White Paper

Title

Blunting the Phisher's Spear: A risk-based approach for defining user training and awarding administrative privileges

Arun Vishwanath, Associate Professor, SUNY Buffalo, http://arunvishwanath.us

Track

Human Factors- Human Factors- Governance, Risk, and Compliance

Abstract

People today are the weakest links in cybersecurity. Solving the "people problem" of cyber security requires us to understand why people fall victim to spear phishing. Unfortunately, the only proactive solution against spear phishing is to train and educate people. But, judging from the number of continued breaches, training appears to be limited in its effectiveness. Today's leading cybersecurity training programs focus on hooking people in repeated simulated spear phishing attacks and then showing them the nuances in the emails they missed. This "gotcha game" presumes that users merely lack knowledge, and if they are told often enough and repeatedly shown what they lack, they would become better at spear phishing detection. We propose a radical change to this "one-size-fits all" approach. Recent human factors research—the Suspicion, Cognition, Automaticity Model (SCAM)—identifies two independent sets of factors that lead to individual phishing victimization: users' cognitive processing schemas premised on their perceptions about the safety of online behaviors, and their habits and patterns borne out of repeated, ritualistic behaviors influenced by work culture and the types of devices people use to connect and communicate. Using the SCAM, we propose the development of an employee Cyber Risk Index (CRI). Similar to how financial credit scores work, the CRI will provide security analysts the ability to pinpoint the weak-links in organizations and identify who is likely to fall victim, who needs training, how much training, and also what the training should focus on. The CRI will also allow security analysts to identify which users get administrative access, replacing the current mostly binary, role-based apportioning method, where individuals are given access based on their organizational role and responsibilities, with a system that is based on individuals' quantified cyber risk propensity. The CRI based approach we present will lead to individualized, cognitive-behavioral training and an evidence-based approach to awarding users' admin privileges. These are paradigm-changing solutions that will altogether improve individual cyber resilience and blunt the effectiveness of spear phishing.

White Paper

Blunting the Phisher's Spear: A risk-based approach for defining user training and awarding administrative privileges

Arun Vishwanath, Associate Professor, SUNY Buffalo, http://arunvishwanath.us

1. The problem in context: Cyber safety and why we must protect it?

At the turn of the twentieth century, automobile accidents were routine, so routine that riding a horse or walking was considered safer. Fast-forward to 2016 and in much of the developed world, automobile accidents occur within the margin of error. How was this achieved?

Quite simply, it was done by studying people—the operators—and understanding how they think, behave, and act, and re-designing roads and the mechanics of automobiles around drivers. Today, roads are built with specific pitches and curves so people can see what is ahead; car seats are designed with specific dials, lights, instruments, and equipment; and promotional campaigns are used to teach people to wear seat belts by addressing the reasons why people don't enact such behaviors. These are just a few of many ways in which automobile safety was realized, making road travel one of the safest and most commonly used modes of transportation.

This analogy is particularly fitting when it comes to protecting cyberspace. Much like the early days of automobiles, today's cyberspace is dangerous. Cyber attacks in this space are common and by now breaches have already compromised the personal information of almost everyone in the U.S.

A. The problem in numbers: Some of the well-know breaches in 2014-15 include Target where 70 million customer records were stolen, eBay where 145 million records were stolen, Home Depot 56 million, JPMorgan 76 million, Anthem 80 million, Ashley Madison 37 million, Experian 15 million, and OPM 22 million. These are but a few of the major breaches in the last year that were reported. Many more have likely already taken place but just not discovered—as it takes upwards to 1 year for breach discovery. Hackers have shown how to compromise every known security protection including two-factor authentication, and even biometrics, thanks to the many fingerprints stolen in the OPM hack.

B. What's at stake? In 2013, hackers sent a single tweet from AP's Twitter account claiming a series of explosions had occurred in the White House. In response to this, the Dow lost 144 points and the S&P lost $143 billion, which although later corrected gives us an idea of the enormous financial stakes. Here are some other scary events that occurred in late 2015–early 2016: a major cyber attack crippled Ukraine's electricity grid—an attack that took down their grid, froze the computer terminals of operators trying to restart the grid, and blocked all their telephone lines so consumers couldn't call-in and check what was even going on. Another attack on Israel's electric grid necessitated a temporary shutdown of its northern grid on two of its coldest days. Germany reported extensive damage to an industrial plant from a hacked blast furnace that couldn't be stopped. Closer to home, the Department of Homeland Security received reports of close to 300 infrastructure incursions including one by Iranian hackers on the sluicegate controllers of the Bowman Avenue dam in Rye, New York. Compare the scope and scale of recent attacks to those from even five years back and

one thing becomes apparent: attacks are getting more brazen and their consequences more severe.

C. Who is conducting these attacks? The actors include terrorist groups funded by Iran and Hezbollah, ISIS, the Syrian Electronic Army, hacktivists groups such Anonymous and Green Lizard, and criminal enterprises from Ukraine to Uganda who have found ways to monetize stolen information. In staggering attacks, ISIS hackers placed location beacons on computers and devices of Syrian sympathizers for pinpointing their location presumably for assassination. In another case, a hacker group from Ukraine that was recently indicted netted in excess of 100 million USD over a five-year period by phishing and hacking into several financial news organizations and selling advance trade information to "investors." Compare this with the estimated 120 million USD that Al Qaeda netted over an eight-year period from ransoms! Thus, a successful breach can net enormous capital that could be used for a variety of purposes, all within a shorter period of time, and with limited personal risk. There are also growing cases of ransomware being used by hackers. For instance, the Sheriff's Office of Dickson County, Tennessee, and a hospital in California are among the many reported victims who have reported paying hackers ransoms to unlock documents that had been encrypted using CryptoWall.

2. The problem in context: What are hackers doing?

While the motives for cyber attacks as well as their sources vary considerably, there is a common thread that runs across them. Almost all, from the attack on the AP to the one crippling Ukraine, begin with a spear phishing attack, where the hacker hides a malware payload in the attachment of an email, which when clicked opens a backdoor into computer networks that are then used to hijack system-controllers or exfiltrate data. Some phishing attacks direct individuals to fraudulent websites that run malicious scripts or directly solicit login and other credentials by spoofing a real website. Most, from the attack on DoJ to the one crippling Ukraine, utilize spear phishing, where the hacker hides a malware payload in the attachment of an email, which when clicked opens a backdoor into computer networks that are then used to hijack system-controllers or exfiltrate data. Some phishing attacks direct individuals to fraudulent websites that run malicious scripts or directly solicit login and other credentials by spoofing a real website.

3. The solution today: What is being done?

Presently, the only pro-active solution utilized by organizations all over the world is training employees to improve their resilience against spear phishing. So how successful is this intervention? Much of the training is done by for-profit enterprises that do not provide their software for peer-review. Thus, it is difficult to precisely ascertain the impact of training. We can, however, triangulate and indirectly estimate the value of training using three different sources. One line of evidence comes from the many organizations that mandate training (including OPM, DHS, DoJ, and such) and continue to fall victim to spear phishing. Another line of evidence comes from independent studies conducted by organizations such as the Army Cyber Institute (ACI). The ACI studies involved training cadets and phishing them at various intervals afterwards. Their results again point to training having only a short-term impact and any influence of education wearing-off within a day or so.

Part of the problem with training remains how it is conducted. Training designed to improve people's resilience against spear phishing focuses on telling people where to find the deceptive cues in a phishing email. Often such training uses mock phishing attacks directed at employees within the organization, who if victimized by the attack are taken to a sort of "splash page," which tells them the attack was a simulation and reiterates what they missed in the email that led to their victimization.

There are a number of limitations of such approaches. On the one hand, frequent training gets routinized and, like any other mandated action, employees being to comply merely by completing the task rather than learning from it. This is especially so when the same monotonous training content is repeated and the individual feels that s/he is already well aware of the information being presented. On the other hand, training through simulated attacks makes employees scared of falling victim, especially when there is some punitive effect of falling prey to simulated attacks. Beside this, training using simulated attacks is akin to teaching people how to drive cars better by repeatedly causing accidents—and then telling them why they had an accident. Finally, all the training today assumes that people fall victim because they lack knowledge about phishing. Empirical evidence shows even people with good cyber security knowledge and even those with high awareness (such as the Robert Mueller when he was Director of the FBI) falling victim to spearphishing. This is discussed next.

4. Understanding the people problem: Why do people fall victim to spear phishing?

There appears to be three main reason why people victim to spear phishing. A recent model called Suspicion, Cognition, Automaticity Model (SCAM) elucidates the reasons for it. The model was empirically developed and found to predict 20-40% of the variance in individual victimization likelihood. Figure 1 present the SCAM model.

Fig 1. The Suspicion, Cognition Automaticity Model (Vishwanath et. al. 2016 [1])

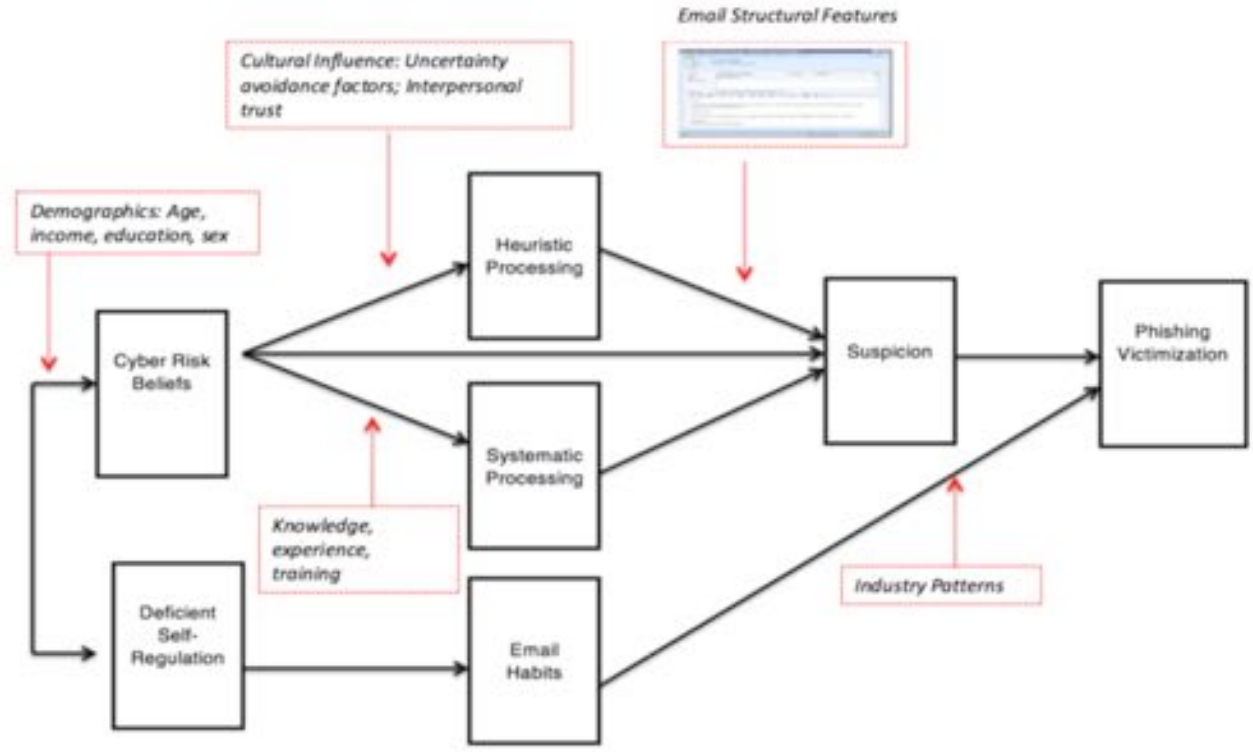

SCAM points to two major factors leading to individual victimization through phishing. The first are cognitive factors beginning with whether people are suspicious about a phishing email. This occurs at the intersection of the structural features of an email and the willingness and ability of the individual to exert sufficient cognitive effort.

This ability to exert cognitive efforts stems from the motivation of the individual to engage in such effort and the knowledge s/he has that helps spot the deception through

the leakage cues in the structural features of the email (sender's reply-to address, IP of sender, and such).

The expectation of training and education is that people are always applying sufficient cognitive effort but merely lacking knowledge of these leakage cues. But, in reality, people seldom commit cognitive resource to analyzing information. People almost always act as cognitive misers, choosing efficiency over detailed processing. Information processing efficiency occurs through the use of cognitive shortcuts or heuristics. This form of processing called heuristic processing allows for quicker and faster information processing. Examples of this include judging the credibility of an email simply because one sees the phrase "Sent from my iPhone" on the message; or the authenticity of a webpage by the color, font, or look and feel of the page.

Outside of processing, a parallel influence on deception stems from people's habits. Habits form as people perform a behavior repeatedly under relatively stable conditions. Such routinization leads to the formation of behavioral action scripts that require little to no cognitive mediation for enactment. These could happen for simple behaviors, such as checking ones phone routinely, to complex behaviors, such as driving to work.

Email use presents a particularly strong case for habituation for the following reasons. Firstly, people frequently interact with emails throughout the day, making it a behavior that is continually enacted. Secondly, checking email is part of the routine of web surfing and often the primary reason for repeatedly going online. Finally, most email exchanges tend to be fairly benign, making it easy to relax cognitive involvement and form formulaic patterns of usage over time. While research has yet to connect such habitual email use to deception, the SCAM contends that the lack of cognitive mediation during habitual email use is perhaps another reason for individuals to ignore cues within a phishing email and fall for the deception. The impulsive nature of habitual enactment means that individuals' cyber-risk beliefs may not be activated and they might fail to consider the risks in opening or responding to the email. As a result, neither are deceptive cues recognized, nor do they trigger suspicion. Consequently, individuals under the influence of email habits are less likely to be suspicious of phishing emails and more likely to be deceived. Such routinization is particularly strong when people have certain personality factors that inhibit the regulation of behavior, chiefly an inability to regulate behaviors. The self-regulatory mechanism involves self-control by being observant about one's media-related actions and contrasting it against what is acceptable. A failure of this control mechanism leads to deficient self-regulation. Thus, a deficiency in the ability to control email use leads to individuals' formation of email habits.

By providing a framework that incorporates both the cognitive as well as the behavioral influences on phishing-based victimization, SCAM provides a comprehensive launch pad for understanding the process of individual victimization through spear phishing.

5. The Proposed Solution: How can we leverage this understanding to improve cyber security?

Knowing what we do about the individual level factors that lead to victimization based on the SCAM, there are three changes to existing protocols and practices that the presentation proposes.

A. Developing a Cyber Risk Index (CRI): Today, there is no unified mechanism for assessing an individuals' cyber risk propensity. Most organizations are, therefore, unaware of which user is a likely risk or weak-link in their cybersecurity. Thanks to the lack of unified measurement metrics, there is also no mechanism for comparing the cyber risk thresholds of different departments in an organization or across different organizational/sectors. The presentation will discuss how we can utilize the SCAM to

develop a Cyber Risk Index (CRI), which will provide a simple, unified metric for evaluating an individual's cyber risk propensity. The approach to developing this index will be a data-driven empirical approach that can be done either in isolation, using "red-team" type phishing simulations, or can be done in conjunction with existing training tools and approaches. CRI will provide a unified metric that will not only allow sysadmins to measure individual risk propensity but also allow them to aggregate it across divisions/work-group within and across organizations. Once developed, much like we provide financial credit scores, CRI scores can be aggregated across organizations, providing a quantitative metric for assessing the security risk across different organizational employees, divisions of organization, and organizations across different sectors, on their cybersecurity preparedness.

B. Using CRI to define who gets trained, how, and the types of training: The presentation will discuss how we can utilize the SCAM-based CRIs to improve our targeting of training. Much of todays cybersecurity training focuses on teaching people by making them have accidents, typically by sending them fake phishing emails, and then telling them why they fell for the phish. Often users who have many accidents are give even more such training, or worse yet individuated and shamed into becoming more careful online.

Instead, the current presentation proposes a different alternative to the current one-size fits all, "gotcha" cybersecurity training. We present a system that allows security analysts the ability to assess who among their users are likely to fall victim to different types of spear phishing attacks and, more importantly, why they will fall for these attacks. The "why" is defined based on the SCAM based Cyber Risk Index (CRI) that can help quantify, measure, and assess who needs training, how much, and the specific area that they need to be trained on.

The presentation will discuss a series of algorithms based on the SCAM for pinpointing the types of training people need. The decision algorithm will be as follows:

Is it the individual's cognition that led to victimization? —> If Yes, then, was it their lack of good heuristics? —>If Yes, then, provide them with efficient heuristics.

Is it their lack systematic processing? —>If Yes, then, is it a lack of motivation, apathy, or knowledge? —> If knowledge, then, educate them; if apathy/motivation, then, frame the need for engagement.

Is it their cyber risk beliefs? —> If Yes, then, find misperceptions and provide them with education that corrects the beliefs.

Is it their habits/patterns of working? —> If Yes, then, change the patterns or inject better, more supportive work flow patterns.

C. Using CRI to create a behavior-based admin authorization system: Organizations today allow for information access based on an old-school, binary system, where information access is provided based on an employee's role or status in the organization. But phishing is an equal-opportunity attack and does not discern between roles or functions. Even knowledgeable Internet users, trained IT professionals, and security experts (such as Robert Mueller when he was Director of the FBI [2]) have been successfully phished.

What we need is another layer of protection, one that effectively silos employee access to only specific areas within computer networks so that a single action or compromised device ported in by a person is incapable of crippling the entire grid. Unfortunately, the current system of according IT privileges was developed before the advent of mass hacking attacks and uses a simple binary system of providing access based on employee roles or positions in the organizational.

We propose a quantifiable, behavior-based, admin authorization system. This will be built using the CRI, which would help index who among the employees are most likely to fall victim to an attack, who is improving with defined training, and which

users/organizations units are the weakest-links. Based on comparisons of CRI scores overtime, sysadmins can reward users, based on their demonstrated cyber hygiene, with different degrees of access to system files and network locations. We believe that the CRI-based approach is dynamic, because it can be updated routinely, and can therefore accommodate changing roles of people, their different patterns of work, and the many devices they use to access and work with information.

Overall, the three pronged approach advocated here would lead to a smarter, dynamic IT management system, one that is resilient enough to address phishing and other threats; one that is able to create smarter training and education tools that take into account people's needs rather than try the "one shoe fits all" approach; and one that identifies the weak-links in organizational security and rewards only those users demonstrate good cyber hygiene.

These ideas challenge our existing approaches of training people and providing access to privileged information. It also provides an evidence-based analysis of what people do and strategies for improving our overall cyber safety. Overall, thus, the talk presents a series of fresh, innovative, paradigm-changing alternatives to the current practices that do little to stop spear phishing.

Sources

[1] Vishwanath, A., Harrison, B., & Ng, Y. J. (2016). Suspicion, Cognition, and Automaticity Model of Phishing Susceptibility. Communication Research, 0093650215627483: http://crx.sagepub.com/cgi/reprint/0093650215627483v1.pdf?ijkey=batj1qn6zPTHFiR&keytype=finite

[2] http://voices.washingtonpost.com/securityfix/2009/10/fbi_director_on_internet_banki.html

Attendee Takeaways

1. We will discuss the latest human factors research that explains why people fall victim to different forms of spear phishing attacks.
2. We will present a novel approach that leverages this research and develops a Cyber Risk Index (CRI)— a mechanism for iteratively pinpointing individual cyber risk propensity.
3. We present a new and unique mechanism for providing individualized training and rewarding users with different levels of administrative privileges based on their CRI-based risk quotients.

What is new?

The CRI-based approach to according training and awarding access has never been tried before. It leverages a recent research understanding of why people fall victim to cyber attacks, and challenges the current cyber security training paradigm that presume that everyone falls victim because they just don't know enough. This system can replace the current practices that were built before the advent of mass hacking attacks. If implemented it could lead to smarter training, education, and sys admin processes that are data-driven, evidence based, and comparable across individuals, divisions, and organizations over time.

Why Black Hat?

The CRI is novel, simple, and powerful, and changes the existing paradigm of cybersecurity training, which presumes that everyone falls for a cyber attack for the same reason, and repeatedly trains them using the same tried and failed system. The CRI-based approach will also change how sys-admins provide users administrative privileges, so only those who demonstrate cyber hygiene are rewarded with access to different systems and networks. The proposed approaches will, therefore, change the dominant system of user training and privileges apportioning and significantly improve individual and organizational cyber resilience—making Back Hat the perfect platform for making the info-sec community aware of this paradigm-changing approach.