Journal of Cybersecurity Education, Research and Practice

Volume 2020 | Number 2 Article 5

January 2021

Gophish: Implementing a Real-World Phishing Exercise to Teach Social Engineering

Andy Luse
Oklahoma State University, andyluse@okstate.edu

Jim Burkman
Oklahoma State University, jim.burkman@okstate.edu

Gophish: Implementing a Real-World Phishing Exercise to Teach Social Engineering

Abstract

Social engineering is a large problem in our modern technological world, but while conceptually understood, it is harder to teach compared to traditional pen testing techniques. This research details a class project where students implemented a phishing exercise against real-world targets. Through cooperation with an external corporate partner, students learned the legal, technical, behavioral, analysis, and reporting aspects of social engineering. The outcome provided both usable data for a real-world corporation as well as valuable educational experience for the students.

Keywords

social engineering, phishing

Cover Page Footnote

Special thanks to Natasha Wyche and Adam Kleiner for their assistance with setting up this project.

INTRODUCTION

Social engineering is one of the biggest security issues facing both individuals and corporations today. Several recent cases have demonstrated the potential for loss due to social engineering attacks including US Department of Justice access (Fruhlinger, 2019), the Yahoo data breach (Williams, 2017), and the DNC email breach (Satter, 2017) just to name a few. The 2019 FBI Crime Report (FBI, 2019) shows that phishing attacks account for the highest number of crimes at double the second place crime. Furthermore, the loss associated with phishing scams totals almost 58 million dollars annually. Phishing emails therefore pose a significant risk for individuals, corporations, and governments that is damaging both monetarily and informationally.

Many academic institutions now teach concepts related to security and penetration testing. Books and resources are available that provide materials and guides related to the art of reconnaissance and attack (Allsopp, 2017; Kim & Faircloth, 2015). The technical aspects of an attack can be implemented in a classroom or lab setting but the social aspect of cybersecurity attacks present unique challenges for the classroom setting.

Research that has investigated educational modules pertaining to social engineering (Kirk, Foreman, Lee, & Beasley, 2019; Weanquoi, Johnson, & Zhang, 2018) reveals that social engineering topics such as phishing focus on student ability to differentiate between safe and unsafe emails rather than creating opportunities for students to run these attacks in the manner of an actual penetration tester.

This study provides a model for a single course project that gives students an experiential learning experience on social engineering. In cooperation with a participating corporation, a phishing exercise is utilized to instruct students in both the technical and behavioral aspects of social engineering. Students develop a statement of purpose document, design the system, setup the environment, and conduct a phishing exercise on actual employees. Students then take the results and present the information to the corporation for usage in security awareness training.

SOCIAL ENGINEERING/PHISHING

Social engineering is, fundamentally, a type of confidence game. As Mouton et al. (2014) point out there are many definitions of social engineering used in academic studies. They ultimately define social engineering as "The science of using social interaction as a means to persuade an individual or an organization to comply with a specific request from an attacker where either the social interaction, the persuasion, or the request involves a computer-related entity." This interaction can

be direct or indirect. The earliest article on social engineering was written in 1987 (Quann & Belford, 1987) and spoke to the exploitation of help desks and support services.

Social engineering, as an umbrella term, has since expanded to include many different types of attacks using a variety of mediums. Krombholz et al. (2015) provide a reasoned taxonomy of these attacks, though specific forms of attacks evolve daily as new technologies and new opportunities for deceit present themselves. They break the attacks down into physical approaches, social approaches, reverse social engineering, technical approaches, and socio-technical approaches. Physical approaches rely on in-person actions by the bad actor and include dumpster diving, looking for passwords written down in workspaces, etc. Social approaches involve direct communication between the bad actor and the victim, typically by phone or face to face. Reverse social engineering relies on the bad actor presenting as a trusted entity and waiting for the victim to contact the bad actor in good faith. Technical approaches fit the common idea of "hacking" as shown in movies and media. Using tools or frameworks like Kali, Maltego, Metasploit, Armitage, NMap, Wireshark, etc. fall into this category.

Socio-technical approaches combine these other attacks and include baiting and phishing. Baiting involves leaving infected media, such as a USB drive, in locations where victims are likely to acquire and use the media. Phishing uses email, instant messaging or other forms of mass communication to reach a large audience in hopes of getting a few victims who are particularly susceptible to being duped. Spear-phishing is similar, but more targeted and typically relies on the messages being created such that they appear to be sent from trusted friends, co-workers, bosses, etc. For the purposes of this study we include spear-phishing under the general term "phishing".

Eighty-eight percent of organizations in a recent InfoSec survey reported that they faced phishing attacks in 2019 (Proofpoint, 2020). Furthermore, recent increases in teleworking brought about by the coronavirus pandemic have brought a sense of urgency to addressing the threats associated with social engineering. A joint alert about COVID-19 being exploited by bad actors using social engineering was issued in April 2020 from United States Department of Homeland Security (DHS) Cybersecurity and Infrastructure Security Agency (CISA) and the United Kingdom's National Cyber Security Centre (NCSC). The alert specifically cautions that "this is a fast moving situation" and that individuals and organizations should "take proactive steps to protect themselves." The alert specially focuses on the threat of phishing for money, credentials, and malware deployment using SMS and email.

Given the clear and present danger represented by social engineering it is paramount that it is addressed in university cybersecurity curricula. While a focus

on countermeasures to social engineering (typically raising awareness) can be found in cybersecurity curricula, an actual hands-on teaching of the topic is largely absent (Kirk et al., 2019; Twitchell, 2006). This is somewhat expected, as teaching general cybersecurity in the classroom can be challenging. Laboratory environments can provide a safe space for hands-on experiences with technical attacks. Related controls and terms and general domain content can be delivered through lecture. Actually getting students to engage in social engineering presents a different set of challenges. Social engineering is fundamentally an act of subterfuge. That requires that one of the two parties to the attack be ignorant as to the hidden agenda of the other party. Having participants in a class try and socially engineer other students in the same class simply won't work because all parties would be aware of the activity. Tasking students to socially engineer other students outside of that class would create ethical, and potentially legal, concerns. Referring to the Krombholz et al. (2015) taxonomy the physical and technical approaches to social engineering might be accommodated in a classroom setting. Students could search for passwords, clues etc. in a prepared physical environment, or use a lab setting to develop experience using technical applications. But how can students learn about social or socio-technical approaches from the perspective of a penetration tester?

There is little in the way of research to guide the creation of an experiential learning experience for students interested in social engineering. A review of cybersecurity education papers at SIGCSE and ITiCSE conferences from 2010 to 2019 showed 17 papers that focused on human aspects but none that offered an implementation for teaching hands-on social or socio-technical social engineering. In those same papers, broader cited reviews of cybersecurity education revealed papers focused on curricula standards rather than any accessible implementation strategy for the social niche of social engineering (Fujs, Mihelič, & Vrhovec, 2019; Jones, Namin, & Armstrong, 2018; Parrish et al., 2018). While there is plenty of support for including social engineering in curriculum, there is a corresponding scarcity of guidance on implementation. Instead, papers that offer a path for teaching social engineering are generally written to the identification of a specific form of attack.

As noted by Kolb and Kolb (2005), experiential learning theory encompasses six different areas; learning is a process, learning is relearning, learning requires conflict resolution, learning is a holistic adaption to the world, learning results from transaction with the environment, and that learning is the process of creating knowledge. More commonly it is the aspect of transacting with the environment that is the focus of educators when they invoke the term "experiential learning", associating it – as we do here – with giving students interactive, hands-on interaction outside of the confines of the academic classroom. This concrete

experience is one of the four stages of the experiential learning process, along with reflective observation, abstract conceptualization and active experimentation.

ACTIVITY

In this study we focus on the creation and delivery of a student project focused on phishing, a socio-technical social engineering attack. Studies on teaching about phishing uniformly focus on raising awareness (Arachchilage, Love, & Scott, 2012; Kumaraguru et al., 2007; Sheng et al., 2007; Stockhardt et al., 2016; Sun & Lee, 2016; Weanquoi et al., 2018) rather than giving students a method for experientially engaging in an actual phishing attack. The student project presented here provides an example of a hands-on attack, resulting analysis, and student creation of a tutorial for replication.

The students first met with a corporate sponsor regarding the phishing exercise. In this instance one student of the group already had a relationship with the CEO. For broader applicability, corporate members of a departmental advisory board, active student employers, alumni, and even the school IT department may all provide avenues for sponsoring this activity. Since the relationship for the exercise exists between the students and the organizational leadership (under the supervision of the instructor) the company leadership is responsible for choosing what employee business data to share with the employees. Similar to a real pen test, the employer also decided what, if any, notice or warnings are provided to the employees. Though IRB policies vary from school to school, student class assignments are not intended to develop or contribute to generalizable knowledge and therefor do not meet the federal definition of research as specified in 45 CFR 46.102 (I). Instructors implementing a classroom project in the form of this study are advised to check with their institutional IRB staff.

Over a series of meetings between the students, sponsor, and professor, the basic scenario for the exercise was drafted. When all parties had agreed on the parameters, a draft agreement was developed by the students that was then signed by the students, the professor, and the sponsor. The sponsor must agree to a one-way hold harmless agreement provided that the exercise does not exceed the agreed upon parameter. Several iterations of the draft were developed before the final version was signed to verify that every parameter of the exercise was explicitly described and acceptable to all parties, just as in a real world pen testing agreement.

The student group then went about setting up the technical backend for their exercise. This exercise requires the purchase of a domain name, a server with email routing, and website hosting available on the Internet and the open source program Gophish or similar application. The server should be a local computer managed by

the instructor, similar to any other hands-on student cybersecurity exercise. Using an externally hosted domain would not be appropriate.

A domain name was registered exactly like the parent organization of the sponsoring company, with the exception of the suffix (i.e. replace the .org with another suffix). Next, the students installed Ubuntu Linux on a local server in the lab. A key consideration was that the server needed to be accessible on the public Internet in order for email routing and website hosting. Given this, the students worked with university IT to obtain a public IP address to use for the server. Next, a Postfix email server application was installed and configured to provide a method for the phishing emails to be sent and any replies received. Apache web server was also installed on the machine to provide a simple form submission page for any employees who actually clicked on the link in the email to be directed to this page. Finally, the Gophish open source platform was installed on the server by the student group (Wright, 2020). Gophish provides an integrated web environment for setting up, conducting, and analyzing phishing exercises.

After setup was completed and tested, the following 3x2 experimental exercise was conducted. As documented in the drafted agreement, three potential vulnerabilities were identified with increasing associated risk.

1. The employee opens the phishing email.
2. The employee clicks on the link in the phishing email.
3. The employee submits information on the webpage.

Two employees were identified with differing levels of authority, including the CEO of the company and an administrative assistant. This provided two differing levels of influence in order to gauge the impact of this influence on the recipients. Target employees were then randomly assigned to receive an email from one of the two source senders, while utilizing blocking in order to ensure the same rough mixture of employees in each group (e.g. same number of mid-level managers, same number of main-line employees, etc.)

The email was designed to be generic in order to be less startling to the users. Below is the prose of the email:

Hello,

We will be sending out thank you's for the speakers who came this year. If you care to participate please follow the link below.

LINK

Thanks,
(Automated Signature)

Additionally, the webpage was also designed in such a way as look professional and hopefully illicit a response. The server stayed up for two weeks to give the email recipients plenty of time to respond. The student group analyzed the results and provided an executive write-up to the sponsoring corporation. This write-up included relevant statistical charts and recommendations for the sponsor.

RESULTS

The Gophish program offers several dashboards to investigate the results of the phishing exercise. As an overview, it can show how many emails were sent, how many were actually opened by users, how many users clicked the link in the email, and how many users submitted data to the fake web form. Figure 1 displays an example of this dashboard used during internal testing.

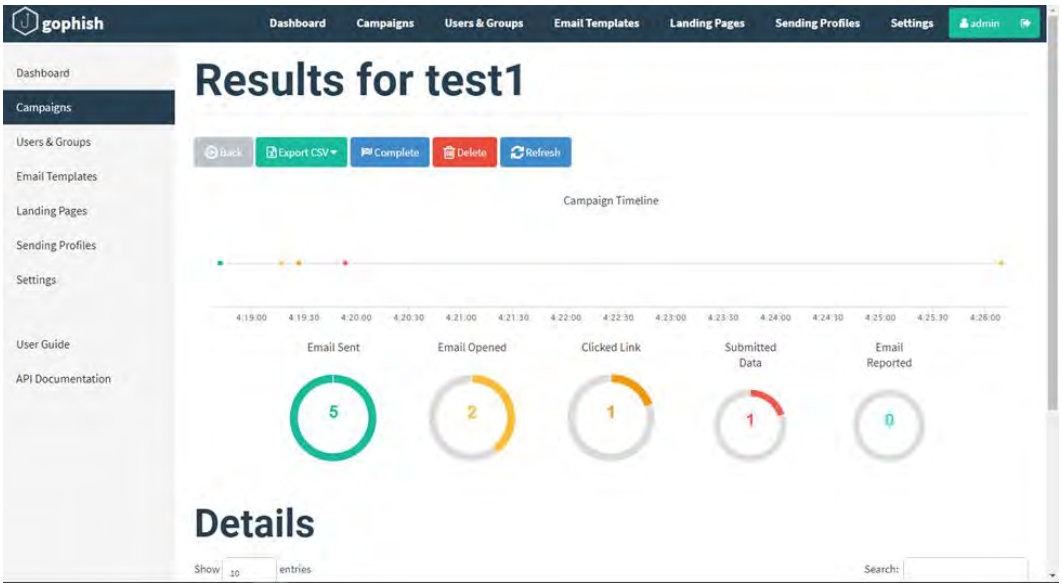

Figure 1. Gophish "Campaigns" overview form dashboard

Gophish will also show the results of individual users so that you can know who specifically opened an email, clicked on a link in the email, or submitted data. Figure 2 shows an example of this type of detailed dashboard, again with test data used for setup.

Figure 2. Individual users and their specific actions.

The students were able to aggregate this data to provide a detailed report to the target organization. Overall they found that while more individuals were likely to open an email, click a link, and submit data when the email looked like it was coming from the boss, several also opened emails and clicked links from one of the lower-level employees (see Figure 3). The students were also able to provide the identifiers of the employees to the boss who could then use this information to deliver extra training to those individuals susceptible to the social engineering attack.

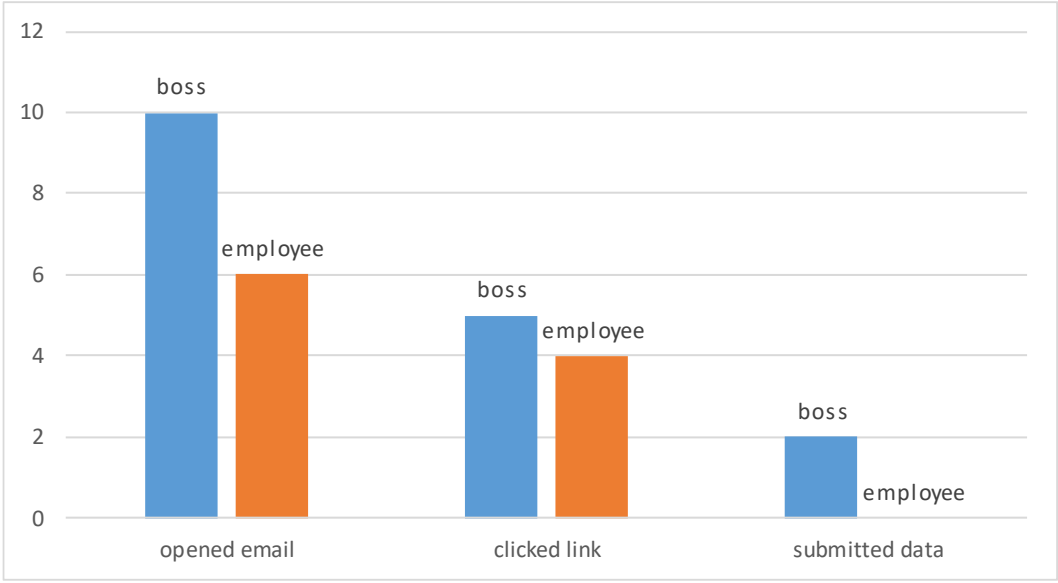

Figure 3. Results for the 2 (boss vs. employee sender) by 3 (opened email, clicked link, submitted data) experiment

COURSE IMPLEMENTATION AND ASSESSMENT

The level of student knowledge and expertise can certainly vary with this project. The setup of the server, the email application, apache, and even the web page development could all be done by faculty prior to the project. The course in which this particular project was used is part of a larger cybersecurity curriculum supporting the participating school's certification as a national Center of Academic Excellence in Cyber Defense (NSA, 2020) so this group of students had prior classroom experience with setting up these technical solutions. Depending on the level of student involvement, a suggested learning goal and related objectives for a similar environment may include:

Learning goal: On conclusion of this activity students will have an experiential understanding of the technical, social, and legal aspects of conducting social engineering tests in the broader context of penetration testing.

Learning Objectives: Students will be able to:

Install Linux on a local computer;

Install Apache on a local computer;

Install an email server application on a local computer;

Obtain and register a domain name;

Create a basic web page;

Configure a local computer to host a web page on the Internet;

Create and understand the elements in an email designed to unobtrusively gain the trust of an end user for pen testing purposes;

Create and understand the elements of a web page designed to unobtrusively gain the trust of an end user for pen testing purposes

Negotiate and create a legal pen testing agreement with limited scope and appropriate hold-harmless conditions.

Gather, analyze and present project outcomes.

Assessment of student learning can be accomplished in multiple ways. First, technical objectives can be assessed by directly observing the success or failure of the solutions. Second, qualitative assessment for the social interaction aspects, focusing on each student's reflective observation of challenges, solutions, and insights gained during the project can also be used.

DISCUSSION

Social engineering is a vital issue in today's technical environment and is exacerbated by the increase of the work-at-home population following the COVID-19 outbreak. Phishing and targeted spear-phishing are two social engineering techniques that are widespread and can result in a significant loss of revenue, resources and private information for both consumers and corporations.

Training and education for phishing largely focuses on raising awareness by having participants learn to recognize the characteristics of phishing emails, such as malformed URLs. Teaching students the other side of an attack is much more problematic due to the inherent unethical nature of such attacks.

This research provides an example of a social engineering educational project. Utilizing a phishing exercise targeted at a willing real-world company, students are provided with applied knowledge pertaining to the legal, technical, and analytical sides of social engineering. The students are first introduced to the legal side by working with an organization to develop a statement of purpose document detailing all aspects of the exercise. Next, the students are introduced to the technical aspects by implementing a functional phishing environment using the Gophish open-source software package. Finally, students are introduced to the analysis and reporting phase by gathering, examining, and reporting on the results of the exercise to corporate stakeholders. This project gives students some degree of experiential learning, mostly focused on the concrete experience and reflective observation.

Several areas for future research exist. The email and data elicited from the employees was designed to be less startling in nature; future research could look at varying levels of information sensitivity in a request to users. This would help address the limited use of active experimentation. Having students suggest or create alternative attacks that use the same human weaknesses as phishing might also address the need to abstract conceptualization in this experiential project. Also, the small nature of the target organization may have impacted the tendency for individuals to respond to the email. Overall, the learning module provides a highly interactive educational tool for instructors to teach both the technical and behavioral aspects of social engineering that could be useful in many programs.

REFERENCES

Allsopp, W. (2017). *Advanced Penetration Testing: Hacking the World's Most Secure Networks*: John Wiley & Sons.

Arachchilage, N., Love, S., & Scott, M. (2012). Designing a mobile game to teach conceptual knowledge of avoiding'phishing attacks'. *International Journal for e-Learning Security, 2*(1), 127-132.

FBI. (2019). *2019 Internet Crime Report*. Retrieved from

Fruhlinger, J. (2019). Social engineering explained: How criminals exploit human behavior. *CSO*. Retrieved from https://www.csoonline.com/article/2124681/what-is-social-engineering.html

Fujs, D., Mihelič, A., & Vrhovec, S. L. (2019). *The power of interpretation: Qualitative methods in cybersecurity research.* Paper presented at the Proceedings of the 14th International Conference on Availability, Reliability and Security.

Jones, K. S., Namin, A. S., & Armstrong, M. E. (2018). The core cyber-defense knowledge, skills, and abilities that cybersecurity students should learn in school: Results from interviews with cybersecurity professionals. *ACM Transactions on Computing Education (TOCE), 18*(3), 1-12.

Kim, P., & Faircloth, J. (2015). The hacker playbook 2. *Secure Planet LLC*.

Kirk, S., Foreman, D., Lee, C., & Beasley, S. W. (2019). Sit Back, Relax, And Tell Me All Your Secrets. *Journal of Cybersecurity Education, Research and Practice, 2019*(2), 4.

Kolb, A. Y., & Kolb, D. A. (2005). Learning styles and learning spaces: Enhancing experiential learning in higher education. *Academy of management learning & education, 4*(2), 193-212.

Krombholz, K., Hobel, H., Huber, M., & Weippl, E. (2015). " Advanced social engineering attacks"; Journal of Information Security and Applications, 22 (2015), S. 113-122.

Kumaraguru, P., Rhee, Y., Acquisti, A., Cranor, L. F., Hong, J., & Nunge, E. (2007). *Protecting people from phishing: the design and evaluation of an embedded training email system.* Paper presented at the Proceedings of the SIGCHI conference on Human factors in computing systems.

Mouton, F., Leenen, L., Malan, M. M., & Venter, H. (2014). *Towards an ontological model defining the social engineering domain.* Paper presented at the IFIP International Conference on Human Choice and Computers.

NSA. (2020). National Centers of Academic Excellence. Retrieved from https://www.nsa.gov/resources/students-educators/centers-academic-excellence/

Parrish, A., Impagliazzo, J., Raj, R. K., Santos, H., Asghar, M. R., Jøsang, A., . . . Stavrou, E. (2018). *Global perspectives on cybersecurity education for 2030: a case for a meta-discipline.* Paper presented at the Proceedings Companion of the 23rd Annual ACM Conference on Innovation and Technology in Computer Science Education.

Proofpoint. (2020). 2020 State of the Phish Report.

Quann, J., & Belford, P. (1987). *The hack attack-increasing computer system awareness of vulnerability threats.* Paper presented at the 3rd Applying Technology to Systems; Aerospace Computer Security Conference.

Satter, R. (2017). Inside story: How Russians hacked the Democrats' emails. *Associated Press*. Retrieved from https://apnews.com/dea73efc01594839957c3c9a6c962b8a/Inside-story:-How-Russians-hacked-the-Democrats%27-emails

Sheng, S., Magnien, B., Kumaraguru, P., Acquisti, A., Cranor, L. F., Hong, J., & Nunge, E. (2007). *Anti-phishing phil: the design and evaluation of a game that teaches people not to fall for phish.* Paper presented at the Proceedings of the 3rd symposium on Usable privacy and security.

Stockhardt, S., Reinheimer, B., Volkamer, M., Mayer, P., Kunz, A., Rack, P., & Lehmann, D. (2016). *Teaching phishing-security: which way is best?* Paper presented at the IFIP International Conference on ICT Systems Security and Privacy Protection.

Sun, J. C.-Y., & Lee, K.-H. (2016). Which teaching strategy is better for enhancing anti-phishing learning motivation and achievement? The concept maps on tablet PCs or worksheets? *Journal of Educational Technology & Society, 19*(4), 87-99.

Twitchell, D. P. (2006). *Social engineering in information assurance curricula.* Paper presented at the Proceedings of the 3rd annual conference on Information security curriculum development.

Weanquoi, P., Johnson, J., & Zhang, J. (2018). Using a game to improve phishing awareness. *Journal of Cybersecurity Education, Research and Practice, 2018*(2), 2.

Williams, M. (2017). Inside the Russian hack of Yahoo: How they did it. *CSO*. Retrieved from https://www.csoonline.com/article/3180762/inside-the-russian-hack-of-yahoo-how-they-did-it.html

Wright, J. (2020). Gophish. Retrieved from https://getgophish.com/