A Chat at the Old Phishin' Hole

Richard Clayton, Drew Dean, Markus Jakobsson, Steven Myers, and Stuart Subblebine

Phishing is an attack in which victims are lured by official looking email to a fraudulent web-site that appears to be that of a legitimate service provider. The email also provides victims with a convincing reason to log-on to the site. If users are fooled into logging-on, then the attacker is provided with the victims' authentication information for the legitimate service provider, often along with personal information, such as their credit-card data, checking account information or social security data. Successful phishing attacks can result not only in identity and asset theft, but also in more subtle attacks that need not be directly directly harmful to the victim but which have negative consequences for society (for example: money laundering).

Professional studies that have attempted to estimate the direct losses due to phishing in 2004 have come up with widely varying figures: from $150-million to $2.4-billion U.S. dollars. However, all the studies agree that the costs will continue to rise in the foreseeable future unless something is done to educate users and/or technologies are introduced to defeat or limit such attacks. Further, these estimates measure only the direct costs, and do attempt to measure the indirect costs that result from the loss of consumer confidence in the Internet infrastructure and all of the services it can be used to provide. Our panel will look at a broad number of issues relating to the past, present and future of phishing, in order to better understand this growing problem.

We will address topics that include the notion that phishing is a special case of "web-spoofing", an attack that was predicted and researched academically as early as 1996. We will look at the mutual progression of the research and practice of such attacks, and what we can learn from both. We will discuss the fact that phishing is currently a problem, and look at what information consumers are being given to mitigate their risk of exposure; we'll ask if the advice is practical and effective. We will see how the percentage of successful phishing attacks could dramatically increase if phishing attacks begin to make use of contextual information about their victims. It will be argued that such attacks are easily automated, begging the question of how long it will take for such context sensitive attacks to appear in the wild. We will see that phishing-graphs can be used not only to model phishing attacks, but also to quantify the feasibility and economic costs of attacks. We will discuss the issue of mutual authentication, and how it relates to phishing attacks. It will be argued that easy to use mutual authentication protocols could mitigate many of the risks of phishing, and we will discuss one such protocol. Finally, we will deliberate on the likelihood of the advent of a silver-bullet technology that will solve all of our phishing problems.

Who'd phish from the summit of Kilimanjaro?

Richard Clayton

University of Cambridge, Computer Laboratory. <richard.clayton@cl.cam.ac.uk>

Phishing emails are now so convincing that even experts cannot tell what is or is not genuine[1]; though one of my own quiz errors resulted from failing to believe that genuine marketeers could possibly be so clueless! Thus I believe that education of end users will be almost entirely ineffective and education of marketing departments – to remove "click on this" (and HTML generally) from the genuine material – is going to take some time.

Providing end users with one-time passwords (pads of single-use numbers, SecureID tokens, PINs sent by mobile phone) can ensure that phishing only works when there is a real-time, Man-in-the-Middle (MITM), attack, which will immediately deter the bad guys whose technical expertise runs solely to copying websites. However, formal protocol analysis shows that only a handful of the "bag of bits" being passed around can be considered to be authenticated – and the MITM will be able to steal what they wish.

Insisting on SSL (`https`) connections will prevent the use of random URLs for phishing websites and bring the focus back to control of the DNS. However, once the second level (`fakebankname.com`) is secured then the attackers will just move down a level (to `bankname.plausible-second-word.com`). I predict a lot of wasteful activity before the nature of DNS delegation is fully understood.

Insisting on client certificates prevents MITM, but also stops me paying my gas bill from a holiday cybercafé – which is bad for business. But why do I need the same authority to pay the bill as to change the name of the gas company? A range of authentication systems is needed, chosen as the risk varies. The banks could learn from the activity monitoring systems of the credit card companies, and thus ensure that extra authentication is seldom necessary or onerous. For example, a check can be made on the IP address of incoming connections. If the session arrives from a cybercafẽ in Latvia or a web hosting rack in suburban Moscow then Mr. Jones in Acacia Avenue is not connecting directly... if he really does want to set up a new payee then perhaps he could ring the branch manager directly to confirm that he's taking an East European holiday?

To conclude; I can see no silver bullet (I can imagine success for phishing emails that ask for client certificates), and most of the proposed argento-ammunition is useless once the end-user machine is compromised. Nevertheless, a blend of security improvements will freeze out all but the most competent criminals. Society may need a general solution to online security, but the banks only have to persuade the bad guys to move on to more attractive targets. However, the fixes must *not* be introduced one by one, allowing each to be overcome individually. What's needed is a 'Kilimanjaro effect', where the security suddenly dominates the landscape and it will always seem to be a long way to the summit.

[1] MailFrontier Phishing IQ Test II `http://survey.mailfrontier.com/survey`

Helping the Phish Detect the Lure

Steven Myers

School of Informatics,
Indiana University at Bloomington,
Bloomington, IN 47406, USA
samyers@indiana.edu

When a client attempts to interact with an online service provider that performs any form of financial transaction, the service provider requires the client to authenticate itself. This is normally done by having the client provide a username and password that were previously agreed upon, through some procedure, the first time the client attempted to use the services provided by the provider. Asymmetrically, the client does not ask the provider for the same form of authentication. That is, the customer of the bank does not ask the web-page to somehow prove that it is really the bank's web-page. This asymmetry seems to come mostly from an attempt to port security models from the physical to the digital world: I would never expect a physical bank branch to authenticate itself to me through any form other than its branding. However, that is not to say customers don't implicitly authenticate their bank-branches, they do! However, it is a rather implicit authentication that is based on the use of branding and law-enforcement by the banks. Unfortunately, many of the security assumptions that hold in the physical world do not hold in the digital world: the costs of setting up an authentic looking but fraudulent web-page are low; the pay-off for successful phishing attacks is high; and digital law enforcement is weak to non-existent in the digital realm and so the risks are minimal. This makes phishing an attractive type of fraud, and has led to its growing popularity.

In order to reduce the ability of phishers to launch successful attacks, we suggest that users request authentication from their service providers. In other words, we suggest that the client and service provider engage in mutual authentication. While such authentication is easily achievable with public-key cryptography and certificates, this solution is not appealing due to the historical difficulty users have had in understanding these concepts: currently many users automatically accept most certificates that are brought to their attention by web-browsers, regardless of their validity or origin.

We will discuss a protocol for mutual authentication that relies solely on a client being able to remember a password to authenticate him or herself to the service provider, and the ability to recognize —and not recall, as in the case of a password— a unique series of images (or other forms of stimuli, such as sound and touch) corresponding to the appropriate service provider. The client only needs to be educated to realize that if his or her appropriate sequence of images does not appear, then the site is not legitimate and should not be used, nor should any personal information be provided to it. Further, the protocol has the property that it is secure against man-in-the-middle attacks in the random-oracle model.

Modeling and Preventing Phishing Attacks

Markus Jakobsson

School of Informatics,
Indiana University at Bloomington,
Bloomington, IN 47406
www.markus-jakobsson.com

Abstract. A *first contribution* of this paper is a theoretical yet practically applicable model covering a large set of phishing attacks, aimed towards developing an understanding of threats relating to phishing. We model an attack by a *phishing graph* in which nodes correspond to knowledge or access rights, and (directed) edges correspond to means of obtaining information or access rights from already possessed information or access rights – whether this involves interaction with the victim or not. Edges may also be associated with probabilities, costs, or other measures of the hardness of traversing the graph. This allows us to quantify the effort of traversing a graph from some starting node (corresponding to publicly available information) to a target node that corresponds to access to a resource of the attacker's choice. We discuss how to perform economic analysis on the viability of attacks. A quantification of the economical viability of various attacks allows a pinpointing of weak links for which improved security mechanisms would improve overall system security.

A *second contribution* of this paper is the description of what we term a *context aware* phishing attack. This is a particularly threatening attack in that it is likely to be successful *not only* against the most gullible computer users (as is supported by experimental results we present.) A context aware attack is mounted using messages that somehow – from their context – are expected (or even welcomed) by the victim. To draw a parallel from the physical world, most current phishing attacks can be described as somebody who knocks on your door and says you have a problem with your phone, and that if you let him in, he will repair it. A context aware phishing attack, on the other hand, can be described by somebody who first cuts your phone lines as they enter your home, waits for you to contact the phone company to ask them to come and fix the problem – and *then* knocks on your door and says he is from the phone company. We can see that observing or manipulating the context allows an attacker to make his victim lower his guards. As a more technical example, we show how to obtain PayPal passwords from eBay users that do not take unusual measures *particularly intended* to avoid this attack.

Finally, a *third contribution* is a discussion of how to address the threats we describe – both in their specific and generic shapes.

A full version of this paper can be downloaded from www.markus-jakobsson.com