole range of current literature in this area and still do justice to our main theme, but here are a few interesting references. Others are included in [2].

'The emperor's new security indicators' [11] is an attempt to evaluate the usefulness of website authentication measures. The study concluded that:

- The absence of HTTPS indicators did not dissuade participants from entering their passwords when asked to carry out common online banking tasks.
- The absence of site authentication images did not dissuade them, either.
- The use of site authentication images can cause site users to ignore other security indicators.
- Participants who were role playing were more apt to disregard attack clues than those who were using their own passwords to access an account.

However, it's unclear how closely the behaviour displayed by participants in the study environment reflected their behaviour in real life online transactions. Whalen and Inkpen [12], in a study referenced in [11], also suggested that study participants were more careful with their own data than with 'made-up' data, and Schechter *et al.* did acknowledge that design aspects of their own study could have caused participants to behave less securely than normal.

'Why phishing works' [13] focuses on data collection, specifically via spoofed websites. This study suggested that a well-spoofed site could take in over 90% of the participants, even though they were aware of phishing, if not of the technological issues.

- Nearly a quarter of participants were influenced only by the content of the website in evaluating its authenticity;

they were more 'persuaded' by design (favicons, animated and static graphics etc.) than by SSL or certification indicators/non-indicators.

- These indicators of trustworthiness were poorly understood or totally unnoticed. The authors were able to fool even their most knowledgeable subjects, using simple spoofing techniques to counterfeit such indicators.
- Legitimate sites that enforced restricted access from SSL-protected pages were actually perceived as *less* trustworthy.

'Protecting people from phishing: the design and evaluation of an embedded training email system' [14] illustrates interesting alternative approaches to quizzes, using: (1) a simple text and graphics 'intervention' illustrating self-protection; (2) a similar intervention, but in cartoon form. In fact, as long ago as 2000, one of the authors implemented a company-wide modification of the email client to include a 'Send to virus alert team' button along with the usual 'Reply', 'Forward' etc. This was geared more toward cutting down hoax traffic (by filtering through the security team). However, because it was backed with a general program of education around security and malware, along with an extensive security-dedicated intranet, it was quite effective in that people often checked out phishes and scams by using that button.

'Best practices for businesses to avoid being phished' is a document being developed by the Anti-Phishing Working Group, the Mail Anti-Abuse Working Group, and the US Homeland Security Identity Theft Technology Consortium. Rather than relying purely on educating the user, it takes the approach of educating the kind of business that is liable to be phished in the kind of best practices that make it less easy for the phishing gang. We have a good deal of sympathy for that approach. The continued use of such poor practices as phish-like text, inadequate personalization, and unnecessary URL redirects into a very different domain, is referred to by James and others [3] as 'consumer miseducation', since it primes potential victims to accept bad practice as 'legitimate'. However, part of the task of educating the banks (etc.) is to persuade them to take on the responsibility of educating, in turn, their customers.

## DE-GULLING THE GULLIBLE – TEACHING SCEPTICISM

Phishing quizzes are an increasingly popular approach to end-user education. These are usually:

- Multiple-choice questionnaires aimed at raising consciousness about phishing issues.
- Email or website recognition tests where the participant assesses whether sample messages or (less often) sites are genuine. These are, however, of highly variable quality. Sometimes the testing site's own analysis of 'suspicious' attributes is inadequate or misleading.

Multiple choice questions are as good or bad (and as up-to-date) as to the knowledge of the compilers, and we don't consider them at length here. We would point out, though, that general questions along the lines of 'How many phishing mails are sent out every month?' have little mitigating impact on the participant's vulnerability to those mails.

The most common type of phishing quiz we've encountered is the type where the subject is shown a number of sample emails and invited to categorize them as either phishing mails or legitimate communications. Informal discussion with an arbitrary sample of other security professionals suggests that they generally:

- Pick up all the real phishes.
- Correctly assess some mails as legitimate.
- 'Fail' to recognize some legitimate mails as such. We believe that this often results because, lacking sufficient contextual information to assess their legitimacy, they err on the side of caution.

Anecdotal evidence suggests that even the general public score better on phish recognition than they do on legitimate, but phish-like mails. But is that *their* problem, or that of the institution that sends out phish-like emails? We have come to the following conclusions, based on fairly informal research into web-based phishing quizzes currently found on the web:

- Quizzes based on categorizing sample emails as phish or legitimate are based on or give rise to the assumption that the participant can make an accurate assessment, irrespective of the legitimacy of the mail, simply by viewing a screenshot. However, they often supply insufficient information to make an accurate decision. It's still common for quizzes not to indicate whether embedded URLs were exactly as shown, obfuscated, or otherwise deceptive – as when the apparent and real target link are quite different. Thus, the subject loses the advantage of an important visual cue for identifying some kinds of phish.
- The use of static screenshots of sample messages deprives the subject of other visual cues such as access to HTML source code or knowledge of whether the message has been sent to a 'legitimate' email address – often phishes are so convincing, that only knowing that you don't use a particular email address or that particular service will help you to identify a phish. Quizzes rarely explicitly address a heuristic – 'Do I have a business relationship with the apparent sender of this message?' – that may be key to the individual recipient, but is less helpful to support staff, email administrators and so on.
- Quizzes don't support (or, at any rate, encourage) the use of tools like *whois* to check the bona fides of a referenced site, so how do you reach a conclusion on whether a site that doesn't use the organization's primary domain is nevertheless genuine? Indeed, if the message purports to come from an institution you don't know or deal with personally, how can you be sure what their primary domain is? The quiz usually makes the implicit assumption of an existing relationship, for the purpose of the quiz ('Imagine that you are a customer…') but doesn't give that contextual information.
- How do you legislate for other attacks such as DNS misdirection, cybersquatting or typosquatting? We have seen quiz samples where the apparent and real target URL were the same.
- Real phish emails are relatively easy to categorize as such for a practised observer. It's not always so easy for even a hardened phish-watcher to confirm that mail is genuine without using other resources. Sometimes it's

easier to guess if you've done a few quizzes, but that's about second guessing the quizzer, not about being security-literate. The point here isn't really about perfect scores in 'off the top of the head' discrimination exercises: on the whole, it's about evaluation based on incomplete data, and the 'correct' answer in such a scenario is always to assume the worst.

- Where legitimate institutions send emails that don't conform with best practice, they actually inadvertently groom the customer on behalf of the scammer. Quiz sites may prefer examples of such mails: mails that conform to good practice such as including the recipient's name are probably easier to categorize, but 'too easy to guess'. In the end, perhaps it's a question of what point the quiz site is trying to make when it includes genuine mails. Of course, the site *has* to include some genuine mails, or it wouldn't be much of a quiz, but that might mean that a quiz isn't the best approach in this case.
- Alternatively, the site may be making the same point about customer grooming. However, it's rare for a site to make this point explicit with reference to a quiz sample, perhaps because of a reluctance to offend the institution from which the sample mail was sent. If a poorly formatted, depersonalized, phish-like message is used as an example of a genuine mail, it may be categorized as fake. When this happens, the participant is 'penalized' or at any rate marked down for being suspicious of a mail that illustrates bad practice. Clearly, the 'wrongness' in this instance should be ascribed to the provider, not to the person taking the quiz.

We still see quiz examples that are based on poor practice. One recent 'genuine' example has no obvious personalization, doesn't refer to any means of accessing the account in question except through an embedded link, and includes such classic phish text as 'Please do not attempt to respond to this email'. The quiz answer relating to this question is basically that 'You can only tell if this is legitimate *if* it's an institution you have an account with *and* the situation they're flagging is one you know applies to your account'. A phish mail example explains in the answer that the scammer used an unspecified zero-day attack to misdirect the recipient to a spoofed site instead of gimmicking the URL so that the real target diverged from the apparent URL.

These examples make essential points, but not as constructively as we might wish.

## CONCLUSION

What are the advantages of the phishing quiz to the participant?

- It has a perceived and actual social benefit. It's bound to help in terms of raising general awareness of the problem, and unless it's horrendously misconceived and/or badly implemented, the subject should learn *something* from it. If they finish it, that is; there is anecdotal evidence that a significant percentage of people who take phishing quizzes don't complete even short ones (8–10 questions). There may be unexplored ergonomic and psychodynamic issues here: for instance, after answering a number of questions and receiving no feedback, it may be that people get bored or discouraged, and don't care enough to complete the quiz.
- It's more fun than most security-related activities, and appeals to the competitive instinct. This is likely to impact on the efficacy of the learning mechanism: '…the more interactive and interesting the training, the more likely the individual is to actually learn and retain something new' [15].

What are the advantages to the quizzer?

- If a quiz succeeds in raising awareness and reducing susceptibility to phishing techniques, that benefit may be shared by the quizzing organization (happier customers, more trust in the internet as a medium for financial transactions, and so on). Even security vendors whose revenue stream is dependent on technological solutions are rarely so cynical as to oppose supplementary educational solutions, or so naïve as to believe that piecemeal educational initiatives pose a major threat to the sales potential of technological solutions.
- Even if it doesn't work, the organization gains brownie points for being socially responsible and doing its bit to address the phishing problem. We don't, of course, suggest that quizzing organizations don't expect their quizzes to have any beneficial impact on the problem, though we do, clearly, have reservations as to how profound this impact is likely to be.
- The kind of quiz we've discussed here is easy to compile, especially the email discrimination tests. Too easy, perhaps: many of the tests we've seen don't show evidence of expert input either from anti-phishing gurus or from educationalists.

Here's the real issue. What are the advantages to the phishing gangs?

- Quizzes present a simplified view of the issues. They focus on one or two aspects of the phishing/black hat economy problem (especially phishing emails) that are particularly susceptible to a 'Janet and John' (US readers may be more familiar with 'Alice and Jerry' or 'Dick and Jane') approach to education. Other aspects (mule recruitment, for instance) could be addressed with a similar approach, but this doesn't seem to happen. Aspects that can't conveniently be addressed with a screenshot are less likely to be addressed, and if they *are* (say as a supplementary sidebar or FAQ entry), they are less likely to be absorbed.
- Most quizzes concentrate on a very limited subset of sample types and issues. They don't usually consider mechanisms in depth, focusing on (some) symptoms rather than the disease.
- A phisher interested in researching the psychodynamics of phish response gets an overview of what the establishment is teaching the masses about phish recognition. The less adequate the teaching, the easier it is to avoid any specified heuristics.
- The subject who does well in some phish quizzes (phizzes?) may be misled into overestimating the significance of the result and a false sense of security.

There is no single solution to the phishing problem, and if there were, it probably wouldn't be education – well, maybe if it was global and well-implemented, but that hasn't happened yet – and technical solutions, important though they are, are outside the scope of this paper.

Too often in security, we see a problem exacerbated by well-meant but ill-founded advice from sources that the everyday user might assume to be authoritative: for example, some of the phished institutions, government agencies, the media and law enforcement agencies. Phished institutions must conform to (and other agencies must promote) best practice:

- Communicating with their customers using personalized messages, expressed in ways that make it harder for phishing gangs to make fraudulent messages look genuine. Can we perhaps suggest exclusive use of snail mail or properly secured electronic channels for sensitive communications? (However, these measures are only useful if the customer is aware that they are in place.)
- Never using email to ask for personal identification information or to link directly to sites, especially on secondary domains or third-party sites.
- Making it easier for customers to get reliable advice and information from customer support facilities in cases of doubt.

Certainly, anyone presuming to give advice on good practice should:

- Be more specific than 'Be careful' and 'Don't go to suspicious websites'.
- Try not to mislead with poor advice or partial information that may be inadequate in some contexts.

What's the *intention* of a phishing quiz? Even a poorly designed quiz raises awareness of the problem, but may be worse than useless if it reinforces wrong assumptions on the part of the quiz participant. Some quizzes seem to promote a service: 'Discrimination is too difficult for your tiny brain; buy our product, or even use our free toolbar/site verification service/whatever'. That's not wrong in itself; a vendor is in the business of selling products or services. If the product or service in question is free, it seems even more churlish to criticize, but there is a problem in that this message fosters dependence, not awareness; worse, that dependence is on a technical solution that is likely to rely on detecting specific instances of malice, rather than a generic class of detection.

A quiz that simply tells you whether you assessed (or guessed) correctly without any further explanation is, it seems to us, of little use. If it's found on a vendor site, it even carries that same implicit 'use our spam service' message.

The best quizzes are, in our humble opinion, those that leave the participant knowing more than they did when they started. The following 'useful things to know' are summarized from a detailed section on recognition heuristics in another of our papers [2]:

- If you don't have a pre-existing relationship with the apparent sender of the message, they shouldn't be sending you requests for sensitive information about an account you don't have (or anything else!).
- Use a specific (dedicated) email address for internet transactions; discard mail to other addresses.
- Untoward urgency ('You must log in within 24 hours or your account will be terminated') is usually intended to panic you into responding inappropriately.
- Requests for sensitive data (credit card numbers, account details, social security numbers, PINs – the more detailed, the more suspicious) sent by email and channelled through direct web links are either malicious or bad practice.
- The more data requested, the more suspicious; these data amount to a substantial definition of financial/social identity. However, an attacker can acquire byte-size lumps of apparently insignificant data over time to aggregate into a full-strength ID theft package.
- If no one complains about bad practice, it won't stop.
- If you respond, do so 'out-of-band': e.g. go to the legitimate site directly (not following links from email), or contact the customer services department or local branch to verify authenticity. There have been (incredibly rare) cases where scams have been elaborate (or bank staff ill-informed or ill-advised) enough for even these measures to be compromised. However, unless you're a gullible billionaire caught up in an elaborate negotiation with the wife of the ex-president of Nigeria, you should be OK.
- Impersonal is suspicious. 'Dear Citibank customer' or 'Dear fredbloggs@bigfoot.com' doesn't qualify as personalized. Even 'Dear John' or 'Dear Donald Trump' isn't proof of personalization: there are many ways to link a name and an email address, and sometimes the process can be automated. If another identifier is used (e.g. an account number or eBay registered name), check it isn't just made up.
- Multiple addressees, a generic mailing list addressee (e.g. 'Client-list') or no addressee (i.e. a blind copy) all suggest random/multiple mailings.
- Any message apparently from someone you already deal with (IRS, your bank, *eBay*) that requires you to re-authenticate online from a link in the message is either fraudulent or incompetent. Embedded phone numbers are also suspicious. Always use known valid numbers and addresses, and pre-established login procedures.
- Pidgin English or poor spelling is suspicious, but impeccable presentation doesn't prove legitimacy.
- There are many techniques for misdirection to a malicious site (URL obfuscation, typosquatting etc.) that echo poor practice by legitimate sites (secondary domains, outsourced web pages, tiny URLs, overlength URLS): verify or discard.
- Look for trust indicators such as https:// and digital certificates, but verify them. In particular, padlock icons are not proof of authenticity.
- Technical tricks to evade standard detection technologies [2] such as image spam, hashbuster graphics or text, obfuscating text and tags, font colour tricks, divergent URLs and so on are a good indication of malice (or at least of spamminess).
- Some of these indicators, however, are of more use and interest to the security professional than to the everyday user and quiz participant.

A final thought: is it helpful to talk about 'ID theft IQ', or 'Phishing IQ?' An educationalist correspondent has pointed out [16] that phishing is a matter of concern 'because users were and are vulnerable because the internet is in their own homes.' It's also been pointed out that phishing risks are not

entirely confined to those old enough to hold a bank account [17].

An 'Intelligence Quotient' is an attempt to index intellectual development and/or ability relative to the rest of a population, but the kind of adaptive social 'intelligence' that 'real' educationalists favour nowadays is actually closer to what this kind of quiz could measure, if implemented with sufficient rigour.

To paraphrase Andrew Klein [18], it's not enough to point out to people that they guessed wrong. You have to show them where to look for better indicators, and whom they should blame for the bad practices that condition them to guess wrong. The trick is not only to raise awareness, but to encourage appropriate responses.

## REFERENCES


[1] Overton, M. An African A-F-F-AIR…. Virus Bulletin, April 2007. Harley, D. Stalkers on your Desktop. The AVIEN Guide to Managing Malware in the Enterprise. Syngress, 2007.

[2] Harley D.; Lee, A. A pretty kettle of phish. http://www.eset.com/download/whitepapers.php.

[3] Abad, C. The economy of phishing: a survey of the operations of the phishing market. http://firstmonday.org/issues/issue10_9/abad. James, L. Phishing Exposed. Syngress, 2005. Harley, D.; Gonzales, E.; Dunham, K.; Melnick, J. Crème de la CyberCrime. The AVIEN Guide to Managing Malware in the Enterprise. Syngress, 2007.

[4] Barrett, D. Bandits on the Information Superhighway. O'Reilly, 1996.

[5] Harley, D.; Slade, R.; Gattiker, U. Metaviruses, hoaxes and related nuisances. Viruses Revealed. Osborne, 2001.

[6] Slade, R. Dictionary of Information Security. Syngress, 2006.

[7] Lee, A. Eset warns about IRS Phishing Scam (http://88.208.205.84/joomla/index.php?option=com_content&task=view&id=1341&Itemid=5); Washkuch F. Jr. Phishing scam uses social security ploy. http://www.scmagazine.com/uk/news/article/604182/phishing-scam-uses-social-security-ploy/.

[8] Mustaca, S. Present and future phishing techniques. Virus Bulletin, September 2006.

[9] Peel, M. Nigeria-related financial crime and its links with Britain. http://www.chathamhouse.org.uk/pdf/research/africa/Nigeria1106.pdf.

[10] Schiller, C.; Binkley J. *et al*. Botnets: the Killer Web App. Syngress 2006.

[11] Schechter, S.E.; Dhamija, R.; Ozment, A.; Fischer, I. The emperor's new security indicators: an evaluation of website authentication and the effect of role playing on usability studies. http://usablesecurity.org/emperor/emperor.pdf.

[12] Whalen, T.; Inkpen, K. Gathering evidence: use of visual security vues in web browsing in graphics interface (2005).

[13] Dhamija, R.; Tygar, J.; Hearst, M. Why phishing works, in human factors in computing systems. 2006.

[14] Kumaraguru, P.; Rhee, Y. W.; Acquisti, A.; Cranor, L.; Hong, J.; Nunge, E. Protecting people from phishing: the design and evaluation of an embedded training email system. CyLab, Carnegie Mellon University 2006.

[15] McDowall, J. Fraud Resource Group. Personal communication, 2007.

[16] Barnes, D.; Harley, D. ICT in education network discussion. http://ictineducation.ning.com/forum/topic/show?id=686727%3ATopic%3A261&page=1&commentId=686727%3AComment%3A601.

[17] Harley, D.; Willems, E.; Harley, J. Teach your children well: ICT security and the younger generation. Proceedings of the Virus Bulletin International Conference, 2005.

[18] Klein, A. SonicWall. Personal communication, 2007.


## GLOSSARY

Bcc (blind carbon copy)

Email receiver field: recipients listed in the 'Bcc' field are not shown in the copies sent to the primary recipient(s) in the 'To' field, or to the secondary recipient(s) in the 'Cc' (carbon copy) field. This terminology dates back to typewritten business letter practice, of course.

Botnet

Network of bot-compromised systems under the control of a bot herder.

Carding

Sometimes used as a synonym for phishing. Also applied to the fraudulent use of an (often stolen) credit card, resulting in direct loss to the *retailer*, rather than the loss of the *card owner's* money or identity. May also apply to a scammer's checking that a stolen card is still valid.

Cousin domains

Domains registered with names that incorporate the names of targeted institutions or a close variant with the intention of setting up a phishing website.

Cross-site scripting (XSS)

A common partial misnomer for an attack where a client-side or server-side vulnerability is used to facilitate an attack against a client application. The kind of techniques sometimes attributed or related to XSS include injecting an arbitrary malicious script into a web transaction, injection of malicious data specially formatted to force the server to misinterpret the input, and even forcing [9] insecure redirect mechanisms.

Cybersquatting

Registration of a domain with the intention of benefiting in some way from a perceived but deceptive or malicious association between the registrant and a legitimate site or business.

Favicon

Abbreviation for 'favourites icon': icon associated with a particular website or webpage and displayed in many current web browsers.

Hashbuster

Some spam filters use a database of 'hashes' to identify spam messages: these are a kind of 'fingerprint' of a message. It has long been common for spammers (among others) to include random text in the subject or body of a message, so as to generate random changes from one spam iteration to another, thus throwing off filters that rely on checksums or hashes. Similar techniques have been applied to image spam.

Keylogger

As applied to phishing, a form of spyware or trojan that records a computer user's keystrokes without his or her knowledge and passes the information on to a criminal, bot herder etc.

Mule

In phishing, usually refers to someone who is involved with money laundering by receiving and forwarding fraudulently acquired funds, goods or services.

Pharming

DNS spoofing is a term applied to the malicious, covert redirection [1] of a web browser from a legitimate site to a different, illegitimate IP address/ web page. This simple technique is effective, because it works even when the user directly enters the correct URL into a browser.

Pump and dump

A form of stock fraud in which the value of stock is artificially inflated so that dishonest speculators can make a profit by selling off when the price is high. This works well for the scammer, but not for the (usually small) company, or for the scam victims whose contribution to the raising of stock value is rewarded by a plummet in value.

Social engineering

Term applied to a wide range of techniques for causing a desired change in behaviour or gaining some advantage by psychological manipulation of an individual or group.

Spear phishing

Phishes aim to hook the users of specific services by pretending to come from a service provider, but the bait is usually distributed more or less randomly – after all, a phishing gang isn't usually able to tell whether the recipient of the message is a customer of that service. Sometimes, though, deceptive mails can be highly targeted as, for example, in some instances of industrial or economic espionage.

Spyware

Generic term for a range of malware such as keyloggers, remote access trojans, backdoor trojans and so on. Malware used for frankly criminal activities such as phishing may also be referred to as crimeware.

Tsunami scams

A range of charity scams and hoaxes allegedly raising funds for victims of the 2004 tsunami. Examples include many 419s and phishing mails.

Typosquatting

Variations on the cybersquatting theme include using slightly misspelt names like 'Barc1ays.com', which may look authentic to a careless observer, but may also catch a careless typist looking for the real site.

Vishing

The use of VoIP as a vector for phishing attacks: approaches used include directing the victim to a spoofed phone number to verify sensitive data, as well as directly approaching the victim, indicating that phone numbers in emails are no safer than URLs.