Some Psychological Factors of Successful Phishing

Don Mosley, Graduate Student in Information Security, East Carolina University

Abstract – **The successful "phishing" attack relies on the victim's willingness to divulge sensitive personal data to a non-legitimate source in response to an email request or an invitation to a web site. This paper will look at some of the psychological mechanisms involved in these types of scams and what the future might hold.**

There is no doubt that the rise of the Internet has added a new dimension to our lives and made drastic changes to some activities. Going to a movie tonight? Check the listings on the web, it's quicker than finding it in the newspaper. Who was the actress who played Mindy on that TV show? Hit the Internet Movie Database web site (http://www.imdb.com) and see the career history of Pam Dawber. Notice that line of filled grocery carts waiting by the door at the supermarket? Those are for folks who did their grocery shopping on the web and just have time to run by and pick them up. Uncle Albert's birthday is TODAY? Send him a cute card through email and he'll never know you forgot. Need the fourth book of the Earthsea Trilogy? Check a dozen online booksellers without driving anywhere. Did your kid mess up in school today? Look for the teacher's email.

Just as the World Wide Web and email have altered our personal lives and most avenues of commerce, another business activity has been transformed – crime. The Internet has opened up new mechanisms for the perpetrators of crime. I'm sure you remember scenes from movies of fifteen years ago where the brilliant criminal lugs his computer to a phone booth, dials up the bank, and transfers $100 million dollars to his secret Swiss account – it's not likely that such exploits could happen today. Banks now have elaborate safeguards in place to render this type of direct assault all but impossible. Criminals have turned to exploiting the one item that has always been the weakest link in any security endeavor. People.

This paper will explore the phenomenon of Phishing, why we are susceptible, and what measures can counteract it.

Phishing is a fraudulent email that attempts to get you to divulge personal data that can then be used for illegitimate purposes. "The word 'phishing' originally comes from the analogy that early Internet criminals used email lures to 'phish' for passwords and financial data from a sea of Internet users. The use of 'ph' in the terminology is partly lost in the annals of time, but most likely linked to popular hacker naming conventions such as 'Phreaks' which traces back to early hackers who were involved in 'phreaking' – the hacking of telephone systems."[1]

The total volume of email, estimated to be 17.5 billion messages per day (as of 31 Jan. 2005)[2] is about 67% spam[3]. That makes 11.7 billion spam messages per day. Emails designated as "phishing" are a subcategory of spam in general. In a survey conducted by First Data in late 2004, 36% of consumers reported having received phishing emails and 5% reported that they supplied the requested personal data. Of those who fell for the scam, 45% reported that they had subsequently suffered from fraudulent transactions.[4]

The Anti-Phishing Working Group is an association focused on eliminating fraud and identity theft caused by phishing and email spoofing. They compile a monthly report of statistics related to phishing. Here are some data from the January 2006 report:

- Number of unique phishing reports received in January: **17,877** (up 39% from January 2005)
- Number of unique phishing sites received in January: **9715**
- Number of brands hijacked by phishing campaigns in January: **101**
- Number of brands comprising the top 80% of phishing campaigns in January: **6**
- Country hosting the most phishing websites in January: **United States**
- Contain some form of target name in URL: **45 %**
- No hostname just IP address: **30 %**
- Average time online for site: **5.0 days**
- Longest time online for site: **31 days**[5]

We'll look deeper into some of these numbers later on. It's obvious that phishing is currently a major problem and that spam in general is consuming a huge portion of Internet bandwidth and resources.

Let's consider some of the reasons people fall victim to phishing scams.

1. Trust of authority

   Humanity has been “designed” to inherently trust authority. Our current society has been built up since the dawn of human ancestors to trust, obey, and follow some type of authority figure whether this be a god, a strong alpha male, parents, the school principal, a rock star, or a political leader. Obedience to commands occurs almost at the genetic level. When Daddy yells “Bring me the newspaper!” the kid’s muscles start moving before he has time to think about responding. When a phishing email arrives marked as “High Priority” that threatens to close our bank account unless we update our data immediately, it engages the same authority response mechanisms that we've obeyed for millennia. In our modern culture, the old markers of authority – physical strength, aggressiveness, ruthlessness – have largely given way to signs of economic power. “He's richer than I am, so he must be a better man”. If you equate market capitalization with GDP then Bank of America is the 28th most powerful country in the world[6]. If you receive a personal email purported to come from BOA questioning the validity of your account data, you will have a strong compulsion to respond, and respond quickly.

2. Textual and graphics presentation lacks traditional clues of validity

   Our socio-economic culture has given us certain modes of discrimination regarding the validity of business institutions. We don't hesitate to deposit our paycheck at the corner bank because it's been there for years, our parents deposited checks there, and it “looks” like a bank. We'll go into the huge brick building at the mall with the large yellow sign “Best Buy” to buy a new TV set instead of to the guy selling TVs out of his van in an abandoned parking lot. We know we can trust our doctor's advice on health matters because he “looks” like a doctor. When we're on a car trip, we are more likely to stop for gas at a modern, well lit, 24 pump mega-station than a run down, 2 pump gas station built in the 40's. Most people feel that they can tell an honest man by looking him in the eye. You can spot a “professional” panhandler before he gets to the fourth word in his spiel. Without clues from the verbal and physical realms, our ability to determine the validity of business transactions is diminished. This is a cornerstone of the direct mail advertising business. If a piece of mail resembles some type of official correspondence, you are much more likely to open it. Car dealers send sales flyers in manila envelopes stamped “Official Business” that look like the envelopes tax refund checks are mailed in. Banks send credit card offers in large cardboard envelopes that are almost indistinguishable from FedEx overnight packages. Political advertisements are adorned with all manner of patriotic symbols to help us link the candidate with our nationalistic feelings.

3. Email and web pages can look real

Want to email your mother a picture of the new car you just bought? Hit the automaker's web site and cut & paste the jpg image from your browser. It's so much faster and better looking than snapping a picture yourself. Want to create a fraudulent web site that looks exactly like the home page of Bank of America? Just visit their web site and copy the elements from your browser.

“The use of symbols laden with familiarity and repute lends legitimacy (or the illusion of legitimacy) to information—whether accurate or fraudulent—that is placed on the imitating page. Deception is possible because the symbols that represent a trusted company are no more 'real' than the symbols that are reproduced for a fictitious company.”[7]

Certain elements of dynamic web content can be difficult to copy directly but are often easy enough to fake, especially when 100% accuracy is not required. Email messages are usually easier to replicate than web pages since their elements are predominately text or static HTML and associated images. Hyperlinks are easily subverted since the visible tag does not have to match the URL that your click will actually redirect your browser to. The link can look like “http://bankofamerica.com/login” but the URL could actually link to “http://bankofcrime.com/got_your_login”

The possession of the symbols of authority is so often mistaken for authority itself. The literature of criminology is loaded with stories of bank robbers dressed in the uniforms of security guards, prisoners walking out of jails by wearing what look like police uniforms, of criminals fleeing the crime scene by sticking a flashing blue light on the top of their car. Countless victims have been intimidated by toy guns or crude sculptures of guns. Just look at the story of Frank Abagnale as popularized in the recent movie “Catch Me if You Can” to see how easy it was to be accepted as an airline pilot or medical doctor simply by acquiring the symbols of these professions.[8]

4. Clues to the fraudulent nature of phishing scams are often below the threshold of the average recipient

Prior to the internet, the typical scam or swindle scenario usually involved face to face or verbal interaction between swindler and victim. “Hey buddy – want to buy a Rolex?” “Ma'am, I was driving through your neighborhood and noticed that your shingles were about to blow off.” With careful attention to the peripheral clues present in the details surrounding the transaction, the victim has a much greater chance of recognizing and avoiding the swindle.[9] Is the “roofer” driving a car with no ladders or roofing materials? Why is the cheap Rolex watch not actually ticking? Does the presenter of this great deal actually look trustworthy? Phishing emails are totally devoid of the traditional clues that assist in our judgment of the offers veracity. The pre-electronic world of scams frequently involved a direct appeal to the victim's greed, a deal too good to pass up. These offers would have a tendency to trigger deeper scrutiny as the intended victim would try to figure out what the “catch” was.[10] There is another type of email scam that makes this appeal to greed, the “Nigerian 419” scam where the victim is offered a few million dollars to participate in the end stages of a money laundering scheme.[11] The scams presented by phishing emails are almost always couched in terms of “verify your existing data or be terminated” or “please help us update our records” or “confirm your account credentials to protect from fraud”. Our innate resistance to being motivated by greed does not come into play.

The clues that can be gleaned from a typical phishing email tend to be of a technical nature that might not be obvious to the average email user. A user is going to look at the tag of an embedded link in a message instead of the browser's address bar after he clicks the link. A typical desktop arrangement would have the address bar positioned very near the top of the screen, out of the users field of concentration on the center portion of the screen. The tag, which we will assume has the proper host and domain, can have a larger font or bold characters or even be a small decorative image that will make it stand out within the message body. The address bar which would contain the URL with a bogus domain or an IP address shows up in plain text only and is surrounded by more distracting visual elements in the adjacent tool bars. The typical email user doesn't have a clue what an IP address is or why it's a bad sign in a URL. This user would not hesitate to trust a URL of the form “http://bank-america-loginverification.net”.

In the early days of the phishing plague, the emails were frequently composed by non-native English speakers and often contained blatant errors in spelling, grammar, and punctuation. As the messages have become more polished, the compositional errors have fallen below the average threshold for triggering alarms. The average American citizen is not known for his prowess at spelling. “The typical American has the spelling ability of a 6th grader.”[12] The United States is currently the leader in hosting phishing websites at 37% of the worldwide total.[13] Chances are now good that your phishing email will be written by “the average American”.

A recent study shows that a well designed phishing web site was able to fool 90% of the study subjects. 68% of the subjects totally ignored and clicked through all the warning dialog boxes a browser pops up when presented with a fraudulent SSL certificate.[14]

Probably the hardest concept for the average email user to grasp is the fallaciousness of the core message presented by the typical phishing email. “Why do I need to email my account number and PIN to the bank? Don't they already know them?” “If PayPal is going to send me an email demanding that I verify my account credentials, why don't the put my name in the body of the message?” A few moments of reflection would make you realize that the bank is going to send you a physical letter if there is any problem with your account, a registered letter if the problem is serious. Email does not have the reliability or auditability that a bank would require for any customer account transactions. You could also figure out why that email from PayPal didn't contain your full name – the real sender never had your name, just an email address plucked from a compiled list of millions. Or an address generated by iterating through every possible text string combination pre-pended to your mail provider's domain name.

What does the future hold for the nature of phishing attacks?

The creators of phishing scams may not be rocket scientists but they have exhibited a great deal of cleverness. Just as the successful practitioners of face-to-face scams in the old days were masters of deciphering and manipulating human nature, modern phishers are showing the same resourceful traits. They are quick to abandon methods with low return rates and to enhance and expand those tactics with greater than average success rates.

The use of "spear phishing" techniques will increase. A regular phishing attack broadcasts email to many thousands or millions of email addresses. In the spear phishing variant, the sender has a narrow list of recipients such as email accounts within a specific company or the congregation of a single church. This targeted list can contain other personal data (such as your full name) in addition to the email address. Imagine what your reaction would be when you receive an email that says it comes from the HR department with the message "Dear Mary Jones – please correct last weeks time card or your paycheck cannot be processed. Hit this link to login". Assuming that the mail headers have been sufficiently altered, this kind of message will be almost impossible to ignore.[15]

The community of phishing perpetrators will increase their mutual cooperation and start a commerce in emailing lists, very much like the legitimate business in mailing lists that provides the underpinning for the direct mail industry. The email lists would grow and propagate in near real time since positive responses would show up in a matter of hours. You could steal a fellow's bank login credentials, drain his account, and sell his name to a dozen other phishers before he even realized that he'd been victimized. A list of known-gullible phishing victims would be a valuable commodity.

Even as the providers of anti-phishing tools are increasing the power and effectiveness of their blocking software, the phishers are ramping up the sophistication of their attacks. One method sure to increase is to send an HTML encoded email that actually displays as plain text. Hidden within the HTML is a large portion of random text with the font set to display in white. The text is aimed at getting past spam blockers looking at a specific set of text strings. Because it's white on a white background, the recipient never sees it. Because it looks like a plain text message, the user is off guard to thinking of looking for any HTML tricks.[16]

Web pages can be designed to turn off the browser's own address bar and to display a carefully crafted and positioned graphic image that appears to be a legitimate address bar. The viewer sees a URL for the bank's login page and the bogus URL is hidden. Some web browsers can be tricked into launching a bogus popup box. Even though the user is at the legitimate web site of the bank, this popup will claim to be a "Security Confirmation" box for the user to login to. Their credentials are then harvested by the popup code and sent to the phisher.[17]

There are a couple of methods coming into vogue in the phishing community to more thoroughly alleviate the problem of displaying bogus URL addresses. One method is to alter the addressing data supplied by DNS servers to the client machines. DNS cache poisoning occurs when the local cache, say at the local firewall or ISP router, of IP addresses resolved by primary DNS servers gets corrupted to point to the IP address of the bogus web server. The browser address bar will say "http://bankofamerica.com/login" but instead of pointing to the IP address of the real BOA network, the name will point to the criminal's web server on the bogus network. Network managers are taking steps to prevent this type of infrastructure attack.[18] A much more insidious method of DNS hijacking is to get the web user to execute a piece of malware that resets a registry key that tells the operating system to look first in the local hosts file before doing a DNS query to resolve a name to an IP address and then loads this file with bogus data. Thus, the local file "C:\Windows\system32\drivers\etc\hosts" will contain the mapping:

```
102.54.94.97     www.bankofamerica.com
```

and the victim's web browser looks like it's going to BOA but will hit the phisher's server at the 102.... address instead of the real BOA address normally supplied by the DNS service. This ruse is not easily detectable by the average user.

The use of Unicode and other extended character encodings in domain name registrations will aide the phisher in URL spoofing. The English character a has a completely different encoding than the German character ä such that the domain **bankofämerica.com** could be legitimately registered as a distinct domain not affiliated with **bankofamerica.com** The average web user would be unlikely to notice the umlaut above the 'a' in the browser address bar. Other character sets in use throughout the world can have certain characters with visible differences even less noticeable than this example.

What can be done to better defend against phishing attacks?

Certainly we should continue and expand efforts at educating the end user, the recipient of the phishing email. We should also be realistic and realize that any education effort will only meet with limited success. There will always be a certain percentage of the population that will respond to these email scams, probably an increasing percentage as the sophistication of the presentation increases.

There are existing and proposed technological measures that could cut down on the volume and effectiveness of phishing attacks. We have to realize that all the underlying infrastructure that global email, the World Wide Web, the Internet, and local area networks is built on was conceived, designed, and implemented before technologists gave any thought to the concepts of security. All of our current methods of authentication, authorization, and accounting are just "bags hung on the side of the machine"[19]. There is always a chance that someone will devise a workable and uncrackable security add-on that renders phishing obsolete but odds are not good. Short of completely redesigning networking and all the tools built on top, the best we can hope for is a series of stop-gap measures that hope to stay one jump ahead of the criminals.

There are mechanisms that mail servers and relays use for “black listing” and “white listing” the sources of email messages. In the early days of spam and phishing, the perpetrators would acquire a legitimate account with an ISP, hang their servers directly on the internet, and start spewing millions of junk emails through this pipe. The blacklist concept uses methods to recognize the network originations of bogus email and to trigger the local mail server to reject or drop all messages originating from that source. This method could escalate to the point where the entire domain of a national ISP would get blacklisted. In practice, blacklists had a fairly high rate of blocking legitimate email. The white list had similar functionality but typically relied on the end user to set the filters. The email client would query the end user on each message as to whether it was from a valid source or not. Valid email would be added to the white list and all subsequent messages from the same source would pass through without question. Variations of this mechanism would send an automated email back to the sender asking to confirm that they really sent the mail. A positive response would then allow the original message to get through. This would eliminate spam since it typically comes from a bogus sender. It would also tend to aggravate legitimate senders. There are other methods floating around of authenticating email servers and validating the messages they originate but none of the methods is without flaw and none have been widely implemented yet. There are various proposals to charge a fee to the senders of email. The theory is that charging a few pennies to the average user who sends less than ten messages a day would not be noticed but the spammer who sends ten million messages a day would be put out of business. This would also affect the “legitimate” senders of bulk email. There is no agreement on how to best to collect and distribute such fees.

There are various forms of email blockers. One method is to outsource your organization's entire email system to a commercial service provider. These providers have a variety of tools to detect and block broadcast spam and phishing messages. Some of these rely on a network of “sensors” - if they detect ten thousand BOA emails with the same content arriving at the incoming mail server of company A then they can add this pattern to filters of company B. As phishers turn more toward targeted spear phishing campaigns, where they send out a thousand messages to specific recipients instead of a million messages to every email address under the sun, these methods of bulk detection will become less effective.

One of the common fraud prevention techniques that's on just about everyone's phishing prevention list – don't click the link in the email, type it into your browser by hand. This is a very good suggestion except that the majority of web users are lousy typists. It is very easy to mis-type a domain name into the browser's address bar. With the example we've been using: you intend to type “www.bankofamerica.com” but you actually type in “www.bamkofamerica.com”. A phisher could register the domain bamkofamerica.com , set up web pages that look like the real BOA site, and capture your login credentials. An organization with a strong web presence needs to control as many permutations of its real domain name as possible. This would include registering in the alternate top level domains (.com, .org, .net, and others).

The one weapon we could deploy against phishing scams that has the highest probability of success also has the lowest possibility of implementation. People need to be less trusting. If we began to view every piece of unsolicited commercial email with a larger degree of skepticism, then their success rate would plummet. Perhaps the coming generations of email users who have grown up with the Internet will have a much lower susceptibility to phishing scams. All of the technology fixes, short of a total redesign of the infrastructure, will eventually be bypassed by the scammers. So there will always be some portion of phishing emails that will arrive at their destination. If we could decrease the success rate of those that do get through to below some low threshold point, the phishers would close up shop and move on to some other more lucrative field of crime. Phishing has a very low production cost so even a modest return rate makes it pay off.

[1] Gunter Ollmann, "The Phishing Guide", NGSSoftware Insight Security Research, 9 Sept. 2004, http://www.ngsconsulting.com/papers/NISR-WP-Phishing.pdf

[2] IronPort Threat Operations Center, "Global Email Volume by Week", 31 Jan, 2005, http://www.ironport.com/toc/

[3] MX Logic, Inc., "MX Logic Reports Spam ...", 22 Sept. 2005, http://www.mxlogic.com/news_events/press_releases/09_22_05_SpamStats.html

[4] First Data Corp., "New Identity Theft Survey ...", 17 May 2005, http://news.firstdata.com/media/ReleaseDetail.cfm?ReleaseID=163659

[5] Anti-Phishing Working Group (APWG), "Phishing Activity Trends Report January, 2006", http://www.antiphishing.org/reports/apwg_report_jan_2006.pdf

[6] Wikipedia, "List of Countries by GDP", IMF data from 2004, http://en.wikipedia.org/wiki/List_of_countries_by_GDP

[7] Annette Beresford, "Foucault's Theory of Governance and the Deterrence of Internet Fraud", March 2003, Journal of Administration and Society

[8] Maureen Peck, "To Err is Human", Feb. 2006, Journal of Financial Planning

[9] J. Langenderfer; T. Shrimp, "Consumer Vulnerability to Scams, Swindles, and Fraud ...", Jul 2001, "Psychology & Marketing"

[10] Langenderfer and Shrimp

[11] Ayoku Ojedokun, "The Evolving Sophistication of Internet Abuses in Africa", The International Information & Library Review (2005) 37, 11–17

[12] Donna Howard, English Teacher, Grace Christian School, March 2006, private email

[13] APWG, Jan. 2006

[14] Rachna Dhamija, J.D. Tygar, Marti Hearst, "Why Phishing Works", Conference on Human Factors in Computing Systems, CHI 2006

[15] Timothy L. O'Brien, "New 'phishing' tactics test security", Dec. 5, 2005, New York Times

[16] Ollman, "The Phishing Guide ..."

[17] Jennifer Lynch, "Identity Theft in Cyberspace: Crime Control Methods and their Effectiveness in Combating Phishing Attacks", Berkley Technology Law Journal, 2005 vol. 20

[18] D. Miyamoto, H. Hazeyama, Y. Kadobayashi, "SPS: A Simple Filtering Algorithm to Thwart Phishing Attacks", Lecture Notes in Computer Science, Volume 3837 / 2005

[19] Tracy Kidder, *Soul of a New Machine*, (Random House, 1997)