Issues in Information Systems
Volume 24, Issue 1, pp. 260-269, 2023

DOI: https://doi.org/10.48009/1_iis_2023_122

Lost and found: testing cybersecurity preparedness at a university

Gideon Hockenbarger, *Washburn University, gidjhock@gmail.com*
Nan Sun, *Washburn University, nan.sun@washburn.edu*
John Haverty, *Washburn University, john.haverty@washburn.edu*
Homer Manila, *Washburn University, homer.manila@washburn.edu*

Abstract

This study was conducted to assess the cybersecurity preparedness of a college-level population. As the world becomes more dependent on computers, it becomes more essential and more difficult to protect them. While bad actors are certainly capable of infiltrating a system without inside help, the number of attacks initiated by simple poor computer safety hygiene far outweighs all other attack methods. Given that universities have digital access to not only the personally identifying information (PII), but also the academic and professional work of thousands of people, they are particularly tempting targets. This experiment was undertaken to learn how likely a university is to fall prey to a cyberattack, in an effort to guard against real attacks in the future. The results demonstrated a disconnect between cybersecurity best practices and whom the practices are built to protect, with the target university's population being snared by a simple, low effort, and easily avoidable attack vector: "infected" USB drives.

Keywords: cybersecurity, penetration test, college security, social engineering

Introduction

Computing technology has increasingly become an integral part of society. Hospitals run thousands of pieces of advanced electronic equipment and maintain huge stores of patient information; energy companies collect data in the trillions to facilitate their services; companies keep mind-boggling quantities of data on servers; the work of a person's whole life can be contained on their laptop or smartphone; and colleges use virtual systems to continue their students' academic and personal development. When these technologies are threatened, it ought to be a grave matter, and one to take seriously, because all it takes to hand over control of a person's or organization's systems is clicking one bad link.

One does not have to look far to find evidence of the very real danger cyberthreats pose. According to data accumulated by the FBI, over $10 billion in damage was caused by cybercrime just in 2022 (Petrosyan, 2023). This is part of a trend that has risen markedly over the past 5 years, and sharply since 2020. These threats range from more common dangers, like locking down a user's machine with ransomware or stealing personal passwords, to events with drastic consequences, like impeding the ability of a hospital to provide life-saving services (Collier, 2022).

In this study, our goal was to explore how the population of a medium-sized university in the Midwestern U.S. would react to a simulated cyberattack. By measuring their response, we might determine how well current cybersecurity education is working, and where there may be gaps in the defenses. This research aimed to answer two questions:

1. How cybersecurity-conscious is the population of a college campus, and more specifically, how likely are they to fall prey to a cyberattack?
2. Where might there be areas in need of reinforcement?

As the only truly effective way to test security is to catch the target by surprise, we undertook a widespread penetration test of the college's population via "infected" USB keys.

This paper will cover some existing literature in this space, then outline our attack process. Following that are the collected data and associated analysis, wrapping up with a look at what was learned and what could be improved for future research.

Literature Review

In this section, we provide a brief look into available literature covering modern cybersecurity risks and an existing study aimed at resolving those risks, along with brief insights from other literature. These attacks are commonly focused on companies, where there are hundreds or thousands of potential entry points and far larger coffers to be drained. To prove this point, a study performed by Positive Technologies found that in an average of two days, 93% of companies could be breached by an external attacker, and 100% could be taken over by an internal attacker (2021, p. 4). While their methods didn't rely on social engineering, it has proven to be very effective for circumventing security measures. According to the FBI's Internet Crime Report, phishing—a method that inherently relies on people unwittingly click a link planted by a bad actor—was the highest-reported method of cyber-crime, outnumbering the next-highest method more than 5-to-1 (FBI, 2022, p. 21).

While most people may be aware of the dangers of suspicious links or responding to unprompted emails, there are more subtle ways to steal data, ones that target people's curiosity and good will. One such method, called USB baiting, involves "[installing] malware onto USB sticks and [leaving] them in strategic places, hoping that someone will pick the USB up and plug it into a corporate environment…" (Cisco, 2023, para. 19). They use either infected USB drives or human interface devices (HIDs). These are made to look like USBs and can execute scripts once plugged in. While seemingly ridiculous, this method has nonetheless proven effective; the most popular example of this is the infection of Iran's nuclear development program with the Stuxnet virus via USB drive (Fruhlinger, 2022).

University of Illinois Study

From social engineering texts (Hadnagy, 2011) to *Real Genius* (Coolidge, 1985), this idea of security breach via a strategically placed infected USB key has lived as a cybersecurity anecdote for some time. But in 2016, a study by the University of Illinois sought to put this to the test by targeting the school's three campuses with "infected" drives (Tischer et al., 2016). These drives contained links disguised as files leading to a website built and controlled by the researchers containing a survey. The goal was to get users to pick up a drive, plug it in, and click on one of these links, whereupon the server would record a handful of non-identifying information from the user's browser and take them to the researcher's site. Of 297 drives dropped, 290 (98%) were picked up, 135 (45%) were interacted with, and 62 (21%) usable survey responses were received (Tischer et al., pp. 4-6). There was no statistically significant difference in the type of drive picked up or the demographics of those retrieving them, but results suggested that users were plugging in drives and clicking on them out of a desire to return the drive to its original owner.

Other Research

There is one concept accepted by a wide array of cybersecurity publications: informing users is an essential element of protecting them. As suggested by the high rates of successful phishing attacks, the greatest risk to users is themselves. Therefore, finding a way to successfully impart sound cybersecurity doctrine is the most effective way to protect systems. How exactly to do this isn't clear, though. What is clear is that material needs to feel personal; needs to "be [designed] to gradually change people's behaviors" rather than check boxes on corporate training requirements; and needs to create simple and actionable rules for recipients to follow (Bada, et al., 2015; Nachin, et al., 2019; Wilson & Hash, 2003). To that end, this study was performed to gain a sense of our target's security awareness; with this in hand, the university could create a testing or training program geared towards its population.

Methodology

To achieve our goals, flash drives were loaded with links to a website managed by the researchers. Clicking a link would send a notice to the server attached to the website that contained: the drive's ID, the time the link was clicked, the IP address of the system contacting the server, and what link was clicked. The site's landing page held a statement of the study's purpose, a consent form, and a survey. This survey would be taken on a voluntary basis. Once complete, users would submit it to the back-end server and have the chance to see what information about their browser was obtained when they clicked the link.

There were several possible routes to "infecting" a machine. To avoid the risk of causing damage to users' information and devices, it was decided a non-intrusive method would be best—that is, one that would neither install anything to the user's machine, nor download anything from it. For the sake of expenses and time, it was also decided that methods using HIDs, such as a Flipper or Rubber Ducky device (Zhukov, n.d.), were impractical. And to achieve easy universal interaction that wasn't browser- or OS-dependent, it was decided that a web-based approach gave us the best chance of finding a one-size-fits-all attack vector. Combining these, it made the most sense to follow Tischer, et al.'s design, paring it down to make it more practical for a single researcher to accomplish over six weeks. As such, the process would involve recording minimal user information from the browser—the user's IP address, the time a link was clicked, the ID of the drive accessing the site, and the link that was clicked—followed by a request that users complete a survey to give us workable data.

With the oversight of the target university's Information Technology Services (ITS) department and the approval of the institutional review board (IRB), a survey was conducted. This was divided into four sections: Interaction, Demographic, Security Assessment, and Reaction questions. Interaction questions asked users why they picked up the drive and how they felt about their security awareness; Demographic questions requested basic demographic information recommended by Tischer, et al.; Security Assessment questions were from the SeBIS cybersecurity assessment (Egelman & Peer, 2015); Reaction questions asked how people reacted to the drives and the experiment, and if they'd like to leave us with any further comments.

To host this survey, a website was built on a private university server run by one of the university's CIS (Computer Information Sciences) professors. The website featured a landing page, the survey, and an exit page. The landing page informed users of the study's purpose, inviting them to participate in the research by clicking a consent form agreement. If they agreed, they were taken to the survey. Upon completion, they were taken to a page containing a brief informational section covering basic cybersecurity concepts, as well

as a display of information able to be collected from their browser; this included the user's OS and browser, as well as the date/time and IP address the site was accessed from. Additionally, the site hosted seven identical PHP scripts of different names (see Table 1 on page 5) that would be linked from the attack flash drives. These scripts were the first step in gathering data and bringing participants to the landing page, as all seven links redirected there after collecting data and setting session variables.

Once the website was complete, links were loaded onto 120 generic 1GB flash drives. Each drive was given an ID based on the day and location it would be dropped. A combination of links to the website were loaded onto each drive, ensuring that every drive received at least one each of an "identifying link"—i.e., links to PHP scripts disguised as files that could reasonably help a person identify the owner, such as a resume or academic paper—and a "non-identifying link".

The links were disguised as a variety of files to gauge why a person would click on a file; for instance, drives contained links disguised as resumes and family photos. Which link they clicked might help determine if a person would search a drive's files out of a desire to find the owner, or purely for the sake of their own curiosity. Additionally, there were four types of drives distributed (see Figure 1):

1. **Blank** – Flash drives without any ornamentation or modification; files on these drives could be disguised as anything.
2. **Labeled** – Flash drives with labels such as "Final Exam", "Class Notes", or "Private"; files on these drives had names relevant to their labels.
3. **Keys** – Flash drives with keys and keychains attached; files on these drives could be disguised as anything as well.
4. **Name Tags** – Flash drives with nametag keychains, each featuring one of five male or five female names and an associated email address. These names were generated using lists of the most popular male, female, and surnames in the U.S., and the email addresses were to accounts controlled and monitored by the researcher. Files on these drives could be disguised as anything, but care was taken to name at least one file on each drive using the name on the tag, so as to entice participants to click those first.
5. **"Other"** – Flash drives unique from and larger than the others, with the name of a well-known cybersecurity/information storage company printed on one side. These drives were made to test the reaction of areas that had been informed of the project prior to its beginning. There were only 4 of these drives deployed, 1 each in the CIS departmental office, the campus police department, the main help desk, and the ITS help desk (no image is provided for anonymity).

Figure 1: Examples of Flash Drives Used.
Left-to-right: Blank, Labeled, Keys, Name Tag

There were two purposes to this approach. First, by analyzing the number of each drive type plugged in/returned, we hoped to identify patterns of behavior valuable to future security training. Second, by diversifying the appearance of the drives, we hoped to reduce suspicion from the campus' population.

30 drives of each type (120 total) were then distributed across ten locations on the college's campus on three days over a two-week period. There was one drive that was lost and could not be recovered, but every other drive was deployed and monitored. Each day, four drives would be left in each general area. To keep knowledge of the experiment to a minimum, only six locations were established as collection points for drives: the campus police station, the university's Student One Stop, the ITS help desk, the library front desk, the campus recreation center's front desk, and the student information center. These locations were checked at least once a week for returned drives.

To manage distribution of these drives, five researchers were assigned to two locations each on the university's main campus. Three days—a Monday, a Wednesday, and the following Wednesday—were designated as "drop days", again spread out to avoid suspicion. Researchers dropped one drive of each of the four types at each of their locations; this way, 40 drives were dropped per day. Researchers recorded the time and location of each drop, along with the drive's ID and any abnormalities, in a spreadsheet. Drives were dropped in a variety of locations, including cafeterias, student work areas, classrooms, bathrooms, departmental offices, parking lots, etc. Drives were checked as regularly as possible, and by the end of the experiment period, all but two had been taken.

Results

By the end of the two-week collection period, a total of 64 clicks had been recorded across 27 identified drives. However, 12 of the clicks had no associated drive ID, and therefore couldn't be included in some counts. This means there were between 28 and 39 drives plugged in, an approximate 23-33% plug-in rate. Only five survey responses were recorded. This lack of sufficient responses meant significant in-depth analysis of the results was not possible. However, we could still investigate what was collected from user clicks since the scripts collected no compromising data. Of the 64 clicks, 25 (39%) were of documents disguised as resumes, with a further 13 (20%) being academic papers. Given that either of these document types could have reasonably given a user the name and/or contact information for the owner, it can be said with reasonable certainty that nearly 60% of the time, a user is interacting with a found device with the intention of finding its owner. Additionally, 15 of the identifiable drives (55% of identifiable drives) had multiple clicks recorded.

Table 1: Distribution of File Types Clicked

File Clicked	Count	Percentage
Resume.pdf.php	25	39.06%
Academic Paper.php	13	20.31%
Personal Photo.php	12	18.75%
dPersonal Paper.php	5	7.81%
Miscellaneous Personal Media.php	4	6.25%
Game.php	3	4.69%
Professional Photo.php	2	3.13%
Grand Total	64	100.00%

Table 2: Frequency of Clicks Per DriveID

driveID	Count of driveID	driveID	Count of driveID
0	12	86	2
18	4	118	2
61	4	3	1
62	3	8	1
70	3	9	1
78	3	13	1
85	3	19	1
95	3	26	1
100	3	31	1
5	2	71	1
11	2	90	1
37	2	97	1
59	2	99	1
74	2	102	1
		Grand Total	**64**

Of the 119 dropped drives, 57 had been received by the six recovery locations or turned in directly to researchers by the time of the experiment's closure; a further 6 were recovered in the following week. The location with the highest rate of turn-ins was the campus police station, which hosts the campus's lost-and-found, at 14 drives. The location with the lowest rate of turn-ins was the ITS help desk at just 3. Interestingly, only 9 of these 57 drives had recorded clicks, meaning the overwhelming majority were returned unused. Drives with name tags attached saw the highest rate of return; 23 of these 30 drives were returned, over 75%.

The lowest rate of return was for those without any labels, with only 8 of the 30 making their way back. Of the drives that were plugged in and able to be identified, 10 of the 27 (37%) were drives with only labels on them, making it the drive type with the highest distribution of clicks; only 2 drives with names and email addresses were plugged in and interacted with. Additionally, 11 emails were received by both the researcher's emails and the monitored addresses; however, due to technical issues, some addresses bounced emails, potentially making the number higher.

Table 3: Distribution of Drives Collected Per Collection Location

Collection Location	# Drives Collected	Distribution of Returned Drives by Location				
		Blank	Labeled	Keys	Name Tag	Other
Campus Police Station	14	0	2	4	7	1
Library Front Desk	11	2	3	2	4	0
Student One Stop	8	1	0	3	3	1
Rec Front Desk	8	3	1	1	3	0
Returned Directly	8	0	3	1	3	1
Student Information Center	5	1	1	1	2	0
ITS Help Desk	3	1	1	0	1	0

Table 4: Distribution of Drives Collected from Total Dropped Per Drive Type

Blank		Labeled		Keys		Name Tag		Other	
8/30	26.67%	10/30	33.33%	12/30	40%	23/30	76.67%	3/4	75%

Table 5: Distribution of Identifiable Records by Drive Type

Blank		Labeled		Keys		Name Tag		Other	
8/27	29.63%	10/27	37.04%	7/27	25.93%	2/27	7.41%	0/4	0%

The average time between a drive being dropped and it registering a click was approximately 41 hours; however, as shown in Figure 2, the distribution of these times skews right, and the average was highly influenced by 2 outliers whose times exceeded 190 hours. Removing those outliers gives a more reasonable approximate mean of 29 hours, with 19 of the 25 identifiable drives (76% of the set of identifiable drives after outliers are removed) falling below this average. The longest a drive sat before registering a click was 193 hours and 33 minutes, or just over 8 days; the shortest time was just 9 minutes.

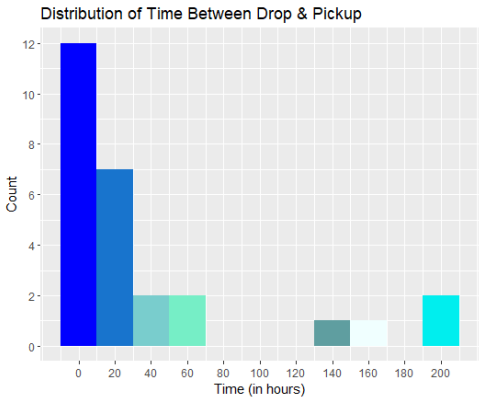

Figure 2: Distribution Plot of Time between Time Drive Dropped and Click Recorded (Only includes identifiable records)

While there weren't enough responses to conduct much significant analysis, they are covered generally here for completeness. Of the five responses received, there was a decent distribution across the demographics surveyed. Notable was the nearly even split between male and female, with a wide distribution across age, income level, and role at the university. From the interaction questions, the average respondent was confident in their security knowledge and ability to respond to cybersecurity threats; additionally, they were equally curious about the contents of and desirous of finding the owner of their USB drive. Only two of the five respondents made any effort to protect their devices when interacting with the drives: one opened the link in a text editor, one scanned the drive with anti-virus software, and both plugged the drive into an old computer they wouldn't mind getting damaged.

Discussion

Only receiving five survey responses was far from ideal, but the fact the experiment saw such a high interaction rate was surprising. Not only did we get 64 clicks, but over half the identifiable drives that were plugged in had multiple files clicked, suggesting users may have been uncertain if the web site's landing

page was legitimate. While we cannot know for sure the motivation of those who clicked the link but did not participate in the study, we can assume by the distribution of file types clicked that the majority were attempting to locate the owners of the drive. This assumption is reinforced by a more than 75% turn in rate for drives with name tags, as well as the few survey responses we did receive.

More surprising than the 23-33% pickup-and-click rate was the nearly 50% drive return rate we saw across the length of the study. Within 48 hours of the first drops, 21 drives—over 50% of the first wave—had been returned. On several occasions, a researcher deposited a drive only to notice someone walking by with it in hand shortly after the drop had been made; on at least 2 occasions, researchers had a drive returned to them directly by someone who saw them drop it. On the other hand, the fact that nearly half the drives are still unaccounted for is concerning; as many as 56 people may have driven that, if this were an actual attack, could be carrying malware or other hazards.

There were positive reactions of note from the university's students and faculty. In one instance, a professor was given a drive labeled with their class name by one of their students. The professor then plugged the drive into a sacrificial Linux system, opened the file in a text editor, and upon realizing its purpose, completed the survey before returning the drive directly. In another instance, a professor received not one, but three separate drives from their students. This professor expressed a great deal of concern that multiple students had lost notes related to their classes, going so far as to send out emails to their entire class body asking if anyone recognized the drives. But importantly, *they never plugged a drive into their computer*. The police department also received multiple emails from office staff around campus asking what to do with drives they had received from other sources. All of these are positive behaviors. The refusal to plug in drives—at least, before taking precautions—is encouraging.

Conversely, there were some less than ideal reactions. One drive's attached keys and keychain were found where they had been left beside a computer, but the drive had been ripped free and stolen; another drive was taken, registered a click, then returned to the same location it was taken from. This latter event is concerning because in the case of an actual attack, this behavior could lead to multiple infections from a single source.

Also of note was the response one departmental office had to the experiment. This office reached out to ITS to inquire about the nature of the drives only after having found 4 drives and registering clicks on 2 of them. When ITS informed them of the experiment's nature, this office argued its staff should have been informed. Further, the department was concerned they had been included in an experiment of this nature without their consent.

This highlights several relevant points. First, the fact staff plugged in and interacted with multiple unknown flash drives suggests proper protocol may be unclear or is not reinforced. Second, they contacted ITS only *after* causing a potential breach, rather than before, which is concerning because cybersecurity must be preventative rather than reactive. Third—and in relation to the previous two—this shows that most often, people are their own worst enemy. In attacks such as this one, which rely on an insider exposing the network, one bad click or flash drive plugged in out of curiosity can be enough to give an attacker full access to your system. And fourth, the reaction of office staff to learning this was an experiment highlights why it is essential these tests be done, and be done as discreetly as possible. Due to the nature of the experiment, consent couldn't have been given before they had clicked a file on the flash drive. The purpose of these experiments is to gauge a person's natural reaction to an attack, but to do it in a safe environment. If we had informed the office's staff it was an experiment, they never would have plugged the drives in. But they *didn't* know it was an experiment, and they *did* plug the drives in, suggesting their natural reactions

would have resulted in a network breach. In a real situation, this would have been catastrophic; given it happened in a controlled environment, it helps the university's ITS department recognize weak points in need of reinforcement.

Conclusions, Implications. Limitations, and Future Research

The purpose of this study was twofold: First, see which groups are most vulnerable to these attacks to find where common security gaps may exist. Second, test the cybersecurity awareness of a mid-sized university in the U.S., with the aim of highlighting the need for cybersecurity measures on every level. Our results, while unable to unveil any specific patterns or permit data analysis, still prove there is much work to do in cybersecurity education. While it is likely impossible to create an airtight security program, the fact that at least 28 people were willing to plug in a USB drive they had no knowledge of is alarming—as is the fact that nearly half the drives weren't recovered. It illustrates there are still people at the university level unaware of—or unconcerned about—the grave dangers a lax cybersecurity stance exposes them to. On the other hand, the reactions from much of the faculty and staff demonstrate the university's staff training is effective. Many targeted drives—ones labeled with specific class names dropped in classrooms and offices related to those classes—were returned straightaway. This may suggest a regular campaign of education directed at the student body, including tests such as this one, could have a similar impact on campus cybersecurity.

This project was not without limitations. Foremost was the lack of respondents; just over 18% of the identifiable drives—4% of total drives—received a survey response. We might have secured a greater response pool by offering some form of compensation for completing the survey. Alternatively, by dropping more drives, or by targeting the university's other campus(es), we might have widened our target population and collected more data. Minor technical problems also affected our data collection.

Regardless of these issues, we think this research could prove useful in future testing and research. By crafting tests that target specific areas of a college—say, exclusively sports facilities or lab buildings—you could more easily test a single subpopulation without requiring survey responses. Future research might begin with an experiment like this one, followed up by an information campaign, and then another penetration test at the end of the semester; that way, you can measure how well your campaigns work over several semesters.

Acknowledgements

This project was completed with the aid of Dr. Philip Hauptman and Dr. Bruce Mechtly in the CIS Department at Washburn University. Also deserving of many accolades are Anna Phelps, Seth Phelps, Moriah Phelps-Roper, and Luke Phelps-Roper, the four aides who helped with the distribution, monitoring, and recovery of the flash drives. Their willing sacrifice of time and effort made managing 120 flash drives possible.

References

Bada, M., Sasse, A.M., & Nurse, J.R.C. (2015) Cyber Security Awareness Campaigns: Why do they fail to change behaviour? *Working Paper of the Sustainable Society Network+*. 118-131. https://www.researchgate.net/publication/274663655_Cyber_Security_Awareness_Campaigns_Why_do_they_fail_to_change_behaviour

Bradley, T. (2023, March 26). *Shifting Cybersecurity To A Prevention-First Mindset*. Forbes. https://www.forbes.com/sites/tonybradley/2023/03/26/shifting-cybersecurity-to-a-prevention-first-mindset/?sh=63f706de59cc#Presume%20Breach:%20A%20Flawed%20Approach

Cisco (2023). *What is Social Engineering in Cyber Security?*. https://www.cisco.com/c/en/us/products/security/what-is-social-engineering.html

Collier, K. (2022, October 7). *Ransomware attack delays patient care at hospitals across the U.S.* NBC News. https://www.nbcnews.com/tech/security/ransomware-attack-delays-patient-care-hospitals-us-rcna50919

Coolidge, M. (Director). (1985). *Real Genius* [DVD]. United States; TriStar Pictures.

Egelman, S., & Peer, E. (2015). Scaling the security wall. *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, 2873–2882. https://doi.org/10.1145/2702123.2702249

Federal Bureau of Investigation, Internet Crime Complaint Center. (2022). *Internet Crime Report*, 21. https://www.ic3.gov/Media/PDF/AnnualReport/2022_IC3Report.pdf

Fruhlinger, J. (2022, August 31). *Stuxnet explained: The first known cyberweapon*. CSO Online. Retrieved May 2, 2023, from https://www.csoonline.com/article/3218104/stuxnet-explained-the-first-known-cyberweapon.html

Hadnagy, C. (2011). Social Engineering: The art of human hacking. Wiley Publishing, Inc.

Nachin, N., Tangmanee, C., & Piromsopa, K. (2019). How to Increase Cybersecurity Awareness. *Information Systems Audit and Control Association, 2019*(2). https://www.isaca.org/resources/isaca-journal/issues/2019/volume-2/how-to-increase-cybersecurity-awareness

Petrosyan, A. (2023, April 24) *Amount of monetary damage caused by reported cyber crime to the IC3 from 2001 to 2022* [Infographic]. Statista. https://www.statista.com/statistics/267132/total-damage-caused-by-by-cyber-crime-in-the-us/#statisticContainer

Positive Technologies (2021, December 20). *Business in the crosshairs: analyzing attack scenarios.* 1-4. https://www.ptsecurity.com/upload/corporate/ww-en/analytics/Business-at-gunpoint-analyzing-attack-scenarios-eng.pdf

Tischer, M., Durumeric, Z., Foster, S., Duan, S., Mori, A., Bursztein, E., & Bailey, M. (2016). Users really do plug in USB drives they find. *2016 IEEE Symposium on Security and Privacy (SP)*. https://doi.org/10.1109/sp.2016.26

Wilson, M., & Hash, J. (2003). Building an information technology security awareness and training program. *Computer Security*, 23. https://doi.org/10.6028/nist.sp.800-50

Zhukov, A. *Turning a Regular USB Flash Drive into a USB Rubber Ducky*. HackMag. https://hackmag.com/security/rubber-ducky /