Historical Archived IP List: Leveraging AWS to Persist Slack Security Data
(Technical Paper)

Analyzing the Effectiveness of Gamification in Cybersecurity Trainings for Organizations
(STS Paper)

A Thesis Prospectus
In STS 4500
Presented to
The Faculty of the
School of Engineering and Applied Science
University of Virginia
In Partial Fulfillment of the Requirements for the Degree
Bachelor of Science in Computer Science

By
Jason Yu

October 27, 2022

On my honor as a University student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments.

ADVISORS

Roseanne Vrugtman, Department of Computer Science

Bryn E. Seabrook, PhD, Department of Engineering and Society

Prospectus

Introduction

Traditional cybersecurity trainings are notoriously boring and ineffective (Reeves et al., 2021), yet organizations have historically relied on traditional security education and training to mitigate cyberattacks. According to a recent TalentLMS survey on the state of cybersecurity training, 61% of employees who took cybersecurity training failed a basic test (Marousis, 2021). Across the board, employees struggle to retain and apply what they learn in cybersecurity trainings to their everyday work lives. In an attempt to make cybersecurity education more engaging and effective, employers are increasingly turning to gamification methods. Gamification is the use of game mechanics and game thinking to engage users in solving problems and to motivate them by introducing elements of competition and reward (Moore, 2017). Gamification is increasingly popular, but the effectiveness of gamification is not well-documented, so the proposed STS research project will explore the effectiveness of gamified cybersecurity trainings compared to traditional cybersecurity trainings.

The proposed technical project relates to my internship project in the summer of 2022 while working on the product security tooling team at Slack. Slack security harnesses a tool called RAINS (Rapid Analysis, Internal Network Scan) to provide visibility into Slack's AWS (Amazon Web Services) infrastructure in order to alert engineers about unauthenticated services and defend against subdomain takeover attacks. However, RAINS does not keep an historical log of facts and findings, making it difficult to determine the cause of a potential security incident. To solve this problem, I implemented HAIL (Historical Archived IP List), consisting of a database to store RAINS findings and a backend API, allowing users to query past results. I leveraged AWS RDS (Relational Database Service) for the data layer, created the backend service using Flask, and used AWS Lambda and S3 (Simple Storage Service) in conjunction with the Risk and Compliance team's Security Data Warehouse project, allowing users to easily view RAINS results. I also modified the RAINS codebase (written in Go) to call the Flask backend and pass the findings to RDS. By the end of the internship, I successfully deployed HAIL to a production environment and configured metrics and dashboards using Prometheus and Grafana.

Technical Topic: Historical Archived IP List

Slack is a messaging application designed for workplaces and office productivity. Slack customers primarily include businesses, so Slack is a likely target for cyberattacks. To defend against certain attacks, Slack owns a tool called RAINS (Rapid Analysis, Internal Network Scan) that provides visibility into Slack's AWS infrastructure, alerting engineers about cyber risks such as services lacking proper authentication and dangling DNS entries.

Slack manages a number of internal services intended to be accessed only by authorized Slack employees. To ensure security, these sites must be placed behind an authentication portal. If RAINS detects an internal service that does not require authentication, then RAINS sends an alert to Slack security engineers. Unauthenticated internal services allow attackers to steal sensitive information or company secrets.

Slack routinely spins up and tears down thousands of hosts on a daily basis, and each host is associated with a human-readable CNAME (canonical name) via DNS (Domain Name System). A dangling DNS entry occurs when Slack tears down a host but does not remove the corresponding DNS record. In some cases, a dangling CNAME record in one of Slack's subdomains is all that an attacker needs to take control of the content served by the subdomain in question. Such an attack is known as a subdomain takeover, and the potential risks include company defacement and stealing session cookies, allowing attackers to pose as Slack employees.

RAINS solves the problem of detecting unauthenticated services and dangling DNS entries by scanning Slack's AWS infrastructure every 30 minutes. Indeed, Slack has historically paid tens of thousands of dollars through its bug bounty program due to such risks. However, RAINS does not keep a historical log of facts and findings, a limitation that makes it difficult to determine the source of a possible breach. To remedy this problem, HAIL (Historical Archived IP List) is a project designed to store RAINS scan results for future reference. In the project, I used AWS RDS (Relational Database Service) for the data layer, created the backend API (application programmable interface) service using Flask, and used AWS Lambda and S3 (Simple Storage Service) in conjunction with the Risk and Compliance team's Security Data Warehouse project, allowing users to easily view RAINS results. I also modified the RAINS codebase (written in Go) to call the Flask backend and pass the findings to RDS. By the end of the internship, I successfully deployed HAIL to a production environment and configured metrics and dashboards using Prometheus and Grafana.

STS Topic: Effectiveness of Gamified Approaches to Cybersecurity Trainings

Cyberattacks pose major challenges to businesses and organizations. According to PurpleSec's 2021 Cybersecurity Trends Report, over 50% of all cyberattacks are done on small to midsize businesses (SMBs), and enterprises experience approximately 130 security breaches per year, per organization, on average (Firch). Human error is one of the primary factors that enables cyberattacks to be successful (Ahola, 2021). In a security context, human error encompasses unintentional actions (or inaction) by employees and users that cause, spread, or allow a security breach to take place. Human error includes a vast range of actions, from downloading a malware-infected attachment to failing to use a strong password (Leal, 2022). According to research from Elevate Security, human behavior had a direct role in 88% of total losses in the largest cybersecurity incidents over the past five years, and about two-thirds of major data breaches are the result of humans (2021). Similarly, phishing is the most common way in which malware is delivered. In fact, a 2019 report by NortonLifeLock found that 92.4% of malware is delivered as an attachment in a malicious email. To mitigate the human risk element, organizations have historically relied on traditional security awareness education and training (The Defence Works, 2019). Traditional security awareness trainings methods include models in which large groups of employees periodically attend one-day events (Hadley, 2018) and mandatory security awareness trainings involving reading information from slides and taking knowledge-checking quizzes afterwards (Obudulu, 2022). Unfortunately, traditional methods of cybersecurity training have proven to be largely ineffective. According to a recent TalentLMS survey on the state of cybersecurity training, 61% of employees who took cybersecurity training failed a basic test (Marousis, 2021). Not only are cybersecurity trainings relatively ineffective, they are also boring. A 2021 paper found that employees generally have a negative attitude toward cybersecurity trainings, citing a lack of interest in the presentation style and a perceived mismatch between the training content and real-world scenarios (Reeves, et al.). Across the board, employees struggle to retain and apply what they learn in cybersecurity trainings to their everyday work lives.

To make cybersecurity education more engaging and effective, employers are increasingly turning to gamification methods. Gamification is the use of game mechanics and game thinking to engage users in solving problems and to motivate them by introducing elements of competition and reward (Moore, 2017). Some examples include ThreatGEN’s *Red vs. Blue*, a game-based cybersecurity simulation platform that combines a gaming engine with an adversary simulation A.I., theoretically tailoring the content of the game to match the user's skill level (2022) and Centrical's game platform that allows users to complete challenges involving a variety of game narratives (e.g., Hide and Seek, car races, hitting targets) in order to earn coins, badges, and move up on a leaderboard (2022).

The proposed STS research project will explore the effectiveness of gamification methods in cybersecurity trainings compared to traditional cybersecurity training methods. Specifically, this project will explore the factors that allow gamification approaches to be more successful than traditional approaches in order to determine the most effective way to implement such cybersecurity trainings. This question is important because gamified approaches are new, and researchers have not fully studied the effectiveness of gamification. Therefore, in order to make progress and improve the cybersecurity landscape, it is essential to know whether gamification is effective and how to implement gamification to optimize effectiveness.

This research project will use a Social Construction of Technology (SCOT) framework to analyze the relationship between cybersecurity trainings and relevant social groups, which include employees, employers, cyber attackers, cybersecurity professionals, and everyday users. SCOT posits that technology does not determine human action, but rather than human action primarily shapes technology. The SCOT framework is apt for the topic of gamification in cybersecurity because innovations in cybersecurity trainings are relatively recent developments, and designers currently have substantial flexibility in shaping the future of cybersecurity trainings. This observation aligns with SCOT's tenet of interpretive flexibility, which suggests that technology design is an open process that can produce different outcomes depending on the social circumstances of development (Klein & Kleinman, 2002). One common criticism of

SCOT is that it fails to account for the ways in which technology shapes human action (Pinch & Bijker, 1984). However, for this project, it is most important to emphasize the ways in which humans construct and shape technology, especially because the primary research question at hand investigates the effectiveness of a novel cybersecurity training approach.

Research Question and Methods

The STS research portion will aim to answer the question of how effective gamification is in cybersecurity training for organizations. To answer this research question, I plan to use documentary research methods by collecting scholarly articles. Documentary research methodologies are most appropriate for this project because they offer data collected in a controlled environment rather than through mere hearsay or anecdotes. First, to supply background information, I will collect a number of articles detailing the current cybersecurity landscape. These articles will include statistics assessing the effectiveness of traditional cybersecurity training programs as well as interviews that provide insight into employees' perceptions of such trainings (Reeves, et al., 2021). Next, I will compile scholarly articles that detail individual cases of gamified approaches to cybersecurity training, noting relevant similarities and differences between approaches and identifying their benefits and drawbacks (Adams & Makramalia, 2015, Sabillon, et al., 2019). Finally, I will research meta-analyses regarding the effectiveness of gamified approaches to cybersecurity training to help answer the research question more fully. A thematic organization of sources is favorable because the research question is ultimately asking about the disparities between two fundamentally different approaches to cybersecurity trainings. Keywords include "cybersecurity," "training," "gamification," "engagement," and "effectiveness" since these are most relevant to the research topic.

Conclusion

Both the technical report and the STS research paper are related to cybersecurity. The technical report will analyze HAIL, a project completed for the Slack product security team. HAIL keeps a historical record of facts and findings related to Slack's AWS infrastructure, providing valuable insight into Slack's attack surface and security posture. This technical portion will detail the relevant security considerations, the implementation process, and the final outcomes. The STS research paper will focus on the human aspect of cybersecurity, namely cybersecurity trainings for employees to mitigate poor security practices and social engineering attacks. Gamified approaches to cybersecurity training are increasingly popular, but the effectiveness of such methods remains an open question. By examining academic literature and applying the SCOT framework, the STS research portion will deliver findings that help shed light on the effectiveness of gamified approaches in order to ultimately improve the future of cybersecurity trainings.

References

Adams, M., & Makramalia, M. (2015). Cybersecurity Skills Training: An Attacker-Centric Gamified Approach. *Technology Innovation Management Review, 5*(1): 5-14. http://doi.org/10.22215/timreview/861

Ahola, M. (2021, February 1). The Role of Human Error in Successful Cyber Security Breaches. *usecure*. https://blog.usecure.io/the-role-of-human-error-in-successful-cyber-security-breaches

Centrical. (2022). Gamification for Employee Engagement. *Centrical*. https://centrical.com/platform/gamification/

Elevate Security. (2021, May 11). Elevate Security and Cyentia Institute Launch First Annual Study on Employee Cybersecurity Risk in the Workplace, Finds Current Solutions Do Little to Reduce Human Error. *Elevate Security*. https://elevatesecurity.com/elevate-security-and-cyentia-institute-launch-first-annual-study-on-employee-cybersecurity-risk-in-the-workplace-finds-current-solutions-do-little-to-reduce-human-error/

Firch, J. (2021, April 29). 10 Cyber Security Trends You Can't Ignore In 2021. *PurpleSec*. https://purplesec.us/cyber-security-trends-2021/

Hadley, J. (2018, October 31). How Traditional Training Is Weakening Businesses' Cybersecurity. *Forbes*. https://www.forbes.com/sites/jameshadley/2018/10/31/how-traditional-training-is-weakening-businesses-cybersecurity/?sh=1eacd74b4b0c

Klein, H., & Kleinman, D. L. (2002). The Social Construction of Technology: Structural Considerations. *Science, Technology, and Human Values, 27*(1): 28-52. https://doi.org/10.1177/016224390202700102

Leal, A. (2022, August 14). Human Factor in Cybersecurity: The Weakest Link? *KuppingerCole*. https://www.kuppingercole.com/events/csls2022/blog/human-factor-in-cybersecurity-the-weakest-link

Marousis, A. (2021, April 6). Cybersecurity training lags, while hackers capitalize on COVID-19. *TalentLMS*. https://www.talentlms.com/blog/cybersecurity-statistics-survey/

Moore, M. (2017). Bringing Gamification to Cyber Security Training. *University of San Diego*. https://onlinedegrees.sandiego.edu/bringing-gamification-to-cyber-security-training/

NortonLifeLock, Symantec Global Internet Security Threat Report (2019). Online Paper. https://docs.broadcom.com/doc/istr-24-2019-en

Obudulu, O. (2022, July 5). 7 Ways to Transform Your Cybersecurity Training and Influence Lasting Change. *skillsoft*. https://www.skillsoft.com/blog/7-ways-to-transform-your-cybersecurity-training-and-influence-lasting-change

Pinch, T. J., Bijker, W. E. (1984). The Social Construction of Facts and Artefacts: Or How the Sociology of Science and the Sociology of Technology Might Benefit Each Other. *Social Studies of Science, 14*(3): 399-441. http://www.jstor.org/stable/285355.

Reeves, A., Calic, D., & Delfabbro, P. (2021). "Get a Red-Hot Poker and Open Up My Eyes, It's So Boring": Employee Perceptions of Cybersecurity Training. *Computers & Security, 106*(2021): 102281. https://doi.org/10.1016/j.cose.2021.102281

Sabillon, R., Serra-Ruiz, J., Cavaller, V., & Cano, J. (2019). An Effective Cybersecurity Training Model to Support an Organizational Awareness Program: The Cybersecurity Awareness TRAining Model (CATRAM). A Case Study in Canada. *Journal of Cases on Information Technology*, *21*(3): 26-39. https://doi.org/10.4018/JCIT.2019070102

The Defence Works. (2019, February 19). Does Security Awareness Training Work? *The Defence Works*. https://thedefenceworks.com/blog/does-security-awareness-training-work/

ThreatGEN. (2022). Red vs. Blue. *ThreatGEN*. https://threatgen.com/