Phishing Awareness Exercises

Ronald Dodge. Ericka Rovira, Radwick Zachary and Shevchik Joseph, *United States Military Academy*

Abstract – The vulnerability of users to social engineering is well known, however very few techniques have been developed to successfully mitigate the threats users unwittingly expose our infrastructure to. Annual training and awareness campaigns have done little keep users vigilant against the many forms social engineering, especially phishing emails. Phishing is regarded as one of the most effective social engineering attacks. In this paper we describe an effort to increase the awareness of users through a campaign of training, policies, and assessment.

Index terms – social engineering, phishing

I. INTRODUCTION

In 2000, the US Census reported that nearly 60% of American households owned a personal computer and utilized it on a daily basis. This increased technological capacity allowed Americans the ability to communicate with someone through e-mail and perform tasks like paying the bills and shopping. Unfortunately, all of the increased benefits associated with computer technology came at the cost of personal security. [1]

As users continue to utilize computers and access the internet daily, they become more susceptible to a fraudulent scheme known as "phishing." Phishing is a form of electronic deception in which an attacker tries to obtain personal information by mimicking a trustworthy entity. [2] Phishers try to lure victims to falsified websites, usually through spoofed emails, by "employing both social engineering incentives and technical subterfuge to steal consumers' personal identity data and financial account credentials". [3]

Names, social security numbers, financial account passwords, credit card numbers, and bank account information are what phishers desire most, and they try to capture this data and use it for their own personal gain. [3] The damages caused by phishing can range from identity theft to financial loss. Billions of dollars are lost each year from financial institutions and millions are left with compromised identities and destroyed credit. [1] Also, falling victim to a phishing attack can also be seen as a security risk where attackers have a higher probability of accessing secure networks like governmental agencies. [3]

Furthermore, phishers have become very sophisticated over the years through the use of a particular strategy known as "spear phishing." This form of phishing involves acquiring background information of a user or group of user and utilizing this information to trick them into submitting information. [4] For example, if a computer user has a history of buying electronics from Amazon.com, a spear phisher would attempt to send them a spoofed email that offers them a discount on their next Amazon.com purchase if they offer to give certain personal information in exchange.

The study conducted by the Anti-Phishing Working Group in 2011 [4] reported a high of 56,362 unique phishing attempts occurred in a single month. The numbers of phishing attacks have and are expected to increase and become even more sophisticated as technology continues to advance at a rapid rate. As a result of increased sophistication and number of attacks, it has become vitally important to teach and train employees and computer users alike in an effort to reduce the success rates of phishing attempts. Further results indicate that the volume of phishing has continued to grow rapidly, with one estimate citing approximately 8 million daily phishing attempts worldwide.

Although there has been an increased effort to educate and train computer users against phishing attacks in both military and civilian sectors, little success has been achieved using standard annual phishing awareness training. This current study attempts to begin research on the most effective anti-phishing training methods.
In order to explore these research aims, an experiment was developed where phishing emails were created and sent to a controlled user population. The participants in this study were provided different training techniques to evaluate their effectiveness, while simultaneously victims of these attacks were studied to see which personality traits they exhibited. To properly develop this experiment, first an exploratory literature review was conducted to uncover the foundation needed to design an effective study. The paper is organized as following: section 2 provides a review of existing work, section 3 presents the experimental methodology. Section 4 provides the

ISBN 1-933510-96-X/$15.00

results. Section 5 concludes with a discussion of future efforts.

II. Previous Work

In the development of this experimental study it became imperative to first understand phishing in more detail. In the study, *Why Phishing Works?*, conducted in collaboration with both Harvard and University of California Berkley, an experiment was executed to test which specific phishing attack strategies were successful and why? This experiment was conducted using 22 participants who were shown 20 different websites and asked to determine which ones were fraudulent. The results showed that the more sophisticated website fooled 90% of participants and this success rate was attributed to user's lack of computer system knowledge and their overall limited attention capacity for security indicators. [5]

The study also uncovered that 23% of participants did not even look at the address bar, status bar, or the security indicators provided on the website's outer frame. Also, many users were even unaware that the padlock icon located in the browser window was an indicator of whether or not the page they were viewing was delivered securely. These indicators were overlooked because users were first uninformed of potential indicators and also phishing attackers utilized several deceptive and malicious strategies. These malicious strategies included using visual deceptive text, masking underlying text, and mimicking actual trustworthy websites as closely as possible. [5]

One of the first studies conducted to test user susceptibility was conducted from 2003 to 2005 by the United States Military Academy. [6] In this study, the student body (over 4000 students) was sent exercise phishing emails to test the awareness level. The emails were distributed equally between emails with attachments, emails asking for sensitive information, and emails with an embedded URL. The result showed that before the students were provided formal training, their 'failure rate' was approximately 40%. The exercise was run three times; the last two after a training program was implemented. The failure rate declined to an average 24% after the training program.

In a study conducted by Carnegie Mellon University, titled, Who Falls for Phish? A Demographic Analysis of Phishing Susceptibility and Effectiveness of Interventions, 1001 online surveys were administered to study the relationship between demographics and phishing susceptibility. These surveys asked participants questions to determine their background and assess their knowledge about phishing, and then participants continued to complete a scenario in an effort to assess their behavioral susceptibility to phishing. [7]

The results of this study suggested that women are more likely than men to fall victim to a phishing attack. In addition, a younger population (18-25) was suggested to be the most susceptible to online deception. This specific age's vulnerability was a direct result of their increased likelihood to engage in risky behaviors when compared to an older population. [7]

Largely, this study provided great evidence on who to specifically target during a phishing experiment and also gave some insight into a specific personality trait that was identified as being more susceptible to a phishing attack. From this study it was concluded that a younger user population (18-25) should be tested during the execution of the experiment.

In an effort to explore in greater detail which personality traits are more susceptible, an additional study titled, A Personality Based Model for Determining Susceptibility to Phishing Attack, was reviewed and analyzed. This study was conducted by the University of Arkansas at Little Rock. It attempted to determine what qualities make some individuals more susceptible to phishing attacks than others. Phishing susceptibility is defined in this study "as the likelihood that a person will respond to a phishing attack." The overall goal of this study was to find specific traits that lead to this susceptibility. [8]

The research targeted specifically the "Big-Five" personality traits. These characteristics were identified as the main predictors for certain human behaviors, and the study used these traits to explore how humans specifically act when engaged in a phishing attack. The "Big Five" include: openness, conscientiousness, extraversion, agreeableness, and neuroticism. Keying in on agreeableness and extraversion, this research suggested that people exhibiting these traits are more likely to submit to a phishing attack. People who are more agreeable are more compliant and more likely to express trust in the unknown; while people who are extraverted feel the need to gain access into a social group and therefore are more willing to give up personal information for a fear of becoming unsocial. [8]

After reviewing this study, these two traits became this current experiment's focus. Both agreeableness and extraversion were hypothesized as being the two personality traits that will exhibit the most susceptibility to a phishing attack. This hypothesis would be tested after the conclusion of the central experiment to test which training methods were most effective in reducing phishing attacks. In order to correctly design such a study, an evaluation of training methods had to be explored.

ISBN 1-933510-96-X/$15.00

In the study, School of Phish: A real world Evaluation of Anti-Phishing Training, a training module title PhishGuru was analyzed to see if it actually reduced the number of people falling for phishing attacks over a period of time. PhishGuru is an embedded training system that teaches users to avoid falling for a phishing attack by administering a message to the user when they click on the spoofed URL in a simulated phishing attack. This type of training is commonly termed the direct feedback method. [9]

In this experimental study, a total of 515 participants where sent a total of 10 phishing emails over a span of 28 days. Additionally, a control group received the same number of emails over the same length of time but did not receive any feedback what so ever. It was found that participants that received direct feedback training were statistically more effective in deterring attacks over the course of 28 days. [9] This research was beneficial because it gave evidence that supported the use of immediate feedback training as an effective tool to quell successful phishing attempts by attackers.

It was concluded from this research study that immediate feedback is a form of training that should be tested during the implementation of the phishing experiment for this current study. A warning message that educates a user on their susceptibility to an attack can prove to be a useful measure to prevent the success rate of phishing attacks. It can then be compared to the more traditional training modules that are provided as an education resource for users to complete on their own.

In fact, this type of research was conducted in a research study performed by Carnegie Mellon University. In Getting Users to Pay Attention to Anti-Phishing Education: Evaluation of Retention and Transfer, a team of researchers from Carnegie Mellon University compared the two types of educational materials designed to teach users not to fall for phishing attacks. One was the direct feedback training (embedded) described in the previous study which is comparable to PhishGuru, and the second method of training was non-embedded training where the module was sent via email for the user to complete at their leisure. The research question posed in this study was whether or not direct feedback training is capable of allowing a user to retain and transfer their knowledge to stop successful phishing attempts when compared to a non-embedded training module. [9]

The results showed that the participants who were exposed to the embedded direct feedback training spent more time reading and acquiring knowledge, which increases retention, than those that were exposed to non-embedded training. Moreover, those exposed to embedded training preformed significantly better than those in non-embedded training at identifying phishing emails. This research produced very noteworthy findings and will be applied to this current study to test the current annual phishing awareness training given to members of the institution.

Our efforts in this study will serve to validate that an embedded training module (feedback) will be more effective as the present annual phishing awareness training (non-embedded) in improving user vigilance towards phishing attacks. Also as previously stated, it is hypothesized that the personality traits of agreeableness and extraversion are the two traits that will exhibit the most susceptibility to a phishing attack. To explicitly test these hypotheses an experiment was designed and implemented using the knowledge gained from the literature review to conduct controlled phishing attacks against an 18-25 year old population using e-mails that incorporated the malicious techniques utilized by attackers.

The phishing emails utilized in this experiment were created on www.phishme.com. This website allows its users to send phishing attacks comprised of a phishing email, a landing site, and a educational forum to a desired population. The evolutionary prototyping model was used to design the various phishing emails. This method uses a trial and error approach where semi functioning prototypes are modified over the course of the study. [10] A phishing email prototype can be seen in Appendix A of this report. The phishing emails were created and sent out to all of the group members. This occurred before any phishing email reached the target population. The email was tested to ensure that all components were operational.

III. METHODOLOGY

The intent for this study was to assess the effectiveness of formal phishing training over the increased awareness from simply running a phishing exercise multiple times. To meet this goal, the following experiment parameters were established:

1. The email to be sent to all subjects would include an embedded URL that when clicked takes users to a web site where they are asked to enter sensitive information (their network credentials). In all cases the email 'bait' was some sort of free or discounted service that would be appealing to the target population and it leveraged knowledge of the organization (spear phishing).

2. The test population was broken down into three groups.

- Group 1: received the phishing email, however after the user entered data into the website and clicked submit, the page returned a server error and no additional information was provided to the user.

- Group 2: received the phishing email, after the user entered data into the website and clicked submit, the page returned a notice that they fell victim to a phishing attack and provided details as to what the user should have identified in the email.
- Group 3: received the phishing email, after the user entered data into the website and clicked submit, the page returned a notice that they fell victim to a phishing attack and directed the user to take the institutions phishing awareness training..

3. The target population was 892 students selected at random. Specifically, there were 287 members in group one, 298 in group two, and 307 in group three. The age of the participants ranged from 18-26. All participants within this experiment were treated in accordance with the American Psychological Association rules and regulations.

4. Emails to the students would be sent in a manner to minimize students alerting others of the phishing exercise. To accomplish this the emails would be send to all students in the experiment groups as closely together as feasible.

5. Emails would be sent from a third party service provider outside the institutions boundary. The service selected was phishme.com [11].

6. The experiment would consist of two emails. The first phase would establish the baseline awareness level. Participants would have all taken the required institutional phishing awareness training (in September) however no additional training would have been provided. The first email was sent out following the Thanksgiving holiday break on 30 November. The second email was sent out on 9 December. The goal would be to determine if the response rate on the second phishing email was different among the groups. The duration between the emails is short enough to ensure the knowledge gained through the training.

The emails themselves were constructed to provide several clues designed to alert end users. The email is shown below in figure 1. The institutional specific information and personal information has been redacted.

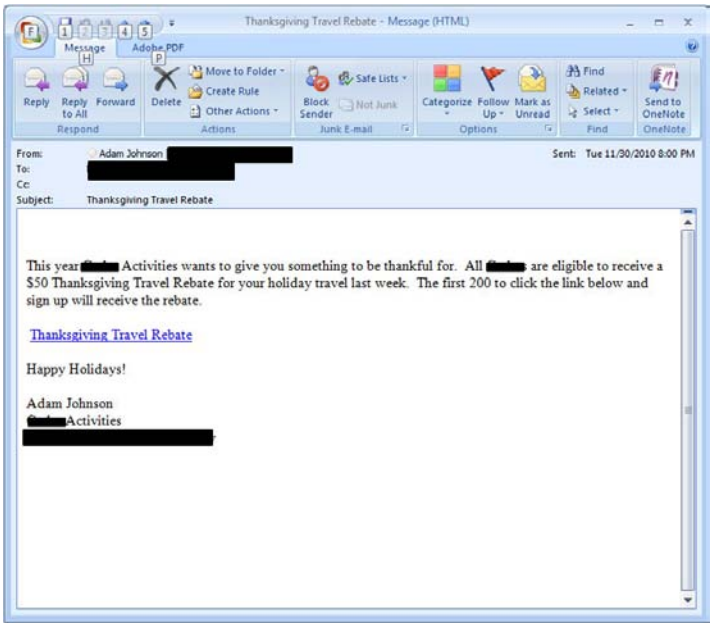

Figure 1: First phishing email sent

The following items details are items that are part of the annual training our users receive and should have alerted them.

1. Source email address: The email is from a person; however the source address is an institutional address. Further the user does not exist in our global list (accessible by all authorized users). The figure has the email address redacted. Here is what the field looked like (modified for blinded version):

> Adam Johnson [studentactivities@school.com]

2. The URL in the email: In figure 1, the URL is shown in the middle using html formatting for email. This is not how email is delivered by default to our users. Instead the 'presentation text' is listed as well as the actual URL string. Below is what the user would have seen (modified for blinded version)

> Thanksgiving Travel Rebate
> <http://studentactivities.travel-refund.com/775892/?login_id=0c573d44-fce7-11df-a9c7-00163e194312&owner=1>

3. The email urgency: While not always a valid sign of phishing email, when an urgent email is received from a source outside the organization, it is highly suspect.

4. Finally, if they looking in the full mail headers; they would have seen (modified for blinded version) that the originating email server is highly suspect:

Received: from LOCAL MAIL SERVER (1xx.xx.xx.xxx) by LOCAL MAIL SERVER (1xx.xx.xx.xxx) with Microsoft SMTP Server (TLS) id 8.2.254.0; Tue, 30 Nov 2010 20:05:21 -0500 Received: from mail.phishme.managedmachine.com (75.127.71.200) by

LOCAL MAIL SERVER (1xx.xx.xx.xxx) with Microsoft SMTP Server id 8.2.176.0; Tue, 30 Nov 2010 20:02:44 -0500 Received: from localhost.localdomain (localhost.localdomain [127.0.0.1]) by mail.phishme.managedmachine.com (Postfix) with ESMTP id 5DC6B10A43 for <USERNAME@SCHOOL.edu>; Wed, 1 Dec 2010 01:05:14 +0000 (UTC)
Date: Wed, 1 Dec 2010 01:05:11 +0000
From: Adam Johnson <studentactivities@SCHOOL.edu>
To: < USERNAME@SCHOOL.edu >
Subject: Thanksgiving Travel Rebate
MIME-Version: 1.0
Content-Type: text/html; charset="utf-8"
X-Priority: 3
X-Pmsid: 775892d4-fce0-11df-97b7-00163e4638cc
Message-ID: <20101201010514.5DC6B10A43@mail.phishme.managedmachine.com>
Return-Path: studentactivities@usma.edu
X-MS-Exchange-Organization-Antispam-Report: IPOnAllowList
X-MS-Exchange-Organization-SCL: -1

IV. Results

The main findings of the experiment are displayed in Figure 2. This figure displays the response rate to the first and second phishing emails. The graph displays the results for each experiment group; those who received no feedback, those who received feedback, and those that were required to retake the annual phishing awareness training (if they 'failed' the first phishing email. Each group is subsequently broken down into the first and second phishing email. The results for each set show the percentage of users who clicked the URL and then entered sensitive data on the landing web page. As can seen, the decrease in susceptibility between the first and second phishing emails was significant.

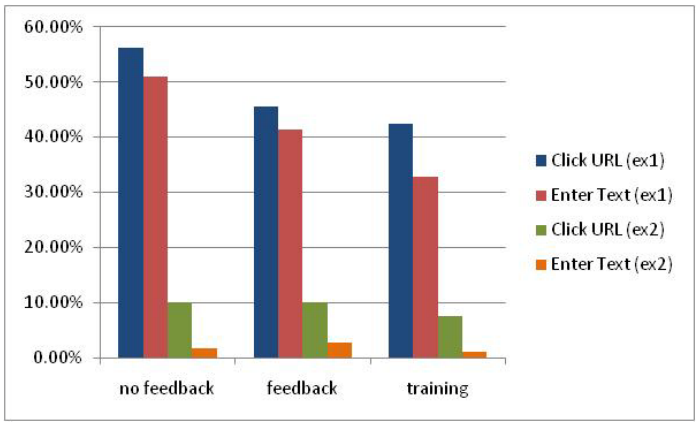

Figure 2: Phishing response results

While not part of the intended metrics to assess; the response timeline was very interesting and worth some discussion. The emails were sent out to all 892 participants between 8:00PM and 8:30PM. This was done to minimize the cross talk effect amongst the students. The students that responded to the phishing emails did so almost entirely (99%) within the first 4 hours.

V. Conclusions

The results indicate that even individuals who receive annual training are susceptible to phishing emails to an unacceptable degree. The results further indicate that given that the subjects all completed annual phishing awareness training in September, that after two months the knowledge retention is insufficient. The exercise results indicate that the phishing exercises provide substantial increased awareness. Given two weeks apart the susceptibility decreased from approximately 40% to approximately 5%.

This study does not attempt to further analyze the causes for the decrease; that analysis will be completed in future work. There are several factors that must be assessed in the future to arrive at firmer conclusions.

First, the annual training effectiveness may have been effective to the same degree as the phishing exercise at the two week mark. We will continue to employ the phishing exercise and assess the response rate for those subjects who were required to re-take the annual training. We will further evaluate this factor by including a new experiment group (a new student group not previously part of the study) in the fall to assess the short term effectiveness of the annual training.

Second, the student body at our institution is very well connected and news travels very quickly. It is possible that the students alerted each other during the second test, polluting the results. We believe two factors mitigate this concern. First, the emails were sent to the entire group in a very small time window. Results indicate (from both emails) that 80% of the students responded within the first hour and the remaining within 4 hours. Second, even if students alerted each other – that is a positive result. In fact, not only does it reduce the immediate effects of phishing, but it re-enforces communication channels for non-test security threats. Specifically, we know of two participants, one in feedback group and one in the training group, attempted to warn their fellow students of approximately 150 students that the email they received looked like a phishing email and that no one should click on it. This behavior shows that the current level of security training within this user population has created some amount of vigilance toward phishing attacks.

Third, the sample size should be increased. In February/March we anticipate increasing the size of the test population to 3000 students.

Finally, we did not discuss the institutional issues around a phishing exercise in this paper. These considerations are important and are discussed in detail in [6]. Several participants publicly complained about the study. One participant in particular posted a picture of the phishing email on Facebook with an angry status update. This occurred despite the fact that the on the feedback web page, the participants were specifically told to refrain from letting others know that the email was a phishing attack. The intent behind that feedback was to ensure that other participants' awareness and vigilance toward phishing emails was not artificially raised.

We believe that the use of exercises in phishing awareness and in the greater information assurance awareness field have a positive effect on awareness. In future studies we intend on providing a more detailed analysis of potential contributing factors, a larger sample size with more iterations of the experiments, and an analysis of personal characteristics (age, gender, academic performance, and personality traits, for example NEO results). Additionally, we are evaluating a survey questionnaire to better document attitudes and other contributing factors.

VI. References

[1] Hicks, D., "*Phishing and Pharming: Helping Consumers Avoid Internet Fraud*," Public & Community Affairs, Federal Reserve Bank Training, 2010

[2] Jagatic, T. N., Johnson, N. A., Jakobsson, M., & Menczer, F. (2007). *Social Phishing*. Communications of the ACM, October 2007, 94-100.

[3] Phifer, L., "*Top Ten Phishing Facts*," Retrieved 29 November 2010, from eSecurity Planet: http://www.esecurityplanet.com/views/article.php/3875866/Top-Ten-Phishing-Facts.htm

[4] *Phishing Activity Trends Report. Anti-Phishing Working Group*, 10 February, 2011, retrieved from http://www.antiphishing.org/reports/apwg_report_q2_2010.pdf

[5] Dhamija, R., Tygar, J. D., & Hearst, M., "*Why Phishing Works,"* Proceedings of the SIGCHI conference on Human Factors in computing systems, pp. 581-590, 2006pp. Montreal, Canada.

[6] Dodge, R., Carver, C., and Ferguson, A., "Phishing for user security awareness," Computers and Security, Computers & Security, Volume 26, pages 73-80, February, 2007

[7] Sheng, S., Holbrook, M., & Kumaraguru, P., "*Who falls for a phish? A demographic analysis of phishing susceptibility and effectiveness of interventions*," 28th International Conference on Human Factos in Computing Systems. 2010, Atlanta, USA.

[8] Parrish, J. L., Bailey, J. L., & Courtney, J. F. , "*A Personality Based Model for Determining Susceptibility to Phishing Attacks*," pg 285-296, 2009, Southwest Decision Sciences Institute proceedings, Little Rock Arkansas

[9] Kumaraguru, P., Rhee, Y., Sheng, S., Hasan, S., Acquisti, A., Cranor, L. F., et al., "*Getting users to pay attention to anti-phishing education: evaluation of retention and transfer*," Anti-Phishing Working Groups 2nd Annual eCrime Researchers Summit, 2007.

[10] Heim, S. G., "*The Resonant Interface*," Pearson Education-Addison Wesley, 2008.

[11] www.pishme.com, accessed 15 February 2011