R&D

Tyler Moore

> “These threats are driven by criminals looking for financial gain, not attention-seeking hackers. Yet we understand very little about how these organisations work… let alone how much money they make

Phishing and the economics of e-crime

The amount of dark matter in the online economy is vast, with phishing attacks alone stealing hundreds of millions every year, writes Tyler Moore of the University of Cambridge's Computer Laboratory

Wickedness on the internet is rife. Millions of computers are infected with malware, while hundreds of thousands more are enslaved as botnets sending spam, launching denial-of-service attacks and hosting dubious websites. Phishing sites impersonate banks, fake storefronts peddle non-existent cameras to defraud unsuspecting consumers and websites recruit mules to launder the stolen proceeds.

Each of these threats are driven by criminals looking for financial gain, not attention-seeking hackers. Yet we understand very little about how these organisations work, whether the attacks are in fact related, or how many people carry them out, let alone how much money they make. Thus it is necessary to examine the underlying economics to better understand not only their motive, but also how best to eradicate the threats. Along with my colleague Richard Clayton, we have sought to answer these questions in the context of phishing.[1]

Phishing is the process of enticing people into visiting fraudulent websites and persuading them to enter identity information such as usernames and passwords. This information is then used to impersonate the victim so as to empty their bank account, run fraudulent auctions, launder money, apply for credit cards, and so on. Although most current phishing attacks target the banks, phishing websites regularly appear for businesses as diverse as online auctions (eBay), payment sites (PayPal), share dealers (E*Trade), gambling websites (PartyPoker), social-networking sites (MySpace) and online retailers (Amazon).

The mechanics of phishing

To carry out phishing scams, attackers transmit large numbers of spam emails with links to websites under their control. When a user clicks on the link, he or she is then presented with an accurate imitation of the legitimate company's pages (often including all the links to warnings about fraud), and thus reassured fills in his or her personal details. Although a handful of sites validate these details immediately, it is more common for any response at all to be accepted.

The compromised details are usually emailed to a webmail address, but are sometimes stored in plain text files at the spoof website, awaiting direct collection by the fraudster. Once they have received the compromised details they will discard the obviously fake and then sell on the details to cashiers who will empty the bank accounts, perhaps transferring the money via a so-called mule who has been recruited via further spam email seeking 'financial consultants' to accept and relay payments for a commission. The spoof website is sometimes hosted on 'free' webspace, where just anyone can register and upload pages, but it is more usually placed on a compromised machine; perhaps a residential machine, but often a server in a data centre.

The banks (and other organisations being impersonated) are dealing with the fake websites through 'take-down' procedures, so that there is nothing there for a misled visitor to see. The bank sends a take-down request to the operator of the free webspace, or in the case of a compromised machine, to the relevant internet service provider who will temporarily remove it from the internet or otherwise ensure that the offending web pages are disabled. Where a domain name has been registered by a phishing attacker, the defenders will ask the domain name registrar to suspend the offending domain. However, not all ISPs and registrars are equally co-operative and knowing that a phishing site exists does not automatically cause its removal.

To determine the effectiveness of the take-down strategy, we monitored the availability of several thousand phishing websites in spring 2007 using reports from PhishTank.[2] In the process we learnt a great deal about the number of attacks taking place and the effectiveness of take-down strategies. A great disparity in take-down performance was identified, and examples of attacker innovation that slow down removal were found.

The results show that a typical phishing website can be visited for an average of 62 hours. But this average is skewed by a number of very long-lived sites, of up to 17 weeks. While in the minority, there are too many long-lived sites to be written off as unimportant outliers. Indeed, we can fit the average lifetime to a long-tailed lognormal distribution. Similarly skewed distributions have been found for the average time between software failures.

Rock-phish attacks

Not all phishing attacks work in the manner just described. The 'rock-phish' gang[3] has adapted its attack strategy to evade detection and maximise phishing-site availability. It has separated out the elements of the attack while adding redundancy in the face of take-down requests.

The gang first purchases a number of domain names with short, generally meaningless, names such as *lof80.info*. The email spam then contains a long URL such as *http://www.bank.com.id123.lof80.info/vr* where the first part of the URL is intended to make the site appear genuine and a mechanism such as `wildcard DNS' can be used to resolve all such variants to a particular IP address.

It then maps each of the domain names to a dynamic pool of compromised machines according to a gang-controlled name server. Each compromised machine runs a proxy system that relays requests to a back-end server system. This server is loaded with a large number (up to 20 at a time) of fake bank websites, all of which are available from any of the rock-phish machines. However, which bank site is reached depends solely upon the url-path, after the first /. (Because the gang use proxies, the real servers – that hold all the web pages and collate the stolen information – can be located almost anywhere.)

We analysed rock-phishing sites during a period of eight weeks between February and April 2007. During this time, we collected 18 680 PhishTank reports which we categorised as rock-phish – 52.6% of all PhishTank reports for the time period. While these reports are intended to be unique, we identified many duplicates due to the use of unique URLs as described above. This yielded a significant saving in effort, since just 421 canonical rock-phish domain names were observed. Rock-phish sites used 125 IP addresses that were found to be operational for any duration. In all, the rock-phish sites impersonated 21 different banks and three other organisations.

> "We also retrieved text files that recorded victim responses from a number of sites. Here we found that around half were clearly fake (names like "Die Spammer"), while the rest appeared legitimate

For traditional phishing sites, removing either the hosting website or the domain (if only used for phishing) is sufficient to remove a phishing site. However, rock-phish sites continue to work for a particular domain that is mentioned in a spam email, provided that they can be resolved to at least one working IP address. Whenever one site is removed, the name server resolves to machines still hosting a working copy of the proxy. While proxy machines and domains were removed constantly by the banks, they were replenished frequently enough to keep a number of sites working every day. Hence, the rock-phish strategy has effectively undermined the bank's take-down response.

Rock-phish domains and IPs also last longer than ordinary phishing sites: rock-phish domains last for 95 hours on average while rock IPs last 172 hours, compared to 62 hours for regular phishing sites. These longer lifetimes occur despite impersonating around 20 banks simultaneously, which should draw the attention of more banks. One explanation for the longer lifetimes is that their attack method is not widely understood, leading to sluggish responses. Splitting up the components of the phishing attack (domains, compromised machines and hosting servers) obfuscates the phishing behaviour so that each individual decision maker (the domain registrar, ISP system administrator) cannot recognise the nature of the attack as easily when an impersonated domain name is used (such as barclaysbankk.com), or HTML for a bank site is found in a hidden sub-directory on a hijacked machine.

Fast-flux domains

During data collection, we witnessed a further innovation by the gang dubbed 'fast-flux' by the anti-phishing community. It arranged for its domains to resolve to a set of five IP addresses for a short period, then switched to another five. This of course 'eats up' many hundreds of IP addresses a week (4572 addresses during our eight-week collection period), but the agility makes it almost entirely impractical to 'take down' the hosting machines. The gang is likely to have large numbers of compromised machines available (probably in the form of botnets), since if they are not used to serve up phishing websites, they are available for sending email spam.

Fast-flux IP addresses remained alive for 139 hours on average, slightly less time than for rock-phish IPs. This is likely a reflection of the nature of the compromised hosts – consumer machines with dynamic IP address assignment – since the sites were not actively taken down. Domains were very long-lived (252 hours on average). This is because many fast-flux sites were not actually phishing sites at all. Instead, many were hosting mule-recruitment sites or selling diet pills and Viagra. This provides further evidence of the overlap present in the dark-side economy.

Estimating phishing's cost

In order to gain a better understanding of how many users respond to phishing attacks, we gathered data about how many visitors a typical phishing website received, as well as what proportion of responses are legitimate. Publicly-available web page usage statistics, collated by the sites where the phishing pages are residing, were collected. Webalizer4 is a particularly popular package, which is often set up by default in a world-readable state on the type of web servers that seem to be regularly compromised. These statistical reports provide daily updates as to which URLs are visited, and these can be used to determine the total number of visitors and how many reached the 'thank you' page that is generally provided once personal data has been uploaded. For around 30 sites, a record of the number of visits to the 'thank you' page for several days while the site was alive was obtained. On average, phishing sites dupe around 20 victims per day until being removed.

We also retrieved text files that recorded victim responses from a number of sites. Approximately half were clearly fake (names like "Die Spammer"), while the rest appeared legitimate. Using this empirical data, it is possible to estimate the cost imposed by phishing attacks. Of course, we are using a number of rather fuzzy estimates, so substantial refinement may be possible in the future as better figures come to light.

We first consider the cost imposed by ordinary (not rock-phish or fast-flux) phishing sites. Data was collected for eight weeks and confirmed 1438 banking phishing sites. Extrapolating, we might expect 9347 sites per year. These particular sites remain operational for around 62 hours on average, which yields approximately 30 victims. The analyst firm Gartner has estimated the cost of identity theft to be $572 per victim (£283, €418).[5] Hence, the estimated annual loss due to ordinary phishing sites is 9347 multiplied by 30, or 280 410 victims. If each loses $57, the total loss is $160.4m (£79m, €117m). Gartner estimates that 3.5 million Americans give away their details annually, which leads to an estimated loss of $2bn.

> **"We estimate, at an absolute minimum, that at least $320m is lost annually due to phishing scams**

We cannot reliably provide an estimate for the costs of rock-phish and fast-flux phishing scams since we do not have similar response data. However, given that the rock-phish gang send a large proportion of all spam, which drives visitor numbers, it is fair to assume that they steal at least as much money as ordinary phishers.

Thus, we estimate, at an absolute minimum, that at least $320m (£159m, €235m) is lost annually due to phishing scams. The disparity with Gartner's total of $2bn is doubtless due to the extremely rough approximations used, both by ourselves and Gartner. But the difference will also be accounted for by the other ways in which personal data can be stolen, for example the theft of merchant databases, and the activities of malware that scans files or operates keyloggers.

Transparency means security

There is significant variation in the lifetime of phishing sites. Some banks perform better than others: the worst performers take nearly one week to remove sites, while the best are removed in under 12 hours. Similarly, some ISPs and registrars do better than others. For example, Yahoo!'s free-hosting sites are removed in around one day, far better than average. However, such differences are difficult to discern at present. Economists refer to this as asymmetric information, where the parties in a position to take action cannot be observed. Asymmetric information prevents efficient outcomes because it encourages free-riding. Why should banks improve their response to phishing if customers cannot tell the difference? Why should ISPs take costly measures to improve take-down times if doing so goes undetected by others?

In this environment, smart attackers will target the sluggish banks and compromise machines hosted by ISPs which are slow to remove them. This is also true for other types of electronic crime, since the internet is comprised of many administrative and legal jurisdictions. When the weakest link prevails, be it an ISP that doesn't respond to take-down requests or a law enforcement agency that does not go after cyber-criminals, attackers move to exploit it. Therefore, it is essential to provide measurements that compare performance for overcoming the information gap. For example, a league table of ISP response to phishing take-down requests could identify laggards to raise the overall security level.

Conclusion

Attacks on the internet have matured to the point that they are driven by greed. As such, economics is increasingly relevant for understanding their effect and behaviour. Having studied phishing, it is clear from the rock-phish gang that attackers can successfully adapt their strategies to overcome defence mechanisms. Furthermore, they are likely to make a lot of money doing so: we conservatively estimate that phishing rakes in at least $320m per annum. Fortunately, economics can also inform the best response for defenders. Here, one step in the right direction is to improve transparency by providing comparative performance measures for responsible banks and ISPs.

References

1) This article draws upon the following technical paper: T Moore and R Clayton, 'Examining the Impact of Website Take-down on Phishing', APWG eCrime Researcher's Summit, 4-5 October 2007, www.cl.cam.ac.uk/~twm29/ecrime07.pdf
2) www.phishtank.com
3) See www.infoworld.com/article/06/12/12/HNrockphish_1.html for an article describing the gang
4) www.mrunix.net/webalizer/
5) See www.gartner.com/it/page.jsp?id=498245