Unannounced Phishing Exercises and Targeted Training: Results and Lessons Learned

Daniel Bliton, Aimee Norwood, Sean Palka
Booz Allen Hamilton
McLean, VA
Bliton_Daniel@bah.com,
Norwood_Aimee@bah.com,
Palka_Sean@bah.com

ABSTRACT

With cyber security on the minds of many large and small organizations, phishing, a type of social engineering attack, poses an increasingly common threat to every organization's information technology (IT) enterprise and therefore to the organization's ability to perform successfully. Phishing attacks target the weakest link in the information security chain—the individual end users. For example, one phishing attack tried to defraud users into resetting their DoD Common Access Card (CAC) Personal Identification Numbers (PINs) via an external website. Some organizations have attempted to protect themselves by engaging their workforce in phishing attack exercises. Frequently, these training exercises are announced beforehand and do not include remediation—these two factors may impede any organization's ability to improve user behavior and to attain required IT security outcomes in an actual work environment.

This paper describes the methodology, results, and lessons learned from a blind study on the effectiveness of pre-incident training to improve performance against phishing attacks (N = 467). During the study, each of the five treatment and control groups received a different type of training before exposure to an unannounced phishing attack. The study then measured the effectiveness of combining sustained, unannounced, phishing exercises with remedial training. The results show that an approach employing sustained training and exercises can significantly improve learning transfer and on-the-job performance as opposed to traditional training approaches, which had no positive impact on performance. Additionally, the response metrics and feedback from the treatment groups offer key insights into how phishing awareness training and exercises should be implemented for a workforce. Also, a real-world phishing attack during the study provided supporting evidence to the efficacy of the sustained training and exercises approach.

ABOUT THE AUTHORS

Sean Palka is a senior penetration tester with Booz Allen Hamilton. He has performed social engineering and penetration testing against networks and applications for a wide range of government and commercial clients, including several banks. He has created tools for executing complex phishing attacks and tracking of a wide range of response metrics. Sean recently won a "black badge" at DEFCON, the world's longest running and largest underground hacking conference, and is currently working on his doctoral dissertation in Computer Science at George Mason University.

Aimee Norwood is an immersive learning strategist with Booz Allen Hamilton and has been engaged in learning and technology for over 20 years. With a passion for helping people and organizations improve performance, she partners with clients to architect and implement holistic learning solutions.

Dan Bliton is a learning innovator and strategist with Booz Allen Hamilton. He is a passionate learner and has been designing computer-based and Web-based training solutions for over 24 years. Dan is currently deeply engaged with research on effective learning transfer. He is also the creator of the documentary film "The Machinima Primer" which showcased the use of video game technologies for storytelling and the rapid production of movies.

Unannounced Phishing Exercises and Targeted Training: Results and Lessons Learned

Daniel Bliton, Aimee Norwood, Sean Palka
Booz Allen Hamilton
McLean, VA
Bliton_Daniel@bah.com,
Norwood_Aimee@bah.com,
Palka_Sean@bah.com

INTRODUCTION

With cyber security on the minds of many large and small organizations, phishing, a type of social engineering attack, poses an increasingly common threat to every organization's information technology (IT) enterprise and therefore to the organization's ability to perform successfully. Phishing attacks target the weakest link in the information security chain—the individual end users. Some organizations have attempted to protect themselves by engaging their workforce in phishing attack exercises. Frequently, these training exercises are announced beforehand and do not include remediation—these two factors may impede the organization's ability to improve user behavior and more importantly, attain desired IT security outcomes, in the real work environment.

This paper will help answer the research question of the relative effectiveness of mandatory phishing awareness training alone compared to the combination of exercises and targeted training in the improvement of job behaviors (i.e., responding properly to phishing attacks) and the desired IT security outcomes. This paper describes the methodology, results, and lessons learned from a blind study on the effectiveness of pre-incident training (e.g. mandatory training) to improve performance against phishing attacks and the learning transfer effectiveness of combining unannounced phishing exercises with remedial training. Phishing and social engineering exercises are a subset of penetration testing.

Penetration Testing

When deploying a network or software component, a known best practice is to obtain a sufficient degree of certification and accreditation (C&A). A critical component in the development of a secure network includes penetration testing. Penetration testing simulates the types of attacks that an adversary might employ, in order to exploit vulnerabilities and leverage access into an IT system or network. Unfortunately, organizations usually only perform an in-house or checklist-based evaluation of system components to evaluate whether the configurations are secure relative to an organization's approved IT security policies. This approach often overlooks critical vulnerabilities and configuration errors that are well-known to adversaries, and, as a result, often requires reevaluating the IT system at much greater cost after it is compromised.

By employing penetration testing, organizations can analyze their systems from the perspective of an adversary, whereby the testers attempt to identify and exploit vulnerabilities in order to leverage access into the system. By simulating the attacker, a penetration test can provide a more holistic view of the vulnerabilities and their impacts on the system, and ideally describes the security posture of the system using empirical examples of how vulnerabilities could be exploited.

Often, various types of attacks, such as phishing or social engineering attacks may be excluded from testing, due to their direct interaction with the mainstream workforce (where more than the system administrators are involved). And, management may place restrictions on the penetration test itself in order not to disrupt on-going operations or the work environment (Klevinsky, Laliberte, & Gupta, 2002).

Phishing

Phishing and social engineering exercises are a subset of penetration testing. In a phishing attack, users receive an unsolicited e-mail that tries to entice them to perform an action. They may be asked to click on a link, open an attachment, or send information to the attacker by replying to the message. The e-mail may often appear to come from a legitimate source. Typically, the goal is to acquire information (e.g., passwords, account numbers, sensitive data) that the attacker can use for future attacks, or to direct users to malicious sites that infect the computer. These attacks often focus on enticing the human element to take action or respond, rather than on the technical components of a system (e.g., gaining access to a web server or a router). Users are difficult subjects to address from a security standpoint because humans are far less predictable than technical components and possess emotions and feelings.

The Role of Training

Because social engineering and phishing target the human element, training plays a critical role in the defense of security breach for any organization. Efficacy of the training is of the utmost importance. The human side of a network implementation is generally to provide awareness training on the risks of lax vigilance; whereas, software and other components undergo staged development and end-to-end testing processes before deployment to a production computing environment. Determining whether a given software configuration functions properly is fairly straightforward, but assessing the effectiveness of user awareness training is more difficult.

With social engineering exercises, user awareness training is tested using simulated attacks that mimic the types of methods an adversary might employ (e.g., phishing e-mails). As with network and software penetration testing approaches, a typical social engineering exercise will attempt to leverage information gathered to gain additional access to personal data using social engineering tactics. The impact of this type of threat becomes clear when provided hard data. Many system owners underestimate the number of users who are susceptible to a phishing e-mail attack, but when 40% to 50% of targeted users respond to a phishing e-mail, the owners better understand how user vulnerability represents a clear and present risk to the enterprise.

With network or software components there is usually a mitigation process for a given vulnerability. In an effective social engineering attack, the user is often unaware that they have been targeted or exploited. As a result, it is difficult to train a user for a situation in which they might not recognize there is an issue and to provide meaningful feedback when they are not even aware that they did something wrong. Determining an effective way to mitigate the effects of a social engineering attack is a primary focus of this paper.

RESEARCH AND PRACTICAL GOALS

A mitigation strategy for a phishing attack may consist of one or a combination of approaches:

- Communicate the dangers of phishing via organization-wide *strategic communications* (e.g., alerts, bulletins, e-mails)
- Create increased awareness of the dangers of phishing via *mandatory training* (e.g., training provided under normal and expected training conditions)
- Assess user behavior via an announced or unannounced *social engineering training exercise* (with or without immediate feedback)

Research Questions

One of the key goals of the study was to determine the relative effectiveness of pre-incident phishing awareness training alone compared to the combination of exercises and targeted training in the improvement of job behaviors (i.e., not responding to phishing e-mails). The answer to the question about the effectiveness of pre-incident training (e.g. mandatory training) has been difficult to find in current literature/studies (Adams, 2010). Although, a small study at West Point showed that periodic launching of phishing awareness exercises alone should help minimize susceptibility to phishing attacks by users (Ferguson, 2005). Also, another study (28-days long) showed that the combination of exercises and awareness training minimized student's susceptibility to phishing attacks (Kumaraguru, et. al, 2009), but did not compare against the impact of traditional awareness training or awareness communications alone.

The study described in this paper took place over nine months and compares the results of:

- Awareness communications alone (e.g., awareness bulletin e-mails)
- Awareness communications combined with pre-incident training (e.g., traditional mandatory training)
- Unannounced phishing exercises combined with immediate feedback and remedial training.

Study Hypotheses

The following hypotheses were examined:

Hypothesis 1: Users who receive interactive phishing awareness training with examples of phishing e-mails will assign higher reaction (satisfaction) scores to the training than the users who receive training that does not contain interactive phishing examples (e.g., content copied from phishing awareness wiki pages).

Hypothesis 2: Users who receive phishing awareness information from bulletins and other non-training communications only, are more likely to click links in phishing exercise e-mails. In addition, they are less likely to submit the phishing exercise e-mails to the Critical Incident Response Team (CIRT) than users who receive bulletins and phishing awareness training in a typical training environment (e.g., announced mandatory training).

Hypothesis 3: The number of incorrect actions (e.g., clicking on suspicious links in simulated phishing attacks) will decrease and the number of correct actions (e.g., reporting simulated phishing attacks) will increase with exposure to unannounced phishing exercises combined with immediate remedial training combined.

Hypothesis 4: Users who learn how to respond appropriately to phishing exercise e-mails will be able to transfer that knowledge to actual phishing attacks.

METHOD

Participants

Four hundred and sixty seven users volunteered to participate in a study on the effectiveness of different instructional approaches used in a pilot lesson on the overview of cyberspace in exchange for the chance to win a small incentive (i.e., gift cards or water bottles). These users were members of either a cyber security community of practice or a learning and performance community of practice. The volunteer information was randomly sampled to verify that the volunteers represented multiple geographic United States locations across multiple functional expertise teams. Most of the volunteers resided in the Washington D.C. metropolitan area. In this blind study, the volunteers were unaware that:

- The actual focus of the study was phishing awareness training and exercises (instead of the stated focus on the effectiveness of a cyber lesson)
- There were different training lessons presented to different groups
- They would be exposed to unannounced phishing e-mails.

The participants were divided into a total of five groups: a control group, a secondary control group, two treatment (experimental) groups, and additional baseline/control group for anticipated non-respondents from the treatment groups.

- **Control Group 1** received general information and a quiz on cyberspace and cyber security (not related to phishing awareness). They also completed an end-of-lesson reaction survey on their learning experience with the lesson materials (the goal of the study, from the participant's point of view).
- **Control Group 2 (Bulletin Group)** received the same general information on cyberspace as Control 1, but the end-of-lesson reaction survey also included questions about the Phishing Awareness bulletin. During the end-of-lesson evaluation participants were also provided an opportunity to review the bulletin. This group was tracked as a separate control group in the event that simulating recall of the awareness bulletin impacted the response to phishing e-mails.
- **Wiki Training Experimental Group** received basic phishing awareness content that had been ported from an internal phishing information wiki site.
- **Interactive Training Experimental Group** received phishing awareness content and interactive suspicious e-mail item identification activities using proprietary training software.
- **Non-Responsive Control Group** members were initially assigned to one of the control or experimental groups above, but did not complete their assigned lesson and evaluation. To our best knowledge, participants in this group had no prior documented training within the context of the study and could be treated as a separate control group.

Materials

Materials created and applied to the study included:

a) **Phishing awareness e-mail bulletins** provided to all participants
b) **E-Learning lessons** for the control and experimental/treatment groups
c) **End-of-lesson reaction surveys** to assess learner reactions to all the lessons (supporting a Kirkpatrick Level 1 evaluation)
d) **Post-test questions** for the experimental groups to assess level of knowledge of the phishing awareness content (supporting a Kirkpatrick Level 2 evaluation)

e) **Phishing exercise e-mails** sent unannounced to all participants. Participant's actions with the e-mails were tracked to help determine behavior in the work environment (supporting a Kirkpatrick Level 3 evaluation)

f) **Custom phishing awareness training** specifically targeting the phishing exercise e-mails for those participants that responded to the e-mails

a) Phishing Awareness E-Mail Bulletins

To help all participants respond appropriately to a phishing attack and to reduce unnecessary Help Desk calls during the study, an awareness bulletin was created and sent to all participants at the start of the study. Additional bulletins were created and sent to all participants during the study in response to ongoing IT requirements and actual phishing attacks. The bulletins were sent by a member of the senior leadership team.

b) E-Learning Lessons

Four different lessons were developed to support the five different participant groups (there was not a separate lesson for the Non-Responsive Control Group)**.**

1. Lesson with general cyberspace information (not related to phishing awareness) and quiz for Control Group 1 (**see Figure 1**). Content comprised of eight pages of static content (text and images) and five quiz questions published in a traditional e-learning format.
2. Lesson with the same general information on cyberspace as Control Group 1 for Control (Bulletin) Group 2. Content comprised of eight pages of static content (text and images) and five multiple-choice questions that used a mini-game interaction (**see Figure 2**) instead of the quiz format used for Control Group 1. The use of mini-games was thought to impact the user reaction ratings for the lesson, but not the responses to the phishing exercises.
3. Lesson with phishing awareness content that was copied from an internal phishing information wiki site for the Wiki Training Group. Content was copied from the wiki site, minimally edited, and published in a traditional e-learning format (**see Figure 3**). Content comprised of nine pages of static content (text and images). This content was followed by a post-test and then the reaction survey.
4. Lesson with phishing awareness content and suspicious item identification activities using proprietary training software for the Interactive Training Group. Content comprised of one traditional page with five mouse "rollovers" for additional information and one interactive identification activity page with three example phishing e-mails (**see Figure 4**). This content was followed by a post-test and then the reaction survey.

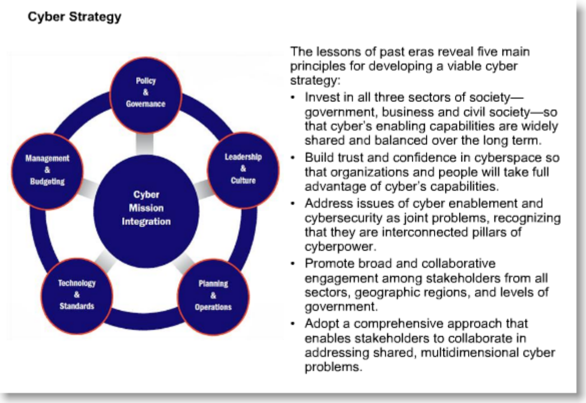

The lessons of past eras reveal five main principles for developing a viable cyber strategy:

- Invest in all three sectors of society—government, business and civil society—so that cyber's enabling capabilities are widely shared and balanced over the long term.
- Build trust and confidence in cyberspace so that organizations and people will take full advantage of cyber's capabilities.
- Address issues of cyber enablement and cybersecurity as joint problems, recognizing that they are interconnected pillars of cyberpower.
- Promote broad and collaborative engagement among stakeholders from all sectors, geographic regions, and levels of government.
- Adopt a comprehensive approach that enables stakeholders to collaborate in addressing shared, multidimensional cyber problems.

Figure 1. Lesson content for Control Group 1 did not cover how to respond to phishing attacks.

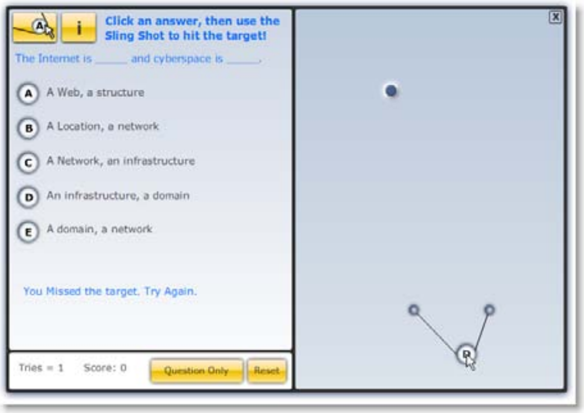

Click an answer, then use the Sling Shot to hit the target!

The Internet is _____ and cyberspace is _____.

A A Web, a structure

B A Location, a network

C A Network, an infrastructure

D An infrastructure, a domain

E A domain, a network

You Missed the target. Try Again.

Figure 2. Lesson content with mini-games for Control (Bulletin) Group 2

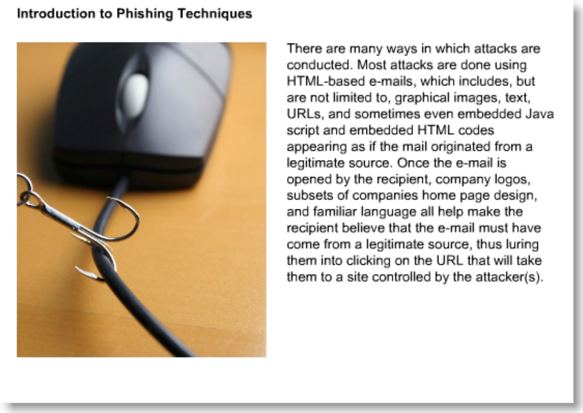

Introduction to Phishing Techniques

There are many ways in which attacks are conducted. Most attacks are done using HTML-based e-mails, which includes, but are not limited to, graphical images, text, URLs, and sometimes even embedded Java script and embedded HTML codes appearing as if the mail originated from a legitimate source. Once the e-mail is opened by the recipient, company logos, subsets of companies home page design, and familiar language all help make the recipient believe that the e-mail must have come from a legitimate source, thus luring them into clicking on the URL that will take them to a site controlled by the attacker(s).

Figure 3. Static phishing awareness content for Wiki Group

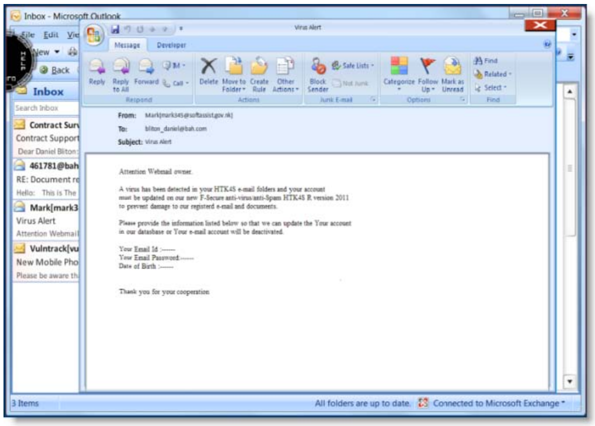

From: Mark[mark345@softassist.gov.nk]
To: bliton_daniel@bah.com
Subject: Virus Alert

Attention Webmail owner.

A virus has been detected in your HTK4S e-mail folders and your account must be updated on our new F-Secure anti-virus/anti-Spam HTK4S R version 2011 to prevent damage to our registerd e-mail and documents.

Please provide the information listed below so that we can update the Your account in our datasbase or Your e-mail account will be deactivated.

Your Email Id :------
Your Email Password:------
Date of Birth :------

Thank you for your cooperation

Figure 4. Suspicious item identification activity for Interactive Training Group

c) End-of-Lesson Reaction Surveys

An anonymous reaction survey was created for each lesson to determine user satisfaction with the lesson. A separate non-anonymous survey was developed to capture the e-mail addresses of all participants who completed a lesson and reaction survey for the purposes of awarding the incentives.

d) Post-Test Questions

A post-test with 11 questions was created to help determine the retention of the phishing awareness content covered in the Wiki Training Group Lesson and the Interactive Training Group lesson.

e) Phishing Exercise E-Mails

Phishing exercise e-mails were created for testing the participants' responses in a non-training environment. From the participants' perspective, these e-mails should not have been viewed as messages from their organization or co-workers, since the e-mails originated from outside the organization (i.e., external domains) and were an unannounced part of the study. The study also collected data from an actual external phishing attack e-mail that was sent to some of the study participants and a number of other users.

The e-mails included in our analysis were as follows:

- E-Mail #1: Scheduled Server Migration
- E-Mail #2: Upgrading to Microsoft Exchange 2011 (external phishing attack e-mail)
- E-Mail #3: New Mobile Phone Tracking Risk
- E-Mail #4: Upcoming Black Hat/DEFCON Conferences.

All the e-mails were created in a similar design/format:

- Fairly well-structured request for action (e.g., click a link) with some grammatical errors
- Referenced generic content (no content specific to the organization)
- Minor attempts at obfuscation (e.g., click "here") with the "here" link being a malicious link
- Did not use the participant's name in the message
- Did not include attachments

E-Mail #1 (exercise e-mail) and E-Mail #2 (external phishing attack e-mail) were very similar in content. Chronologically, the external phishing attack e-mail occurred between the 1st and 2nd of our study e-mails, and (unintentionally) served to illuminate the realism of the content employed in our exercise e-mails.

f) Custom Phishing Awareness Training

Custom phishing awareness training was created for each phishing exercise e-mail. If participants responded to a phishing exercise, they received the interactive phishing training (same as provided to the Interactive Training Group) that modified to include specific instruction and feedback on the phishing e-mail used in the exercise.

Procedure

The exercise component of the phishing study was managed by a web-based tool. For each exercise, participants received simulated phishing e-mails developed using the tool. The administration tool was used to coordinate the development, content management, and response analytics for the phishing exercises.

One of the key aspects of the exercises is determining response statistics. To accurately determine how a participant is responding to the exercise, the tool employs a tagging scheme so that each e-mail is uniquely marked for tracking purposes. When a participant responds to the phishing e-mail, the unique token in the request is interpreted by the tool to track relevant information (e.g., geographic location, forwarded e-mails, potential users who are responding).

A key learning objective of the study was to teach participants to not click suspicious links in e-mails. To truly simulate a phishing attack, the attack team registered external domains that acted as capture agents. The tool also managed and delivered the custom training content. When responses were received by the capture agent, updates were passed to the tool and the participants were directed to targeted remedial training. The tool tracked when a participant started and completed training relative to the captured responses, and tracked (over time) how often a user had been targeted and to what attacks they responded.

Phase 1—Phishing Awareness Bulletin

All participants were e-mailed an official bulletin (using the organization letterhead sent by a senior leader) that provided basic information on phishing attacks and the actions to take if users received a phishing e-mail.

Phase 2—Pre-incident Training

Three months after the awareness bulletin was sent, all participants were e-mailed directions via a Constant Contact® message (the message views and link clicks were automatically tracked). Participants were directed to complete the pilot training lesson via a provided link, answer the post-test content questions, complete the reaction survey (anonymous), and then submit their name for the chance to win a nominal incentive. A total of 281 participants (out of the 467 volunteers) completed the pre-incident training lessons (e.g. typical mandatory training). There was a decrease in participants during the study due to requests to be removed from the study. Group composition was:

1. Control Group 1 (43 users)
2. Control Group 2 with Bulletin (59 users)
3. Wiki Content Group (78 users)
4. Interactive Training Group (101 users)
5. Non-responsive Group (186 users who did not complete enough of the process to be tracked as completing the training and survey).

Phase 3—First Exercise E-mail

One month later, all five groups were sent E-Mail #1 with the topic area of “Scheduled Server Migration.” Participants who clicked any link in the exercise e-mail received immediate feedback in a new window and were directed to take the phishing awareness training.

Clicking the training button displayed a new screen for participants to complete to log into the training. Any participant who responded to an exercise e-mail, but didn’t complete the short training, received a follow-up reminder e-mail to complete the training.

Participants were tracked for clicking links, accessing training, completing training, and total training time.

Phase 4—Additional Awareness Bulletins

Over the next three months two more awareness bulletins were sent to all participants s. No tracking or evaluation was performed.

Phase 5—Unplanned External Phishing E-Mail Attack

Three and a half months after E-Mail #1 was sent, E-Mail #2 was received by 1,234 users, including 11 of the study participants. The e-mail contained the subject line “Upgrading to Microsoft Exchange 2011,” and originated from an external adversary. Tracking was used to determine how many non-study and study participants responded to the e-mail, as the original content was not developed by the authors and did not contain tagging elements.

Phase 6—Second Exercise E-mail

A few weeks later, all five groups were sent E-Mail #3 with the topic of “New Mobile Phone Tracking Risk.” Participants were treated and tracked using the same process as E-mail #1.

Phase 7—Third Exercise E-mail

A few weeks later, all five groups were sent E-Mail #4 with the topic of “Upcoming Black Hat/DEFCON Conferences.” Participants were treated and tracked using the same process as E-Mail #1.

RESULTS

Results from Awareness Bulletin and Pre-Incident Training (Phase 1 and Phase 2)

All participants received the e-mailed awareness bulletin on phishing, (except a few who joined the organization just after the release of the bulletin and were included in the study). Recall of the bulletin was later measured by survey questions to Control Group 2.

All participants received a “next steps” e-mail that directed them to complete a new “pilot” training lesson and the anonymous training evaluation via a provided link.

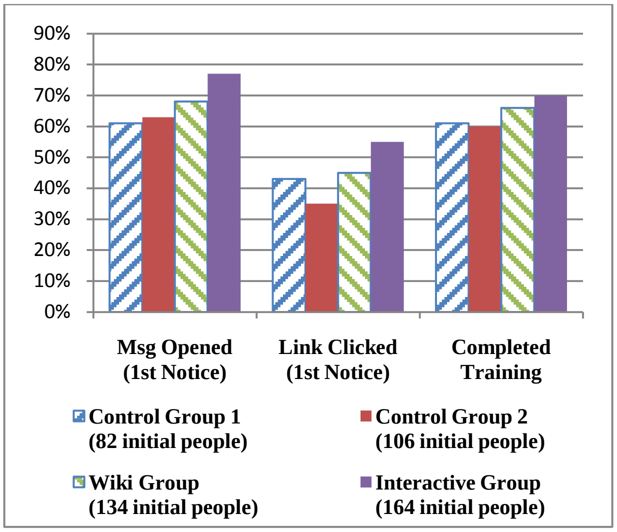

Figure 5. All initial groups responded similarly to a “next steps” e-mail, as measured by Opening Messages, Clicking Links, and Completing Training

The groups' initial actions with the "next steps" e-mail were analyzed to determine if the groups were balanced in their responses (e.g., did they open the message, click the link to training, and eventually complete the training). Although the two training intervention groups (Wiki and Interactive) were more active (e.g., opened and clicked links) with the e-mail (**see Figure 5**), there was no significant difference between the groups (chi-squared analysis, P >0.05).

The interactive phishing awareness training and the wiki content in traditional format received the highest (non-significant difference) average ratings (testing Hypothesis 1) from the training evaluations across the four groups (**see Table 1**). Note that Control Group 2, (the group who received control content with mini-games), provided lower ratings across almost all areas, but this should not have differentiated them from Control Group 1 in their reaction to phishing e-mails.

Table 1. Interactive Training and Wiki Training Received the Highest User Ratings

	Overall Rating (5 max)	**Delivery was engaging (5 max)**	**Able to recognize phishing (5 max)**	**Would recommend to others (Yes)**
Control 1 (N=50)	3.7	3.7	3.0	70.0%
Control 2 (N=64)	3.3	3.4	2.6	57.8%
Wiki (N=88)	3.7	3.6	4.1	85.2%
Interact (N=114)	3.8	4.1	3.9	85.1%

The post-training evaluation was also used to capture Control Group 2's (bulletin) recall of the initial phishing awareness bulletin that was provided to all participants. These participants (N= 64) also noted whether they re-reviewed the bulletin when provided the opportunity. A majority of the group either remembered the awareness bulletin (59.4%) and/or reviewed the bulletin again (6.3%). Since the bulletin contained information about what to do in case of a phishing attack, these participants should have been prepared for the first exercise e-mail (testing of Hypothesis 2).

The two experimental groups (the Wiki Training and the Interactive Training groups) completed post-tests with scores that indicated recall of the training materials (**see Table 2**). A high percentage of correct responses on what actions to take upon receipt of a phishing e-mail, should have resulted in appropriate action upon the receipt of the first exercise e-mail; that is, if learning transfer took place with the traditional (pre-incident) training approach (testing Hypothesis 2).

Table 2. Post-tests Indicate that Both Training Groups Knew What to Do with Suspicious E-Mails

	% correct response to a question on what to do with suspicious e-mails
Wiki Group (88 participants)	87.8% correct
Interactive Training Group (114 participants)	95.6% correct

The experimental groups provided their highest rating for the appropriateness of the amount of time it took to complete the training, which was on average less than 11 minutes.

Results for First Exercise E-mail (Phase 3)

Incorrect response rates to the First Exercise e-mail across all five groups are shown in **Figure 6**. Chi-squared analysis indicated no significant difference (P > 0.05) between the groups (pre-incident training did not impact behavior to an unannounced phishing e-mail).

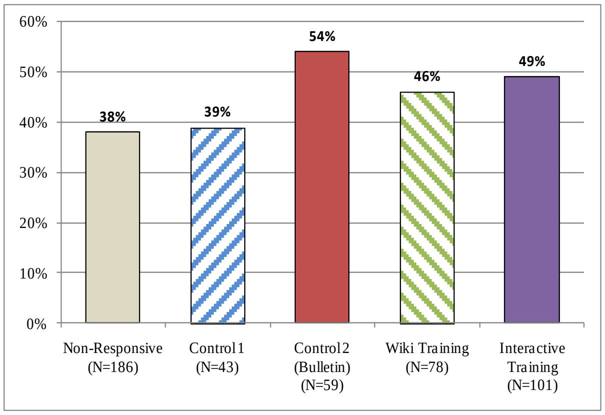

Figure 6. Incorrect actions for Exercise 1 show that the traditionally trained groups (Wiki and Interactive) were not significantly different than the control groups (Control 1 and Control 2)

An average of 45% of the participants clicked a suspicious link in the e-mail. Only two participants completed the desired action of forwarding the e-mail to CIRT.

Results for Additional Bulletins and Outside Phishing E-Mail Attack (Phase 4 and Phase 5)

Two additional awareness bulletins were e-mailed to all participants; no tracking or evaluation was performed. Then, an unanticipated outside phishing attack targeted 1,234 users, including 11 of the study participants with the previously described E-Mail #2 – "Upgrading to Microsoft Exchange 2011." The following results were captured:

- Targeted users: 1,234
- Targeted users from phishing study: 11 (1%)
- Users who clicked on the phishing links: 9 (0.7%)
- Study group users who clicked on the phishing links: 0 (0%)
- Users who reported the e-mail to CIRT (desired response): 14 (1.1%)
- Study group users who reported the e-mail to CIRT (desired response): 3 (27%)

The 14 people who reported the e-mail to CIRT provided additional information on how they knew what action to take. A summary of how they knew:

- Recalled bulletins: 7 (50%)
- Asked a co-worker or the Help Desk: 5 (36%)
- Had received training: 3 (21%)
- Was aware/guessed that a standard CIRT@domain.com should be in place: 2 (14%)

Results for Second and Third Exercise E-Mails (Phase 6 and Phase 7)

All five groups were sent the Second Exercise e-mail (E-Mail #3) and then the Third and final exercise e-mail (E-Mail #4).

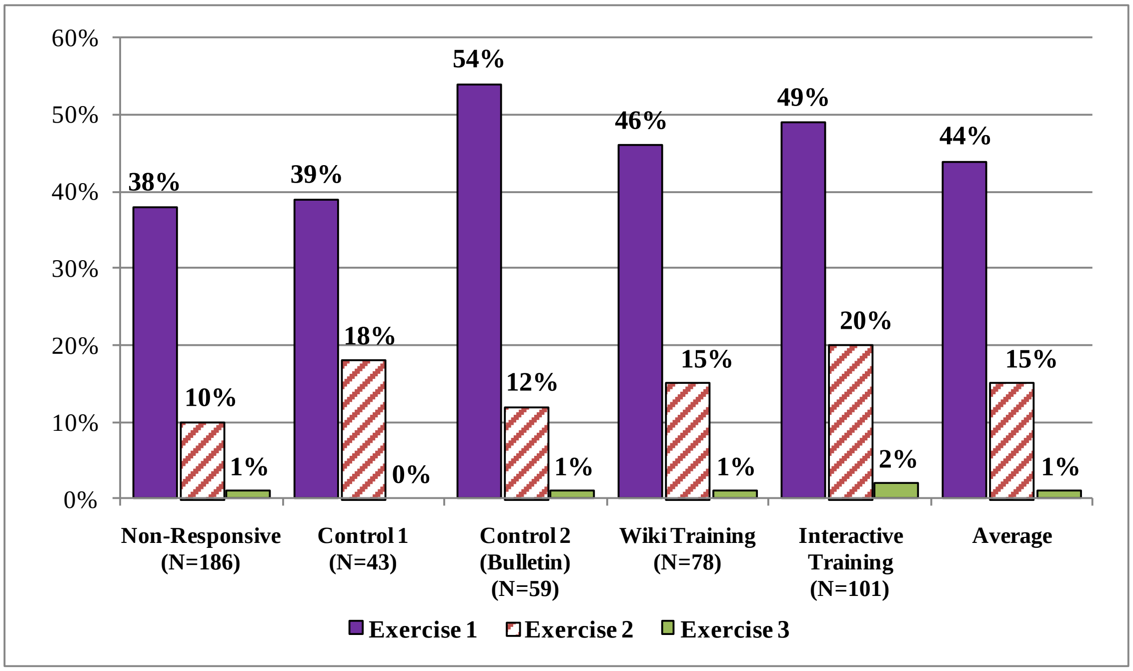

Figure 7. Incorrect response rates decreased significantly for every group for all three exercises

The incorrect response (clicking on a suspicious link) rate decreased from an average of 44.1% for the First Exercise to 1.4% for the Third Exercise (**see Figure 7**). This represented a normalized 96.8% reduction in incorrect response rates from the First Exercise the Third Exercise (**see Table 3**).

Table 3. 95% Confidence Intervals

		95% Confidence	
	Avg	**Low**	**High**
Exercise 1	44.1%	39.7%	48.6%
Exercise 2	14.5%	11.5%	18.2%
Exercise 3	1.4%	0.6%	3.1%

Analysis of confidence limits and chi-squared analysis revealed a statistically significant difference across all three exercises with P < 0.05. There was also a statistically significant difference (chi-squared analysis, P < 0.05) from the First Exercise to the Second Exercise and from the Second Exercise to the Third Exercise.

The desired/correct response (forwarding the phishing e-mail to CIRT) increased from a total of two for the First Exercise to a total of 23 for the Third Exercise. Due to the small number of responses, there was no statistical analysis performed, but there was an increase of correct responses for four out of the five groups (**see Figure 8**).

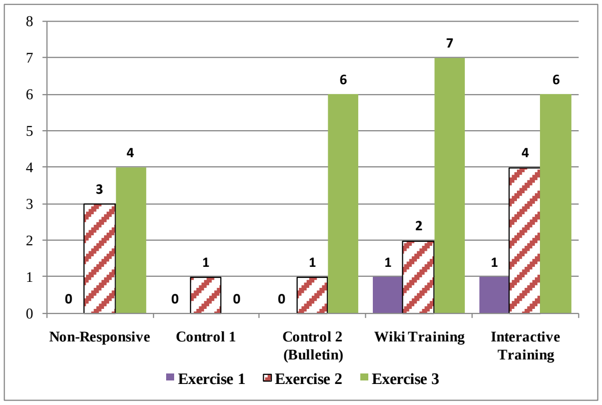

Figure 8. Correct responses increased for every group from the First Exercise to the Third Exercise

DISCUSSION

This study investigated the relative effectiveness of pre-incident training to improve performance against phishing attacks and the effectiveness of combining unannounced phishing exercises with remedial training.

The results of the study data show the following conclusions for the stated hypotheses:

1. **Hypothesis 1:** Not supported. Users assigned slightly higher reaction (satisfaction) scores to the two experimental lessons, but there was no clear distinction between the control and experimental groups.
2. **Hypothesis 2.** Not supported. Pre-incident phishing awareness training (provided to the experimental groups) had no impact on the participants' response to the phishing exercises as compared to the participants that did not receive phishing awareness training.
3. **Hypothesis 3.** Supported. A significant number of incorrect actions (clicking on suspicious links in simulated phishing attacks) decreased and a non-significant number of correct actions (reporting simulated phishing attacks) increased with exposure to three unannounced phishing exercises combined with feedback and immediate remedial training.
4. **Hypothesis 4.** Indicated. A non-significant number of participants that learned how to respond appropriately to phishing exercise e-mails responded appropriately to an actual phishing attack. This indicates that users exposed to unannounced phishing exercises combined with immediate remedial training were able to transfer what they learned during the phishing exercises to actual phishing attacks.

A surprising result of this study (at least for the instructional designers) was that traditional pre-incident training (e.g., mandatory training) had no significant positive impact on the user's response behavior to phishing e-mails. In this study, the two groups that received relevant phishing awareness training (Wiki and Interactive Training Groups) and then were exposed to a simulated phishing attack one month later, did no better than the control groups. In fact, only one control group (the one that received unrelated mini-games) did worse in Exercise 1 than the two treatment groups that received phishing awareness training.

A possible explanation for the lack of impact of either of the two types of phishing awareness training presented is that some groups might have been composed of an uneven mix of users with prior knowledge and experience with phishing e-mails and appropriate reactions. This is an unsupported explanation since no testing was performed to determine the user's existing level of phishing awareness at the start of the study. Also, phishing e-mails are essentially a marketing attempt to entice the user into action, but users are enticed by different types of marketing, which makes it difficult to adequately compare responses across the groups with only three "marketing" e-mails. Still, this study supports a healthy skepticism concerning the use of mandatory training separated from the context of the actual work environment, at least for any awareness training like phishing awareness.

It is relevant to highlight that in this study, providing training alone was ineffective, but providing training combined with unannounced exercises was very effective in changing the behavior of the users.

The study was primarily focused on the effectiveness of combining unannounced phishing exercises with remedial training (which was very similar to the training used in Phase 2). This exercise and training approach proved to be very effective with an average reduction of inappropriate actions by 42.7%. The study's exercise and training approach aligns closely with Robert Gagne's model that students will learn optimally if instructors carefully select and integrate the appropriate combination of nine events or strategies into their lesson plans (Gagne, 1965). Although Gagne noted that each of the nine events of instruction do not need to be present in every learning situation (Gagne, 1992), the phishing exercise and remedial training approach used all nine events. Extensive attention was paid to the first four events with the goal to place the learners in position of "learning at the point of realization." Learning at the point of realization refers to the state when users are open to learning because relevance, knowledge gaps, and immediate needs are identified in an engaging/unexpected and concrete fashion.

Research from Suzanne Hidi (Hidi and Baird, 1988) and Mark Sadoski (Sadoski, 2001) indicates that the unexpectedness/interestingness of the phishing exercises (and related inappropriate user actions) combined with the concreteness of remedial training may have greatly influenced the learners' reception of the learning intervention. Basically, getting caught by a phishing exercise e-mail (i.e., failing a realistic scenario) should gain the interest of the users and help them recognize the relevance of their need for improvement. Gaining interest and placing the user at the point of realization is related to attention, deeper processing, and learning transfer.

One quote from a learner that responded correctly to the external phishing attack highlights that learning at the point of realization may greatly influence the level of learning transfer.

> "*I learned about the CIRT team through the phishing training email sent out a couple months back. It really stuck with me, since I 'failed the test.'* "

Conclusion

The reality is that cyber security is a people problem first and a technology problem second. Phishing attacks are a threat to most organizations' ability to perform successfully and must be met with an ongoing awareness program, but any program that relies primarily on traditional pre-incident awareness training without ongoing exercises (for learner reinforcement) is suspect in its effectiveness.

It is clear that sustained, unannounced, phishing exercises with short and targeted remedial training are very effective in reducing the incorrect responses of the target audience. While the study methodology was focused on phishing, the same unannounced exercise-based e-mail approach with associated training can be applied to other types of cyber awareness challenges such as Personally Identifiable Information (PII) disclosure and Computer Use Policy training.

ACKNOWLEDGEMENTS

Thanks to Semion Bezrukov, Sharon Shin, Cary Harr, Gina LoPresti, Michael Parmentier, and Ron Tharp for their contributions to this paper.

REFERENCES

Adams, E. VA Healthcare System, VA Technology Assessment Program. (2010). *Mandatory training—a systematic review of research and trends in learning organizations*, Boston, MA: Office of Patient Care Services.

Ferguson, A. (2005). Fostering e-mail security awareness: The West Point Carronade. *Educause Quarterly, 28(1)*, 54-57.

Gagné, R. (1965). *The conditions of learning and theory of instruction.* Fort Worth, TX: Holt, Rinehart and Winston, Inc.

Gagné, R. M., Briggs, L. J., and Wager, W. W. (1992). *Principles of instructional design* (4th ed.). Fort Worth, TX.: Harcourt Brace Jovanovich.

Hidi, S., and Baird, W. H. (1988). Strategies for increasing text-based interest and students'recall of expository texts. *Reading Res. Q*. 23: 465–483.

Klevinsky, T. J., T, Laliberte, S., & Gupta, A. (2002). *Hack I.T.: security through penetration testing*. Indianapolis, IN: Pearson

Kumaraguru, P., Cranshaw, J., Acquisti, A., Cranor, L., Hong, J., Blair, M., & Pham, T. (2009). School of phish: a real-word evaluation of anti-phishing training. *Proceedings of the 5th symposium on usable privacy and security* Mountain View, CA:

Sadoski, M. (2001). Resolving the effects of concreteness on interest, comprehension and learning important ideas from text. *Educational Psychology Review*, 13(3), 263e281.