

Use Authorization

@
@ o y @ 0
@
y 0 @ 8 o)
y 0 @

o

)

° `h=@=@/8'0' -k-)) -7-Vo-`
7k° U -‡ \kM'

°

" .

'M .

A thesis
submitted in partial fulfillment
of the requirements for the degree of
Master of Business Administration in the Department of
Business
Idaho State University
Spring 2014

#

o

.

u 8 7

u # M

) # o
U

)) M
U

)) o
8 7 k

Table of Contents

O	7
Abstract	
Chapter 1: Introduction	
Chapter 2: Related Work	
Chapter 3: Literature Review	
3.1 Cyber Kill Chain (Hutchins, Cloppert and Amin 2014)	
3.2 Modeling and Preventing Phishing Attacks (Jakobsson 2005)	
3.3 Phishing, Personality Traits and Facebook (Halevi, Lewis and Memon 2013)	
3.4 Self-Control and Criminal Opportunity: A Prospective Test of the General Theory of Crime (Longshore 1998)	
3.5 How Could I Fall for That? Exploring Phishing Victimization with the Heuristic-Systematic Model (Zhang 2012)	
3.6 The State of Phishing Attacks (Hong 2012)	
3.7 Who Falls for Phish? A Demographic Analysis of Phishing Susceptibility and Effectiveness of Interventions. (Holbrook, et al. 2010)	
3.8 School of Phish: A Real-World Evaluation of Anti-Phishing Training (Ponnurangam Kumaraguru 2009)	
3.9 Data Shield Algorithm (DSA) for Security against Phishing Attacks (Ram Avtar 2011)	
3.10 Special Publication 800-12: An Introduction to Computer Security – The NIST Handbook ((NIST) 1995)	
3.11 Lexical Feature Based Phishing URL Detection Using Online Learning (Aaron Blum 2010)	
3.12 INFORMATION SYSTEMS SECURITY: A COMPREHENSIVE MODEL (McCumber 1991)	
3.13 Psychological Safety and Learning Behavior in Work Teams (Edmondson 1999)	
3.14 On the Folly of Rewarding A, While Hoping for B (Kerr 1975)	
3.15 The Psychology of Security (West 2008)	
3.16 Technical Note – Analysis of a Layered Defense Model (Walter R. Nunn 1982)	
3.17 Defense in Depth: An Impractical Strategy for a Cyber World (Small 2011)	
3.18 Observations on the effects of defense in depth on adversary behavior in cyber warfare (Dorene L. Kewley 2014)	
3.19 Security education against phishing: A modest proposal for a major re-think (Lacovos Kirlappos 2014)	

3.20 Spear-Phishing Email: Most Favored APT Attack Bait (Team 2014)	· ·
3.21 Understanding Scam Victims: Seven Principles for Systems Security (Frank Stajano 2011)	· ·
Chapter 4: Concept	· ·
4.1 A Sequential Chain of Events	· ·
4.2 Layered Defense Model	· ·
4.3 The McCumber Cube	· ·
4.4 Categorical Defense Structure	· ·
4.4.1 Policy	· ·
4.4.2 Education and Training	· ·
4.4.3Technology	· ·
Chapter 5: Technical Defense Zone	· ·
5.1 Detection Toolsets	· ·
5.2 Preventative Toolsets	· ·
5.3 Warning Toolsets	· ·
Chapter 6: Human Defense Zone	· ·
6.1 Psychological Safeguards	· ·
6.2 Awareness	· ·
6.3 Culture and Reporting	· ·
Chapter 7: Further Research and Study	· ·
Chapter 8: Conclusion	· ·
References	· ·

List of Figures

7	·····	·····
7	·····	·····
7	·····	·····
7	·····	·····
7	·····	·····
.		
.		

Abstract

Recognizing phishing as one of the great threats to an institution, many researchers have set out to find a solution or create a tool that will prevent phishing attacks. Others have studied how the users can be prepared to better defend against such attacks, but nobody has done the research that ties it all together into a holistic defense strategy. This paper seeks to create such a holistic approach through the use of a layered defense strategy.

The intent of this research is to bring together the work that has been done by many other sources, to create a unified framework for defending against phishing attacks. Through the fusion of various studies, it is hypothesized that a more robust and complete defense strategy can be created and that through its implementation an organization will reduce the number of successful phishing attacks it experiences each year.

Chapter 1: Introduction

Today is a day when information travels at the speed of light. Financial transactions are handled as 1's and 0's rather than gold or paper. Email has all but decimated a once thriving postal industry. If you can dream it then you can probably get online to purchase it or get the plans to build it. If you want to find information on any topic imaginable, then start with searches using sites like Google or Yahoo. Servers and databases have replaced volumes of paper filed away in boxes and file cabinets. Governments, industry, and even our homes are managed by technology and our information is stored on hard drives at home or in the cloud, all in the name of convenience. But that convenience comes at a price. Because these storage solutions are interconnected over a meshed network we call the internet, we have the ability to make our data available for our use all over the world at any time. The problem is that so does everybody else. While we are getting better at securing our information from someone "knocking the door down and stealing it," the attackers are getting better at convincing us to open the door for them and inviting them in. Phishing is one of the tactics used to do just that.

Phishing is the art of tricking an individual into responding to a communication crafted for a nefarious purpose. Common mediums used in phishing include email, SMS (texting), phone communication, mail, MMS (chat), etc. Phishing specifically targets the users of systems rather than the systems themselves. A commonly seen example of phishing is the attempt to harvest valid credentials from the intended target. For example, an attacker may send an email posing as the victim's bank. In this email they may try to

convince the user that they need to update their account to prevent a future loss of access. The email would also include a link to a page that is crafted to look like the bank's website, but when the victim follows it and enters their banking credentials, the information is sent to the attacker rather than the bank, giving the attacker full access to the victim's account.

What attackers have found is that they can use this tactic not only to grant them access to individual accounts, but can also use it to gain access to entire networks of institutions. Once inside, they establish persistence, or the ability to come and go as they please, and just like that they have maneuvered past the defenses intended to keep them out. These threats, known as "Advanced Persistent Threats" (APTs), are recognized as some of the most prevalent and dangerous threats to the security of systems and data within the modern organization. TrendMicro, a leading security research firm, recently reported that 91% of APT breaches involved a spear-phishing attack, supporting their conclusion that spear-phishing remains one of the most favored tactics for instigating targeted attacks (Team 2014).

So how do we (information security professionals) fight back? What can be done against a threat that attacks the resources of a firm through its people? Recognizing that phishing is one of the great threats to an institution, many researchers have set out to find an answer or create a tool that will prevent phishing attacks. Others have studied how the users can be prepared to better defend against such attacks, but nobody has done the research that ties it all together into a holistic defense strategy. This paper seeks to create such a holistic approach through the use of a layered defense strategy. The intent of this research is to bring together the work that has been done by many other sources, creating

a unified framework for defending against phishing attacks. Through the fusion of various studies, it is hypothesized that a more robust and complete defense strategy can be created, and that through its implementation an organization will reduce the number of successful phishing attacks it experiences each year.

Chapter 2: Related Work

After considerable study, I have determined that no other work has considered the creation of a framework for phishing and its defenses that includes the use of psychological factors into a layered defense model. Other works exist defining a layered defense approach to information systems. Consider the work done by the Lockheed Martin Corporation around a cyber-kill chain approach (Hutchins, Cloppert and Amin 2014), or research from Markus Jakobsson that creates a basic model similar to that of a disaster recovery plan (Jakobsson 2005). In both examples, researchers have identified factors in further identifying potential attacks and even creating some layers of defense, but no research appears to have put together an in-depth review of the technical and human factors behind successful phishing attacks. Much of the research appears to be centered on using technical tools and algorithms to identify and prevent common attacks as they have happened in the past. Some research looks to answer the human factor through the use of training tools with varied approaches to so doing. Yet no one has built a framework putting it all together in a holistic approach.

Chapter 3: Literature Review

3.1 Cyber Kill Chain (Hutchins, Cloppert and Amin 2014)

Information assurance and cyber security professionals are at a disadvantage in their efforts to maintain the security of systems, including their confidentiality, integrity, and availability. It makes sense since, theoretically, the security teams of any organizations would need to prevent each and every attack, numbering in the 10's of 1000's for some, while the attacker really only needs to get it right once. The odds appear to be in favor of the attacker. It is on this premise that the Lockheed Martin Corporation researched and came up with what is commonly known now as the cyber kill chain. Their research is based primarily on the idea that the biggest threats to sensitive systems could be classified as those that seek to establish a persistent presence within an organization's systems; commonly known as advanced persistent threats (APT's). What they found is that attackers always follow a similar pattern when establishing this presence. Such a pattern has seven steps that provide seven different opportunities for organizations to prevent the attack from being successful, thus shifting the advantage from the attackers to the defenders. The idea is that organizations would "study intrusions from the adversaries' perspective. Each discrete phase of the intrusion is mapped to courses of action for detection, mitigation, and response." In other words, it is important not only to discover and mitigate a threat, but rather to study it and discover, as best one can, all phases along the kill chain that were utilized or were intended to be utilized. Then the organization would create detection, mitigation, and responses to all phases discovered, closing doors to re-entry as they go.

The seven phases described by Hutchins et al. include; (1) Reconnaissance, (2) Weaponization, (3) Delivery, (4) Exploitation, (5) Installation, (6) Command and Control, and (7) Actions on Objectives. Traditionally, security events have been handled and are resolved at their point of discovery, leaving all/most other phases untouched, making them reusable for the same actor or other new ones. By analyzing events as far as possible in either direction of the kill chain, greater security can be achieved in a quicker and more efficient manner. For example, a security analyst might discover an intrusion when an alert is generated due to the installation of blacklisted software onto a system. The analyst might clean out that system, restoring it to its clean state, and add the attacking IP address to the system's access control list, effectively blocking that attacker from further attacking that particular system. The problem is that if the exploit used and the delivery method were not found, the same attacker, or others, could follow the same path and simply use a different installation method. Conversely, the same analyst could analyze the event as far up and down the kill chain as possible, closing doors as they went. If they did they might discover that the exploitation came because the system was unpatched, and delivery was received through an open but unused port on the system. They might find that the company website inadvertently listed some details of the internal systems that the attacker used to gain an understanding of what exploit to use. All phases, in this example, can be addressed in a way that closes doors and mitigates multiple threats from future use.

Performing a successful phishing attack is not unlike other, more traditional, attacks on systems. Phishing requires phases and knowledge not only of systems but of people as well. There is a delivery method established, usually email, and a weaponization of that

delivery that allows for penetration into the organization's systems. In the end, it is not just a one-step event but a process of multiple steps/phases to successfully attack a firm's assets through the use of phishing. Similar to the research of Lockheed Martin's analysts, I believe research could be done on phishing attempts to create a framework of phases that would describe a phishing attack. This would include more than the discovery of mitigation strategies for phishing events, but would also include motivating factors, psychological traits of both attackers and victims, and cost to both attackers and victims. A framework would be beneficial to defenders, as they would potentially find alternate and more in-depth methods for reducing the threat posed by this prevalent and highly successful attack method.

3.2 Modeling and Preventing Phishing Attacks (Jakobsson 2005)

In previous research to identify a model for phishing attacks, Jakobsson utilizes a graphing model to represent the multiple paths, or threat vectors, that an attacker might use by means of phishing to gain a presence in a system, obtain personal information, or gain access to resources previously prohibited. He further uses a system of vertices and edges to notate access to information or resources and the likelihood of accomplishing such a task. This "graphing" system is a novel approach to describing the methods of attack when considering phishing and its variants as threats to any system. Jakobsson treats analyzing these attacks in a similar method to the way an organization might create a disaster recovery policy (DRP) (in a very general sense). His approach is to think of any given resource in the context of a protected item whereupon any number of incidents could affect it. It is up to the owner to identify and list the possible disasters (or attack

vectors), along with the associated steps to ensure success and likelihood for success.

These become the vertices of the model while the lines that connect the vertices identifying likelihood for success become the edges. Using this model, the defender can design mitigations to slow and even thwart attempts to phish any given victim.

There are two key aspects to this research that are pertinent and interesting in regard to designing a better framework. 1) This is the first attempt that I am aware of at a forward looking model to discover and create mitigation strategies for threats. 2) This is the first example I have found of a layered defense model. When approaching this model from a DRP perspective, the firm is doing more than finding ways to respond to events as they happen. They are instead looking forward into the future, seeking to identify threats as they currently exist rather than threats after they have begun their exploits. The advantage is that the firm can identify ways to block paths to exploitation before they begin. The team begins the thinking process early to discover weaknesses in their approach to protect data and systems from human vulnerabilities.

A second key aspect to this approach is that it encourages the use of a layered defense process. After recognizing the sources of phishing threats, the model encourages the further discovery and enumeration of the necessary steps to accomplish the overall goal. Once discovered, each door on the path can be given proper defense tactics to prevent the attacker from exploiting such a vulnerability and, in so doing, provide multiple layers of defense. The advantage to this type of defense structure is that by creating multiple barriers to entry, the defender not only blocks the one intended attack but also, likely, prevents other unconsidered attacks by creating many barriers in the path of exploitation that happen to coexist among various attack vectors. In laymen's terms, the defender

unintentionally blocks one attack while trying to block another, an efficient advantage when trying to defend against an innumerable number of attacks from many different attack vectors.

3.3 Phishing, Personality Traits and Facebook (Halevi, Lewis and Memon 2013)

The report done by Halevi et al. touches on a few of the aspects of phishing that are rarely studied and yet might hold some of the keys to a truly effective understanding of why phishing is both prevalent and successful. They have looked into the correlation between personality indicators and successful phishing attempts in an effort to understand the psychological nature of phishing attacks. By linking the “Big Five Framework” (a framework of 5 distinguishable personality traits that all humans are believed to exhibit) as found by the NEO-PI FFM test to the success rate of multiple phishing attempts, they have discovered some leading indicators as to the type of individuals who would typically click on and fall for a phishing attack. Their results indicate that women tend to be more susceptible than men to clicking on links in emails that show common signs of scams including, but not limited to, the usage of misspelled words, a promise of free products, a created sense of urgency, and the usage of links that do not match the actual destination as indicated by hovering over the link. Their results showed that 14% of men clicked on the links, while 53% of women clicked on the same links. While their results clearly demonstrate that gender is an important factor, I would like to see further research on the matter as their sample (17 women and 83 men) was small for this type of research. Other indicators such, as pessimism and neuroticism, showed interesting results and warrant further research to assess their viability.

What makes the research done by Halevi et al. interesting is that they are one of the few examples of researchers looking outside of the mainstream research being done on the matter. While most of the research community appears to be focusing on the detection/prevention and mitigation of phishing attacks, very few are looking to other factors that might give the security industry a more complete picture of phishing and why it is so successful. I believe that answers to solving the problem in a more meaningful way, at this time, lie outside of what mainstream research is doing. For example, phishing has recently (within the last 5 years) become a solid focus for most IT departments and security advisors and yet the success rate and usage of this technique among attackers has continued to grow exponentially. In 2013, security provider Kaspersky Labs estimated that phishing attacks had grown from 19.9 million in 2001 to an incredible 37.3 million in 2013 (Labs 2013). Attackers are using this method more frequently because it is so successful. Something appears to be missing when it comes to defending systems against this threat and I believe that at least part of what is missing resides in the peripheral knowledge that, in part, includes the research into the psychological nature of the victims. Additionally, the psychological nature of the attackers would be an important aspect to gaining a greater understanding of how to combat this problem, as it would give security professionals a deeper understanding of the motivating factors behind their attacking counterparts.

3.4 Self-Control and Criminal Opportunity: A Prospective Test of the General Theory of Crime (Longshore 1998)

An example of psychology, as it pertains to phishing, and how it might be studied can be taken from the general theory of crime. Research done by Douglas Longshore, for example, was aimed at confirming the positive association between self-control and criminal activity. What he found was that a positive link between the two did exist but seemed to be more of a factor in the presence of a second variable. The second variable (opportunity) added nearly 4% of explained variance when used in conjunction with self-control. While these indicators were only small factors in determining the likelihood of an individual engaging in criminal activity, they could be used in conjunction with other studies to create a more holistic view of who would be likely to commit such crimes and how they might proceed in doing so. The presence of this type of information in an organization's security center could give them greater awareness of an attack vector, as well as more opportunities to detect attacks at an earlier stage in the process. For instance, knowing that self-control is linked to crime and opportunity would also tell the security professional that the same trait could be exploited in the firm's staff. Those employees with low self-control are more likely to take advantage of other perceived opportunities, thus clicking on links for free products and/or promotions that, when selected, redirect the user to a "black hole sight" or other form of fraudulent attack. With that level of foresight, the security professional may seek to amend policies dealing with the treatment of such advertisements or use technologies as firewalls or content filters to prevent the intended target from ever receiving such a request.

Additionally, from this study alone we see that there is a positive correlation between crime and opportunity. In the presence of opportunity, attackers can leverage their skill sets to penetrate networks and carry out mission goals to the detriment of the victim organization and its constituents. While security professionals can do little to affect the skill sets of the attacker, they can do much to affect the opportunities presented them. From preventing certain types of emails from entering the system, to limiting the type of information about employees and systems that are published on the company website, opportunities to send malicious correspondence can be limited, effecting the attack surface of the company or organization. I believe that further research into the connection between the human psyche and phishing can produce similar results, leading to better and more encompassing security measures that prevent such attacks.

3.5 How Could I Fall for That? Exploring Phishing Victimization with the Heuristic-Systematic Model (Zhang 2012)

The author of this article proposes that further research should be done that focuses on heuristics of a typical phishing attack. He further proposes that the “Heuristic-Systematic Model” (HSM) be used as the primary base for such research. Again, interest is shown in the peripheral research areas beyond detection and mitigation to further explain why and how these types of attacks are so successful and yet that is where the author leaves us. This article is simply a proposal describing the need for further research that describes a unique point of view. His is one that seeks to explain tactics used against the human condition. Others have sought to explain ways in which they could overcome these human characteristics with technology, but few have sought to understand their nature

first. Zhang points out through the HSM model that people process information and make decisions based upon that information very differently from person to person and from instance to instance. These decisions can be influenced by many factors, including what the author refers to as a “sufficiency threshold” or point at which the decision maker feels they have enough information to effectively make a proper decision. This threshold, however, can itself be influenced by the perceived importance of the decision to be made and the urgency of the required response. Of this Zhang states: “some phishing attacks attempt to exaggerate the urgency of the situation and press the message recipients into action as soon as possible, thus suppressing systematic processing. Under such circumstances, message recipients usually have to rely on heuristic processing to make decisions, which are often incorrect due to bogus heuristic cues.” This example clearly depicts an attacker cueing in on and exploiting a human weakness rather than a system weakness. This begs the question; what other human weaknesses are easily exploited and how could an attacker capitalize on such a weakness to more effectively phish their way into a system?

This proposed study and others like it ask questions outside of what has become the typical approach to social engineering and, more specifically, phishing. It is questions like these, in conjunction with the plethora of research dealing with detection and mitigation, that would lead the information assurance community toward a more complete model for understanding the threats, risks, and mitigations associated with complex phishing attacks. Understanding the heuristics and behaviors of people and building a model around those principles creates a more forward looking and proactive model for defense. Conversely, the accepted model and research of the day, where

efforts are directed primarily, if not wholly, toward detection of and responding to phishing attempts is a reactionary model. Organizations that react rather than proceed in a proactive manner are always trying to catch up to the criminals who seek to thwart their defensive structure, while those who are forward looking in their strategies set the terms for engagement, often at a cost too high for most attackers.

3.6 The State of Phishing Attacks (Hong 2012)

What author Jason Hong has discovered is that an effective strategy against phishing attacks will be more than any one tool or approach, though he lists a few different tools in his article. His point is that combating attacks of this nature will, like all other cyber security related tasks, require a defense in depth (DID) approach. DID represents the idea that layering defenses against any given attack allows for a more robust defense against that same, and many other, attack vectors. For example, in medieval times, castles were protected in the same manner. High walls were not the only layer of defense that an attack force would have to penetrate. Often times an expanse of water surrounded the castle walls, sentries were posted at key points around the castle's perimeter, and guarding forces remained inside and atop the walls to combat forces that tried to enter. So it is with IT security today. While many organizations are looking for the silver bullet to stop phishing attacks from being successful in their organizations, the reality is that a more practical approach, with a greater ability to prevent such intrusions, likely exists in the idea of layering multiple defenses upon one another. Hong illustrates this idea behind three principles he identifies as key to better resist such attacks. (1) Make it invisible. His approach to this principle is to create defenses that are imperceptible to the end user.

Using filters on emails and enabling phishing filters in web browsers are two examples of ways that administrators can make protection invisible to the end user. (2) Create and use better interfaces. In his research, Hong discovered that many of the interfaces used today try to implement warning features to alert users of potential dangers in their online activities. However, these features are often ignored by users for various reasons. (3) Train Users. What makes phishing attacks unique from many other types of technology related attacks is the fact that the vulnerability being exploited is not a technological vulnerability but rather a human one. While some technical controls can, and should, be used; a DID approach signifies that safeguards should be put in place at the human layer. Thus, training should be a relevant part of the defense strategy. Much research and many approaches to training users have been discovered, predominantly by researchers at Carnegie Mellon University. More recently they have proposed a gaming approach through the use of a tool called “Anti-Phishing Phil” and an embedded approach called “PhishGuru.”

A layered defense model is both rational and proven to be successful in other areas of information assurance. Hong’s approach is logical and easily understood, but I’m unconvinced that it encompasses enough breadth and depth to make it a truly effective defense framework. The three principles outlined are a great start to a more inclusive approach. When the research done by (Hutchins, Cloppert and Amin 2014) , regarding the “Cyber Kill Chain”, is merged with that done by Hong, I believe we can approach a deeper and more inclusive framework for defense. Besides interfaces, training, and invisible safeguards, the framework ought to include research done by the attacking party, delivery methods, psychological motivations, human vulnerabilities, and intent.

Further identifying factors such as these may allow defenders to provide more layers of defense, shutting doors at multiple levels of engagement, and increasing the cost of the attack to the point where utilizing such attack methods becomes non-viable to many, if not all, parties.

3.7 Who Falls for Phish? A Demographic Analysis of Phishing Susceptibility and Effectiveness of Interventions. (Holbrook, et al. 2010)

In one of the few examples of secondary research on the threats of phishing, Holbrook et al. seek to explain, in part, the relationships between human factors such as age, gender, etc., and the likelihood of a successful phishing attack. The questions they sought answers to stemmed from previous secondary findings gained from other studies on the topic of phishing. While other studies led researchers to theorize that there was a relationship between human factors and the likelihood of a successful attack, Holbrook et al. designed their study to prove or disprove these connections. What they found was a confirmation of certain connections and yet more questions regarding the explanations of those connections. Previous results indicating that gender was a significant factor in the likelihood of a successful attack was confirmed. Their research showed women are significantly more likely to click on a phishing link or divulge personal information during a phishing attempt. Further, they confirmed a slightly weaker connection to the age of an individual and their likelihood to be successfully attacked. Their research indicated that individuals 18-25 years of age are the most likely to fall victim to such an attack. The final indicator confirmed through the study that training and education of the subject matter also decrease the likelihood of a successful phishing attack.

While further questions were raised and limitations to the viability of the research were admitted, their research is interesting in the fact that secondary indicators could be used to develop better detection tools and better training programs. Understanding the aspects of human characteristics that determine the success rate of directed attack may prove to be key in limiting and making attacks such as phishing, spear phishing, and whaling too costly for nefarious actors to undertake. If that were to become a reality, then one would expect the growing rate of these types of attacks to decline, thus saving institutions across the globe millions of dollars annually. The human factor is just one of many factors that have yet, to the best of my knowledge, to be encompassed in a total defense structure combatting social attacks, or more specifically, the many variants of phishing.

3.8 School of Phish: A Real-World Evaluation of Anti-Phishing Training

(Ponnurangam Kumaraguru 2009)

What makes phishing and its variants different from many of the other cyber-attacks is its relationship to human vulnerabilities? Because the human brain uses so many factors, including context, timing, urgency, compassion, etc. to rationalize a decision before one is made, it is increasingly difficult to predict and prevent all attempts to appeal to those senses in order to coerce information from the victim organization. In an attempt to combat the efforts of social engineering to include phishing, firms have deployed trainings and awareness efforts among their employees. Their intent is to empower employees to make informed decisions before clicking on suspicious links or traveling to suspect websites in the hopes that they would recognize phishing style attacks and avoid them. Technologies have been put in place by web browsers and security applications to

warn individuals of the recognized dangers of certain actions and web sites, but in the flood of visual cues these have become largely ignored by the populace. With this in mind, Kumaraguru et al. have created a training tool that utilizes real world examples in an attempt to better help personnel be cognizant of the dangers as they appear.

Phishguru is a tool that repeatedly, over time, sends out communications containing common phishing tactics to all individuals within the organization. If the targeted individuals fail to recognize the signs and clicks on the links contained therein, they are directed to an interactive training site that trains and teaches the victims to better recognize threats such as phishing, spear-phishing, and whaling. By teaching “in the moment,” the designers of the program hoped to encourage the long-term retention of such a training. Their research showed that there was a significant benefit to the long term retention by those who used the Phishguru system, indicating embedded training systems that train users through real life situations have longer lasting effects than those of traditional training programs.

When looking at a program such as Phishguru in a layered defense approach we see that systems can blend together facets of technology with facets of human cognitive thinking. Training users is one of many ways that an organization today can combat phishing attempts from a personnel perspective. Because it is difficult, if not impossible, to secure information from the human vulnerability through the use of technology alone, training and other non-technical approaches must be utilized in conjunction with their technological counterparts to create a more robust approach to securing today’s networks and information. Creative and effective training programs should be created and utilized. Novel approaches that consistently teach while presenting the information in a

memorable fashion ought to be designed and implemented across organizations. Finding ways to help users recognize and reject attempts to gain information, or access to systems, is crucial in aiding the technological innovations that seek to unobtrusively protect users from phishing.

3.9 Data Shield Algorithm (DSA) for Security against Phishing Attacks (Ram Avtar 2011)

Social engineers, particularly those who use phishing, utilize multiple methods of deception to accomplish their goals of gathering private information from users. Some of these deceptive methods include trickery where the user believes they are pursuing a legitimate path, when really they are being led into a snare set by the attacker. An example might be an email with an embedded link sent by what looks like the victim's banking institution. The warning from the bank might read, "We've noticed irregularities in your recent banking transactions and fear that your banking credentials have been compromised. Please update your username and password immediately to avoid a lockdown of all of your accounts." A link to login to the victim's account is also provided; but what the victim doesn't know is that the link actually takes the user to a falsified web page that is a replica of the bank's login page. The difference is that when the victim enters their credentials their information is harvested for the use of the attacker.

Phony links are a primary way in which a phisher exploits the human element of a technological system. While many safeguards and training programs have been employed by Internet Service Providers (ISPs), web vendors, and organizations across

the globe; many still fall for such attacks or ignore the warning signals provided them. Automated processes that work behind the scenes add another layer of protection to guard against this rampant threat. This appears to be the logic behind the work done by Avtar et al. Their Data Shield Algorithm (DSA) analyzes the hyperlinks that are so commonly used in phishing attacks. Their ultimate goal is to specifically prevent distributed denial of service (DDOS) attacks that are a result of such phishing attempts. Their justification is that a host based DDOS is the inherent outcome of a user that has successfully been phished. This logic is true in a manner of speaking, but greater implications including loss of data, integrity, and confidentiality also accompany these types of attacks. Their solution is to send the hyperlinks against a battery of tests that analyze where they claim to go, where they are actually going, what their Domain Name System (DNS) claims to say about them, what it actually says about them, if their DNS is of reputable nature, and if it appears on any unclean lists. The end product is a sanitization process of all hyperlinks within a website or email correspondence. When the links are flagged as a possible phishing attack, the user is warned or in some cases the site is blocked from ever reaching the user.

Algorithms like DSA could be useful as a standalone product or model but even more so as an element of a layered defense model. The idea that an organization, from the beginning, would lay out a plan to protect users through layers of training, technology they interact with, and technology that works behind the scenes makes a more robust defense structure than one that relies on training or technology alone. DSA, or products similar to it, could be utilized as one of those unseen layers of defense protecting users and networks from data loss and compromising threats.

3.10 Special Publication 800-12: An Introduction to Computer Security – The NIST Handbook ((NIST) 1995)

The United States Federal Government views training and educating their employees as an essential portion of securing their vast resources and networks. The National Institute of Science and Technology (NIST) is primarily responsible for creating frameworks and policies for the use and maintenance of technology utilized within the civilian federal workspace. As such, their Special Publication 800-12 is an introduction to computer security and outlines various principles that civilian agencies; i.e. the Department of Agriculture, Department of the Interior, Department of Health and Human Services, etc. should follow to maintain a minimal level of security within their respective organizations. Chapter 13 of this same publication outlines the essence of training and education programs that would be created in each agency. And though this is created specifically for federal use, many of the concepts translate nicely to industry, education, and all other sectors or groupings of people and resources.

Of particular interest is how NIST breaks down training into three compartments; awareness, training, and education. In NIST's example, awareness is the "What" attribute that informs employees about the existence of security threats and the need to be on guard against such threats. Training represents the "How" of information security giving employees the knowledge or skill sets to avoid and/or mitigate threats as they are encountered. Education is reserved for the "Why" an attacker would seek to infiltrate a firm's systems or exfiltrate the same firm's data including passwords, information, or resources. Each compartment is also associated with an expected level of impact made upon the individual (short-term for awareness to long-term for education). Where I feel

that NIST's document is weak is in their description of whom should receive this compartmentalized training. They dictate that general users ought to receive "Awareness" level training while "Education" style training is reserved for specialized or security staff only. I would say that caution ought to be exercised in the types of "Education" style training that are offered to all users, as to not drown them in specifics and jargon; but if long-term benefits of training and education are the goal then the "whys" of information security should be included in educating the workplace. In my opinion, the "whys" are going to have a real impact on making the company culture one with underpinnings of security in mind.

3.11 Lexical Feature Based Phishing URL Detection Using Online Learning

(Aaron Blum 2010)

Detection of phishing events is a major challenge for firms and organizations seeking to protect their users from direct attack and yet it has proved to be a nearly impossible one. Many different researchers have come up with different methods and different indicators to key on. Each arguing that their method, algorithm, or process yields tremendously positive results and each admitting that more research is needed to make it better. The work done by Aaron Blum et al. is one of those types of research. Their method is a two-fold process focusing on the uniform resource locator (URL) of a website or link and the contents of the page the user is directed to. Their hypothesis is that by utilizing these two indicator values in combination with machine learning techniques, they should be able to accurately predict the likelihood of a site being malicious leading to phishing attacks.

The first part of the detection process begins with analyzing and comparing the values contained within the URL of a given site. This is accomplished by breaking the URL apart into values the author refers to as “tokens.” These tokens are then analyzed against a list of features called “bags.” When processed through this list it is expected that the likelihood that a site is malicious could be determined and used as part of the detection process.

While comparing the values of the URLs with a list of threatening features is a good start to the process, the second part of the detection process solidifies the determinant value. In the second step, the site page itself is analyzed and compared against other known phishing sites looking for key indicators often used by attackers. One of the most interesting aspects to this portion of the process is the author’s use of “Deep MD5 Matching.” A hash value of a page is often used in detection engines to identify malicious pages. In “Deep MD5 Matching” that page is broken up into its various parts and pieces (images, text, banners, etc.). Each piece is then hashed and then aggregated to gain an overall or averaged score. The advantage in using this technique is that it is still possible to identify pages even when the attacker is smart enough to make small changes to the site in order to avoid detection.

Detection is a key factor in preventing a successful attack on people within the organization. Tactics to perform such detection comes in two basic forms: human and technological. This technological approach uses some novel ideas to not only discover attacks based on signature based detection, but also has some limited ability to discover previously unknown or unseen attacks. Breaking the process into small parts for analysis, rather than analysis as a whole, has the advantages of rating the degree of

likelihood of an attack rather than a concrete yes or no determination. Creating a valued degree allows professionals better opportunities to see what is falling through the cracks, thus enabling them to better tune their tools and approaches.

3.12 INFORMATION SYSTEMS SECURITY: A COMPREHENSIVE MODEL

(McCumber 1991)

John McCumber's framework for information security has long been a guide for securing data in its various states. While the field of information assurance was still in its infancy in the early 1990's, McCumber saw it as a necessity at the time and for future generations. He also saw it as more than the systems on which the data reside, and realized that it wasn't the systems that need to be protected, but rather the data held within. This aspect creates the fundamental basis on which his model was created. McCumber's model can be thought of as a 3x3x3 cube where the height is represented by the goals of information assurance, namely "confidentiality", "integrity", and "availability" or the "CIA triad." Its length represents the three states in which data can exist. These include "storage", "processing", and "transmission." Finally, depth is represented by the mitigation tactics utilized to achieve the CIA triad for each of the three states of data. McCumber proposes that all mitigation strategies really fall into three buckets: "policy", "education", and "technology." From the culmination of each of the previous labels, a framework for securing data is formed. This matrix demonstrates that for each state in which data might reside there are three goals to achieve in order to secure it. For each of the three stated goals (confidentiality, integrity, and availability) there should be three layers of defense to protect it. Notice that I said three layers and not

three tools or items. The implication here is that there may be multiple tools for each layer that may need to be implemented in order to provide adequate protection upon the data. For example, a certain database may require antivirus, access control lists (ACL), and an intrusion detection system (IDS) in order to protect the data's confidentiality. All three considered to be a technical control, neither one by itself adequate to fill the technology role sufficiently.

John McCumber's model is really, in its most basic sense, a layered defense model. It assumes that technical defenses are not enough, and thinking of just one state of security is ineffective. The model replaces the single-point defense structure with a holistic approach to maintaining security over the data within the organization with multiple layers from multiple vectors. He has created a structure where awareness and education is thought to be equally as important as firewalls and antivirus. Policies and procedures stand as key factors just the same as an IDS or ACL. It is this layered approach that is most relevant to the work I am doing. It is the idea that we as security professionals need to stop plugging holes in our defenses and start building landscapes of protection against a myriad of threats.

3.13 Psychological Safety and Learning Behavior in Work Teams (Edmondson 1999)

Psychological safety is the idea that any given member of a particular group would feel comfortable enough to express their thoughts and ideas, and/or take interpersonal risks without the fear of punishment or retribution. Amy Edmondson's work in this report shows that there is a quantitative and qualitative link between psychological safety within

a group and the effective output of that same team. She verified her hypotheses by studying multiple teams with various team structures. It was discovered through comparison of the teams that each had a varying degree of team cohesiveness, effectiveness, structure (team leader interaction and leadership approach), contextual support (from leadership), and psychological safety. With these team characteristics as variables, Edmondson sought to determine the relationship between them as they affect team production. What she found is that in teams where contextual support was given and where members felt comfortable with sharing their ideas and making the occasional mistake, the team was empowered to a certain degree and produced at a higher rate in a more efficient manner. Teams that felt pressure to only pursue goals and approaches that were consistent with their thoughts of “what management wanted,” often were slower to achieve their goals and provided less effective results. To illustrate the juxtaposition of the two extremes, two examples stood out to me from Edmondson’s research. 1) A member of a stain team explained: "This is going to sound very childish, but let's say I just did a part and I got drips on it. Now, if they [those next in the production process] told me I got drips on the edge, I say 'thanks' and then I'm glad I can get these drips off. Where it used to be, when that happened, we'd just try to find something wrong that person did we'd keep an eye out for it! It wasn't to be helpful, it was to bring them down to your level or something like that.... Now, we don't think anything of it. We just fix it." Clearly something has changed that now allows the group to work in a more cohesive manner with the goal of producing a better product. A winning combination for both the employee and the company. 2) In contrast, a different group explained their experience with this remark: "We went down a lot of blind alleys.... We would go down a path for a

while, develop details, then abandon it. Each path represented time wasted...." The lack of experimentation behavior appeared to be related to the team's concern that they had to produce a certain solution that "management" wanted.” This team exhibited little psychological safety and thus wasted a lot of organizational time and resources to arrive at their solution.

Company culture and the reporting aspect plays a large role in the effectiveness of a security team and the overall security posture of a company. Employees need to feel it appropriate to share their security concerns with the IT staff. The security teams need to know about the current threats to their environments as they are happening to adjust and reposition their defenses as needed. This is especially true in a layered defense model. If a phishing or similar attack were to make it through the other defenses and still land at the doorstep of the user, the security staff would need to know about it in a timely fashion so that they could react to and mitigate future attempts to perform the same attack. To do this, employees need to feel a level of group cohesiveness and psychological safety that will encourage them to act as participants in the overall security efforts of the organization. Organizations, especially the leadership, should place a premium focus on creating a secure culture for the benefit of not only the organization but also for the benefit of the employees contained therein.

3.14 On the Folly of Rewarding A, While Hoping for B (Kerr 1975)

Research done by Steve Kerr on the follies of rewards systems in business, government, society, etc. comes from the 1970s and yet rings true today just as much as it did back then. The truth of the matter is, that while reading his article I was under the assumption,

and nothing in the article lead me to believe otherwise, that it was written post 2000. I only mention that aspect because it is obvious that society has not gained much ground in the area of aligning our reward systems with what we profess to desire from our employees, politicians, doctors, and other members of society. Kerr's hypothesis is that we profess to demand certain qualities or actions from individuals and truly expect to receive that action from them and yet, for some reason, our reward systems push the same individuals away from the stated desires. For example, Kerr uses politics to accurately describe his position when he speaks towards how elections are won. Each candidate seeks approval from his/her constituents in order to gain their subsequent vote in the elections. As citizens, we seek to obtain operative goals from our candidates including the details of their strategy like how they will fund their goals and timelines, yet candidates are punished for doing so. When details are given, they face a barrage of criticism from the media and the voting base, often leading to a loss of votes. The candidate who stays more general in their campaign messages tend to face less criticism and appeal to a broader base of individuals. The point is that we ask one thing of our candidates but reward another. Kerr points out this folly in many aspects of our societal system including; sports, politics, medicine, business, even orphanages.

For the human aspect of a defense layered approach to become successful, it will inevitably include some form of a reward system. This may be in a positive fashion like the recognition of an employee who successfully reported a dangerous and fraudulent email, or it could be in a negative fashion. For example, failure to attend the mandatory training and awareness meetings could result in the withholding of a raise in the employee's salary or damage their opportunity for promotion. Irrespective of which

system is used, or maybe both are, any course that requires the cooperation of humans is likely to include a form of reward system. That being the case, it is important to discover ways to ensure that those rewards align with the organizational goals. The firm must ensure that if they truly want and expect individuals to report fraudulent emails, even if that reporting happens after the employee has already clicked on the link in question, that they refrain from punishing individuals who do so. While punishment may send the message that the employee should not have haphazardly clicked on the link, it also tells employees that it may be in their best interest to hide the email and hope for the best, an action that is incongruent with the goals and aspirations of the firm and its security posture.

3.15 The Psychology of Security (West 2008)

Often overlooked in favor of the more glamorous dazzle of the blinking lights that create technical defense tools, the psychology of security plays a unique and important factor in the overall security of a file, system, or firm. Ryan West has researched the risk and decision making process of people in various fields of study including medical mistakes and information security. His research here asks questions beyond the technical factors of maintaining a secure system. He looks at how individuals make decisions and weigh the risks against the rewards. He explains the phenomenon of users who feel like “it would never happen to me” because they do not feel like they are at risk. This becomes more of a factor when they are aware of security measures put in place to protect them. They assume a posture of accepting additional risk, feeling encouraged to do so by the protection of tools coupled with the sense of inherent safety that they feel.

West also speaks towards the term “cognitive miser,” where he explains that an individual has limited time and processing ability in each day which requires them to make decisions on what they will dedicate their full attention toward. This is found in the fact that we often skim material sent our way, picking out the important facts rather than reading the entire body of text. From a security standpoint, alerts are often ignored and security policies glanced at rather than read completely for understanding.

It becomes very clear to see that psychology plays a major role in the security aspects of an organization, whether the firm addresses it or not. More often than not the psychological, and/or human factor, is ignored, replaced instead with another technical tool or program designed to mitigate the subjective aspect of human nature rather than addressing it. A complete defense framework should, at the very least, address this aspect and optimally incorporate not only the inherent weaknesses but also the associated strengths that can be gleaned from the human element of the organization.

3.16 Technical Note – Analysis of a Layered Defense Model (Walter R. Nunn 1982)

Originally written for the analysis of a military defense structure, Nunn et al. introduce the concept of predicting the amount of penetration through each layer of defense using the Markov chain principles. The Markov Chain is a mathematical principle that transitions based solely upon the process that precedes it. In a layered defense system, each layer would affect the outcome of the layer immediately following it. The current state is dependent upon the previous state but all other subsequent previous states would be irrelevant to the current. For example, a military base might fortify their borders with

a fence structure, a watch tower inside the fences, and a combat unit behind the towers. Let's say that of the ten troops in an enemy unit, six make it through the fence. Of those six, two make it past the watch tower and are met by the combat unit behind. The idea here is that the friendly combat unit faced two enemy soldiers because the watchtower found four enemies and not because the fence stopped four prior to that. Remove the fence, and the numbers all change. Add another layer of defense, say a mine field before the fence, and you get yet another result, but the fact still remains that the number of troops faced by the internal combat unit is predicated upon those that make it past the watch tower.

The fight against a digital attacker is not unlike the scenario described above. An organization or government agency might be the intended victim of hundreds if not thousands of phishing attacks each year. Using the Markov Chain principles, we can numerically estimate the likelihood of a penetration through each layer of the defense structure. Consequently, we would also expect to see a decrease in the overall likelihood of a full penetration as additional layers are added, or as we increase the effectiveness of the existing layers. Either way, it is demonstrated mathematically that layers of defense prove to be more effective against attacks than any single defense structure by itself. I would further hypothesize that layers from varying aspects of defense are more effective than layers designed around a single aspect. In other words, involving psychology, technology, and training is a more robust model than the utilization of technology alone.

3.17 Defense in Depth: An Impractical Strategy for a Cyber World (Small 2011)

A defense in depth strategy, or layered defense, is not a new concept. In fact it has become somewhat of a standardized model for overall defense in many realms including: military, fire suppression, engineering, and cyber-security. What the author, Prescott Small, points out in this paper is that defense in depth has, over time, become something it was never intended to be. First, he points out that this strategy was originally designed as a military tactic to protect against kinetic threats. These are threats that exist in a physical manner (soldiers, tanks, bullets, etc.). Applying this model directly to a cyber-world made up of ones and zeroes will not produce the same beneficial results. Small further declares that the layered defense model has become a one trick pony if you will, meaning that the so called layers are so homogenous in nature that while we call them separate layers they are really only acting as a single layer to bypass. Small calls for a rethinking of today's defense in depth strategies to include a broader spectrum of considerations and approaches to layers of defense. He further states that one of the layers of defense should encompass a more wholehearted approach to communicating the sterilized attack data of discovered attacks across industry and government. He explains that attackers have been using this sharing of information for years to their benefit. If defenders and security professionals were to adopt the same culture of information sharing they could expect a dramatic increase in the effectiveness of their programs.

Small further alludes to the idea that a layered defense model needs to include more than a defense in depth approach but also should include a defense in breadth aspect. Layering the same types of defense structures are a less effective way of guarding against attacks and gives the false sense of security that attackers would need to get through the

many layers set forth before gaining access to the internal network when in reality the attacker may only need to undermine one, or a few, basic principles to find success.

Coupling defense in depth with the defense in breadth ensures that the layered defenses vary enough in type and implementation to make it more difficult for attackers to enter.

An attacker may have to get through a firewall or antivirus but they would also have to break through a trust relationship or a psychological feature as an example. Multiple layers that leverage different mechanisms for defense.

This idea that layered defense is more than just grouping a bunch of tools together as speed bumps along the way to network access is a key factor in designing an effective defense structure. Phishing attacks are no different in the sense that a single technological tool or a single training structure warning against the attacks are not sufficient. The layering of one over the other along with further layers from even different approaches can help to mitigate the threat in a more effective and efficient manner.

3.18 Observations on the effects of defense in depth on adversary behavior in cyber warfare (Dorene L. Kewley 2014)

Seeking to justify their layered defense hypothesis that “adding layers has at least a cumulative impact on adversary work factor” Kewley and Lowry set up an experiment where a fictitious network was created with a high impact target. Protecting this target were multiple configurations, one building upon the other, that acted as a layered defense structure. A red team, group of hacking individuals who act as attackers, was designated and give the task to acquire resources from the target. They did this in iterations where

the first iteration was with the first configuration only, the second iteration commenced to include the first and the second configuration, and so on and so forth. What was discovered was surprising. Multiple layers of defense, as designated by the layering of configurations, did not prove to increase the work load on the red team. In this case, because the configurations were dependent upon one another, a workaround was created by the red team that allowed them to bypass multiple layers at once. The lesson learned from this experiment was that defense in depth alone is not sufficient to expand the work required by an attacker. Instead the more effective model contains both defense in depth as well as defense in breadth.

Defense in breadth contains the idea that one not only defends at layers against a particular attack class but also defends with layers against multiple attack classes. Attack classes can be anything from a poorly configured device, to a vulnerable open port, to an employee's willingness to click on a link that promises them free gifts in exchange for their support. It represents the multiple attack vectors that an attacker might leverage to gain unauthorized access to a network or resource. So when a defense structure is set up the security professional must include creative thoughts of how resources might be compromised from a myriad of different approaches. Then a layered defense in depth strategy can be placed upon those separate vectors.

This experiment highlights the approach toward a phishing layered defense framework. It shows that there is no secret sauce that acts as the magical elixir preventing penetrations from this kind of attack. No one piece of technology or one method of reward system is sufficient for true defense. The best approach to date is one that considers all areas of the firm; both technical and humanistic, both policy oriented and

digitally protected. Human, machine, and written guidelines working together to form a stronger and more impenetrable shield against the threat of phishing and all of its variants.

3.19 Security education against phishing: A modest proposal for a major re-think (Lacovos Kirlappos 2014)

Security awareness and education has long been held as the recommended method for transferring what the organization or firm expects a user to know to that user in a broad sweeping action. Kirlappos et al. used their research to find out if the spewing of information to users is really the best method of training and raising awareness. What they found was that not only do users engage in risky behavior, but at times they seem to seek it out in hopes of gaining a greater reward. The “Need and Greed” principle was actually researched by another pair of researchers, Stajano & Wilson, but the premise is that an individual will accept greater risk in the presence of opportunity for greater reward. This principle is evident in other areas of society, most notably in the world of finance (i.e. the stock market).

The gap in perception of risk is one of the variables that outlines the differences between the user and the organization. What is understood from this is that training and awareness should be tailored to fit the users, or at least with an understanding of the users’ point of view, otherwise, the users are likely to reject the training or not understand its purpose. For example, a certain company may wish to implement a new strategy for selling their brand of widgets. They would then offer one training to the sales staff outlining the features of the widget and its bestselling points. To the manufacturing arm,

the company would offer a different training all about efficiency and quality of work.

Both trainings point the workers toward the goal of selling more product, but they are tailored to the specific audience. The sales staff would gain very little from a speech on efficiency while the operations staff would fail to see the link to them from a training on selling features. So it is with training users on phishing. Kirlappos et al. propose that if it is the nature of the user to accept certain levels of risk, then our message should not be to avoid any and all risky sites but rather to better assess the risks against the benefits of using such a site. Doing this, along with training users on the key indicators of a legitimate site, as well as the ways that such indicators could be faked, increases the likelihood that the end user will accept and remember the trainings, that make them more effective.

While their research shows a disconnect between perceived risk from an organizational standpoint and perceived risk from the users' standpoint, I am not sure that I wholeheartedly agree with their proposed solutions. I do believe that a tailored message in the form of a training is necessary, but many of their solutions fall outside of what I believe the typical user will accept and likely do. For instance, in an example centered on website payment options, the author talks about increasing the awareness of protection by including tutorials on phrases like "how am I protected." While I agree that detailing the ways that an online retailer is protecting the user's information is good, it is also unlikely that the vast majority of users are going to take the time to read it. Many of the proposed approaches toward addressing the training issues seem to be laborious or time intensive to the user. This aspect seems detrimental to the acceptance of the training by the user. End users are already under massive time constraints and are already filtering out all but

what is most important to them at the time. Under this assumption, it makes sense that any training or awareness tool needs to be concise, to the point, and attention grabbing. It needs to be tailored to the user and not presented with a bunch of frivolous filler. Finally, it needs to be impactful. The message should inspire some sort of “ah-ha” moment that will reside with the user.

3.20 Spear-Phishing Email: Most Favored APT Attack Bait (Team 2014)

This whitepaper demonstrates the favored use of phishing, specifically spear-phishing, by attackers seeking to establish a foothold in an organization’s network. This type of attack, commonly referred to as an Advanced Persistent Threat (APT), is a targeted attack which is why the use of spear-phishing, a targeted style of phishing, is so prevalent.

TrendMicro reports that from February – September of 2012, 91% off all target attacks reported involved spear-phishing emails. This astonishing percentage points to the fact that phishing is one of, if not the most, dangerous threat to a firm’s security. Though phishing is a consideration in most security structures today, it is rarely treated as the dominant threat that it is. Heavy time and consideration is given to items like access lists, firewall rules, and antivirus settings while phishing training and detection toolsets are largely ignored. When due diligence is given to social engineering and phishing, it likely comes as an add-on to a web browser or a minor discussion in a security training program. In reality this type of attack should be given a spot at the head of the table of security concepts, especially when one considers the ease of the attack vector.

TrendMicro further went on to discover that nearly half of the email addresses used in the spear-phishing attacks discovered were only a Google search away. This means what

within minutes, the attacker had a selected target along with a means for delivering a package containing the exploit. A package, that without proper filtering, and even with proper filtering if the attacker is good at what they do, moves right on through the company's defenses including the company's firewall.

This report demonstrates the need for proper consideration of phishing attacks in all of its variations. Furthermore, it demonstrates the need for a framework of defense against phishing as it is likely just the tip of the spear (no pun intended) leading to an APT in which the internal network is breached and attackers have a persistent presence on the inside.

3.21 Understanding Scam Victims: Seven Principles for Systems Security (Frank Stajano 2011)

A single phrase from this article really embodies what I find most interesting about it. The line reads: "an attack is possible because the designers (i.e. security systems designers) thought only about their strategy for responding to threats, without anticipating how real users would react." This seems to be consistent with many of the technically driven, and even policy driven, approaches to dealing with phishing today. Often times, security teams are enamored with detection and mitigation of phishing. They come up with policies and technological tools to combat the threat, truly believing that they are doing all they can to prevent such an attack, but they fail to consider the person sitting at the keyboard reading an email that was specifically crafted for them in an attempt to prey on their desires and vulnerabilities. Stajano and Wilson ask the question: is there a common set of vulnerabilities or tactics that an attacker might use against a human

victim? What they found was a set of seven core tactics used by attackers that seem to be common in all scams: 1) the distraction principle, 2) the social compliance principle, 3) the herd principle, 4) the dishonesty principle, 5) the kindness principle, 6) the time principle, and 7) the need and greed principle.

Each of the seven principles defines a different weakness in human nature that is either innate to the person or which has been acquired through outside social pressures. For instance the herd principle is an example of an individual feeling secure in taking a risky action because it would appear that others around them are comfortable with taking the same risk. However, one that really caught my attention was the need and greed principle. This factor deals with visceral cravings of our own internal desires, be they emotional, physical, drug related, sexual, or otherwise. Of this notion the author states, "If we want to con someone, all we need to know is what they want, even if it does not exist." Of the seven principles listed, this seems to be the one that is the most personal. The most amount of knowledge about the victim would be required but I would venture to guess that, if done correctly, this is the most effective of all the methods. All of us have needs and desires--whether it is a physical object like money, a sexual one, or even the desire to provide a great education to our children--that make us vulnerable. People often take greater risks or ignore warning signs when there exists an opportunity to obtain that which is desired. This research demonstrated that principle through a television show called "The Real Hustle," where a TV crew created elaborate scams based on these seven principles to teach people how to not become a victim. The need for greed factor was often used in conjunction with other principles to extract the intended resource from the victim.

In systems security, a holistic approach to combating any social engineering type of threat would not be complete without a deep discussion of the human aspect of any system. Just like considering the security of a technical system requires an analysis of the existing vulnerabilities to that system, so does the human element require an analysis of the existing vulnerabilities. The seven principles listed here should serve as a foundation for analyzing and placing safeguards to protect the human element. Inevitably, a persistent attacker is going to find a way through the technical defenses put in place and it is going to rest upon the shoulders of the person at the keyboard to recognize the threat and combat it. Seeing the vulnerabilities of the person ahead of time, and addressing them through policies, training, and technology can give the victim an advantage when it becomes time for them defend the network.

Chapter 4: Concept

What I have found in much of my research is that there are a lot of great ideas to address parts of the phishing problem, but these ideas seem to be siloed or encompassing of only parts of the problem. No one has put it all together in a comprehensive, or holistic, structure to address the issue as a whole. From Kumarguru's research on some of the links between human conditions and the likelihood for becoming a victim of a phishing scheme (Holbrook, et al. 2010), to phishing training tools like "Phishguru" that teach users about the dangers of phishing in the moment through interactive training sent after the trainee clicks on a phishing link crafted by the trainer. (Ponnuram Kumaraguru 2009) Further examples include innovative detection engines that utilize different methods and algorithms to detect phishing emails, phishing web pages, and other scams. Some compare URLs as in the research done by Blum (Aaron Blum 2010), while others compare visual cues on the page itself, and still others look at the hyperlinks themselves, sending them through a battery of tests. (Ram Avtar 2011) In all cases the research is limited in its scope to a small aspect of the overall problem. What I am proposing is that consideration be given to the problem as a whole. I am not looking to argue one method over another, or decide whose approach is best when it comes to combating the phishing problem. Rather, I am looking to answer a central question: how can we take all the information and research that has been done and evolve it into a comprehensive and holistic solution for a phishing defense structure? I propose that the best solution lies not in any single tool or school of thought, but in the layering of many of these tools in such a way to address six conceptual areas of focus across two general zones that are discussed later in this paper.

This concept is not about identifying specific pieces of technology, training programs, or policies that represent a ‘holy-grail’ of defense against attack, but rather providing a framework for the design of a defense program. The details of what tools to use and policies that need to be made should be customized to the firm that employs such a framework. This is because every firm is different. From their composition to their primary objectives, they are inherently different and function at different levels. This necessitates that any framework be adaptable to the situation and the organizational structure.

Finally, I recognize that such a framework may need to be adjusted and reconfigured as methods of attack change and assailants find new, and unexplored, methods of attacking the human vulnerability. As such, I call on researchers and professionals to treat this as a “living” framework, one that replaces outdated methods with more effective ones as they become available. For as soon as a defense structure is adopted attackers will seek ways to exploit it or find a work around, and it is under this assumption that the idea of a living framework should be adopted.

4.1 A Sequential Chain of Events

Until recently, breaches in security and attacks of various flavors were considered as singular events. Stuxnet was an attack on nuclear centrifuges in Iran, the 2011 Sony Playstation attack was about harvesting user credit card information and personal identifiable information (PII), and Titan Rain was a massive attack against the US Federal Government that sought to harvest information from key resources. (Thornburgh 2014) As such, the attackers have always been thought of having the upper hand when it

comes to a “cyber-war.” The common thought being that the security enclave of any organization needed to be right 100% of the time, blocking all attempts to attack the resources of the firm. On the other hand, the attacker only needed to be right once. One successful hack through the company firewall or one successful breach of the firm’s security is all it takes for aggressors to be successful in completing their mission. While that train of thought seems logical from a pragmatic standpoint, is it really true? Does an attacker of a high profile target really just sit at the computer one day, send an exploit package across the wire, and begin harvesting information? Is it really just a two-step process: 1) send exploit and 2) harvest information? And is step one where the security team has to get it right?

Even though much of the security profession today understands that the act of penetrating a network is more of a process than an event, they still view and treat the process as one single event they need to mitigate. Often times it is a post mortem response to an attack where an investigation is launched and the perceived entry point is closed, or mitigating defenses and detection tools are placed around it. However, researchers with the Lockheed Martin Company have proposed that we treat any attack as a sequence of events. Commonly known today as the “cyber-kill chain”, their method states that every attack can really be broken down into seven steps: (1) Reconnaissance, (2) Weaponization, (3) Delivery, (4) Exploitation, (5) Installation, (6) Command and Control, and (7) Actions on Objectives. (Hutchins, Cloppert and Amin 2014) When an attack is broken down into these steps they represent seven different opportunities for detection and mitigation. No longer does a security professional need be correct 100%

of the time; instead, they have seven different opportunities to be correct whereas the attacker needs to be correct seven different times.

The implication is that during that post-mortem investigation the security analyst should try to discover, as far up and down the chain as possible, the seven different steps used in the attack. Then the organization would seek measures to counteract further breaches along all seven steps. The advantage to the organization in performing such a task is evident in the way hackers, and people in general, tend to approach problematic issues. If your car starts to squeal when you apply the brakes do you replace the brakes or start over and rebuild the car from scratch? Obviously, we would choose to only replace the part that needs replacing. So it is with attacking a network. When the attacker finds that one door has been closed they don't start over at the "Reconnaissance" phase again, but rather use the knowledge they already have and adjust course from the point of blockage. When a security team understands and acts on this principle, and they use it to close as many doors up and down the kill chain as they can, they effectively force the assailant back to a point where they must start at the beginning. This creates a more robust mitigation strategy that better protects the organization's assets.

A phishing attack, like any other attack, requires many phases before the true goal has been achieved. Always, a familiarity with the organization and the people who operate within it is necessary as they are the primary targets for exploitation. An understanding of the systems used, like email domains and addresses, are required to send the phishing emails. A familiarity with the firm's defensive tools aid in the creation of a successful exploit package. Understanding the human psyche and what makes us vulnerable increases the likelihood that a victim will fall for the scam. Setting up a command and

control node and having a plan for executing the goal, whether it is harvesting information, disrupting systems, defaming websites, or other nefarious intentions, requires research, weaponization, delivery, and persistence. Once again, understanding this principle allows security teams to better set up layers of defense to prevent, at multiple levels, the possibility of a successful phishing attack. While many of my colleagues have been stellar at identifying detection and mitigation strategies for these stages individually, none to my knowledge have recognized strategies that encompass each of the layers along a phishing “cyber-kill chain.”

4.2 Layered Defense Model

The ideas and methodologies behind a layered defense structure are not new, and in fact have been used for hundreds, if not thousands, of years. In medieval times castles were



Figure 1 taken from
<https://en.wikipedia.org/wiki/Castle>

created with multiple layers of defense all designed to keep invading armies out and the inhabitants inside safe. Figure 1 is an aerial photo taken from Caerlaverock Castle in Scotland. This photo embodies what is typically

thought of when we talk about castles and their defenses. From the photo we see high walls

designed to keep intruders from entering. They are made of thick stone to ensure they are not easily demolished, thus exposing the inhabitants. Turrets and perches on top of the walls enable the defenders, typically archers, to fire upon invading armies from their

relatively safe position. The only entry is a narrow bridge leading to a gate that restricts a mass rush on the structure itself. Finally, a moat surrounds the exterior of the castle keeping high rise structures designed to allow attackers to overcome the high walls far enough away from the castle to be useful or effective. Modern physical defenses offer similar methodologies in design and intent. Their physical appearance is different and their technologies improved but the layers remain the same. From barbed wire fences and guard towers to motion sensors and missile defense units, the layers remain in today's guarded structures.

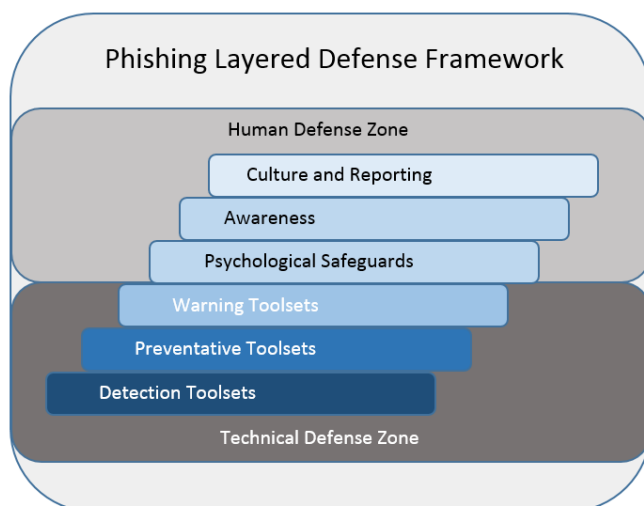
In the modern era, protections and defenses apply not only to physical structures but also to data including proprietary information, national defense data, personal identifiable information (PII), financial data, etc. Once again a layered defense model is commonly used to protect data and the infrastructure that supports it. Firewalls, intrusion detection systems, traffic analyzers, and other tools are commonplace in the typical organization. What is not always considered, however, is the difference between defenses in depth vs. defenses in breadth. Simply layering multiple defense tools on top of one another, if they are all configured and intended to prevent the same single action, is not a sufficient method of providing layered defenses. For instance, implementing a firewall in front of the network and initiating an access control list does provide multiple layers of defense but only against intrusion attacks. They do little to protect the organization from an insider who might send private data out through an email or by copying it to a thumb drive. A more appropriate approach would be to define policies, implement technological safeguards, and introduce user training in a more rounded approach that recognizes the multiple ways and vectors by which an attack can happen.

Defense in breadth becomes more relevant when one considers that “there is no “magic bullet” to protect private networks from attackers. No one vendor, product, or service can protect any environment from every attacker over a period of time. The best that IT Security Professionals can hope for are products and services that are highly effective, and then have overlapping technologies, or defense in breadth, that complement one another. The idea that what is missed with one product is caught by another. (Small 2011) When the breadth of the defenses covers a sufficient base and is complimented by multiple layers of depth, we consider the asset to be under a sufficient amount of protection. Besides the reactive advantages associated with this type of structure, proactive measures present themselves as well. Much like Caerlavrock castle’s narrow bridge and entrance that places the only point of entry at where the castle is best protected, so can well designed defense models direct attacks to where the network is best defended. To this point Kewley states: “defense in depth without defense in breadth can be ineffective for a sophisticated adversary... Appropriately selected depth and breadth layers can cause the attack point to move to more manageable locations where the adversary’s actions can be contained and possibly monitored.” (Dorene L. Kewley 2014) If an attack is going to happen, wouldn’t you like it to take place where you are expecting it and where you are most prepared to defend against it? This proactive benefit is one of many reasons that layered defenses have become a highly implemented strategy in network defense.

While network defenses have benefited from such a scheme, phishing seems to be still approached in a siloed manner, meaning that we are still seeking that “magic bullet” spoken of by Small. Much research and investigation has been done to look into

automated tools that will eliminate the threat presented by such an attack. Automated tools are very effective at protecting standardized automated processes but could be considered only helpful at best when protecting a more inconsistent process like that of the human brain. Our brains do not work in the same manner as a computer. Computers think in 1's and 0's, true or false, a situation is or it isn't. Humans are more rational, spending much of our processing power in the "grey" areas logic. Answers are situational, derived from a mixture of logic, comprehension, and emotion. We rarely, if ever, feel completely positive that an answer is "X" or "Y" but rather view the solution as some mixture of the two. For this reason I propose that any defensive solution ought to include both a technical defense zone as well as a human defense zone. The layers of defense will then be designed and implemented to take advantage of the strengths of the two zones as well as attempting to mitigate the weaknesses of both.

What I propose is a defense structure containing six layers across two primary zones. This framework, called the "Phishing Layered Defense Framework," acts as guide to better consider and approach the task of defending the network from technical attacks based on social weaknesses or vulnerabilities. Figure 2 provides a visual representation



of this framework. It starts with a bottom foundational layer of technical tools designed to detect, prevent, and warn users of the pending danger, and then ends with

Figure 2

three categories intended to empower the user to make better and more informed decisions as they mull through legitimate and illegitimate requests. Notice again that no specific tools or approaches are listed within this framework; it is intended to act as a framework and not as a solution.

A category, however, is really only useful if one understands how to approach and analyze it appropriately. The framework in Figure 2 is intended to provide the defense in breadth as described above. It is expected that the necessary depth will be provided within each category. To do this I further propose a Categorical Defense Structure (CDS) to help the organization consider appropriate layers of depth for each categorical layer of breadth.

4.3 The McCumber Cube

In the early 1990s John McCumber postulated and designed a defense model that changed the way layered defenses were thought of. Prior to his research, it was primarily supposed that the intent of information security was less about the information and more about the system it resided on. The goal of administrators was to protect the database, the server, the network, etc. Protecting the system was the same as protecting the company's assets. Such systems and hardware do indeed represent a significant portion of the company's assets but the real value to any organization, and to the offenders attacking it, resides in the data/information contained within the systems and hardware. The data is what attackers seek to obtain, modify, or destroy, not necessarily the system. It is this concept around which McCumber built his model.

This model, commonly known as the “McCumber Cube,” is really a 3x3x3 matrix that considers the state of the information, the primary security objectives for each, and mitigation strategies (or layers of defense) to achieve each objective (Figure 3). The layers that make up the depth of the model are of the greatest significance to the Phishing Layered Defense Framework and are the backbone of the Categorical Defense Structure, for in these




Figure 3 *Image retrieved from (IBM 2001)

layers lie the mitigation strategies that are the framework or essence of defense in depth. This is because they consider defense from three primary factors: technology, people, and policies.

4.4 Categorical Defense Structure

The Categorical Defense Structure (CDS) is where this phishing framework gains its strength in depth (Figure 3). Whereas the main framework primarily considers various attack vectors, or breadth of attacks, the CDS addresses the depth of the mitigations considered for the various attack vectors. Once again, this framework is designed to be adapted to each environment and each organization, be it a secret and secure environment or a more open and public space. The strength of this approach comes from John McCumber’s cube theory that looks at each step in the phishing defense framework and asks the administrator how they feel they should best secure their environment from a policy perspective, from an education and training perspective, and from a technological perspective. The implication is that any vector is best protected by utilizing a multi-

pronged approach that encompasses the various tools at the fingertips of management and the organization.

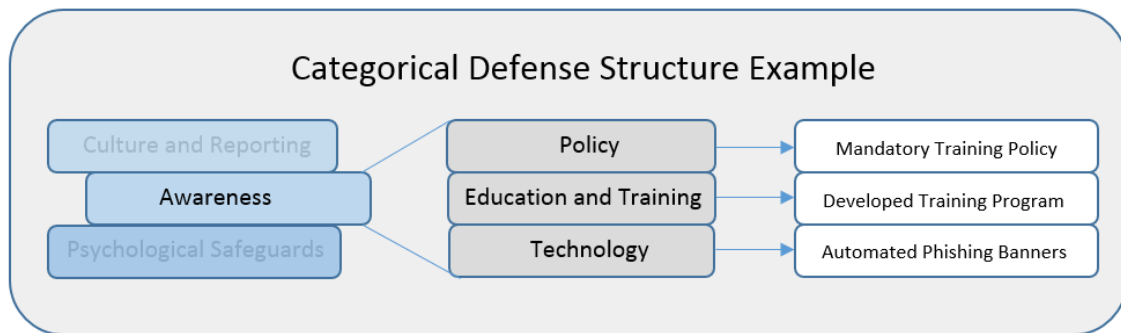


Figure 4

Administrators of such a framework should bear in mind that these three mitigations tactics may not be solved, individually, with a single tool or training. The best practice applied in any of the given strategies may require two or more assets to achieve a desired level of protection. For example, if I were an administrator seeking to establish secure login procedures into my network; I would likely need to establish multiple policies to do so. I might need a password policy that talks about proper length, strength, and usage of any password set in my system. I might also require a policy establishing what privileges users should have into the various resources offered by my organization. I might also require a policy establishing where and when it is acceptable for me to log into those resources. Other more pointed policies might be deemed necessary on top of those already defined for me to achieve my desired level of security. The same constraints may also require the use of multiple technologies to enforce such a mitigation strategy and

multiple avenues for training and education to raise the awareness to a sufficient level. In the end, the answer is not about finding the perfect tool to solve the problem, but rather, finding a great solution among a suite of tools and tactics that achieves the stated goal. It is about finding a depth of solutions that works best for that specific environment in that specific organization at that particular point in time, and adjusting those solutions as time moves forward. It is about solving the issue, not from a single standpoint or point of view, but rather considering all approaches that could be integrated into a more robust solution.

4.4.1 Policy

Policy is unique in that it is neither technical nor human, and yet it governs how both interact with each other and their environment. It acts as the fundamental building block for any security project that is undertaken. “Without one, a company cannot begin to know where or what to purchase or what processes to follow in order to secure their environment. The policy lays out, in detail, what the company’s goal is for security.” (Reyes 2014) Adopting a security policy has the added benefit of providing a framework for the rest of the security implementation. With the goals and framework established, direction is given for the creation, acquisition, and/or implementation of the technologies and training that are required to properly secure the given asset or assets.

A policy is the high level, 10,000 foot view, documentation of the goals and direction of a security program. From it, comes other documents designed to define in more detail the responsibilities and expectations of users and systems. From the security policy comes the creation of standards, guidelines, and procedures. Standards are detailed

documents explaining the mandatory controls that must be placed upon the defined resource. For example, a standard for password complexity would outline the length, history, use of numbers or special characters, and randomness for the creation of a user password. The keyword in the standard definition is the word “mandatory”. Standards must be achieved. Guidelines on the other hand are best practice suggestions. These documents outline what is recommended but are not made mandatory. Procedures are the children of the aforementioned standards. They are the step by step instruction sets for achieving compliance of the standard. Procedures allow technicians and users to perform necessary actions in a uniform and predictable manner ensuring uniform and predictable results. This is necessary to reduce the risk of a user misconfiguring or misusing an asset or resource.

As important as the creation of security policies are the enforcement of such policies. These policies should be monitored for compliance ensuring that what is outlined in the policy is what is being implemented in the organization. “Having a corporate security policy that is not monitored or enforced is tantamount to having laws but no police....the best deterrent to breaking the rules is not the severity of consequences but the likelihood of being caught.” (West 2008)

Policies should be the starting point for the risk mitigation process. As discussed earlier, they act as the foundational base for all other standards and procedures. They lend direction to the technologies that should be acquired and the trainings that should be established. As such, the associated security policy should be the starting point to any layered defense scheme. It is in the creation of this document where plans and strategies are created and discussed by management. In that process, limitations and boundaries are

set and processes are vetted. From the policies and the standards and procedures that follow, a more robust and planned strategy can be implemented and communicated across the organization, ensuring that all staff are in alignment to combat the given threats.

4.4.2 Education and Training

Addressing the need to include the human element of technology into a security program, education and training aims to combat threats of this nature by providing the training necessary to users who might come under attack. This training need not only be defensive in nature, but rather should include proactive measures that teach users to act in responsible ways prior to an attack, thus decreasing the attack surface exposed to a given threat. For example, a certain organization might seek to train their users on key indicators to look for when browsing the web. This might include helping users understand the difference between http and https. It might further include a discussion on certificates and certificate authorities (CAs) as well as common CAs like VeriSign, Comodo, or GeoTrust, to name a few. The danger would be to go into too great of detail and subsequently confuse or lose the audience in details and technical complexities. Such trainings should be simple to understand and direct in nature. Simply give the participants the knowledge they need to act effectively and accordingly and leave the detailed lists of how it all works to the engineers and analysts who need them. Filling presentations with this type of information can lead the user to feeling lost or feeling like the process is too complex to implement. If this happens they are likely to avoid implementing the knowledge altogether. (Lacovos Kirlappos 2014)

If designed and implemented correctly, however, organizations can expect a workforce that is not only aware of a range of threats and vulnerabilities, but also well equipped to properly handle them as the opportunity arises. This doesn't mean that they will necessarily deal with the threat themselves, but that they will know who to contact or what to do as the occasion arises. For example, a user who receives a mysterious email asking for credentials to the user's system will know not only to not respond to such an email, but also to contact the security staff and/or systems administrator for review of the email so that they can propagate a warning out to other users in the organization. Doing so strengthens the defenses of the organization as a whole. Other victims will know not to follow such a request, but also any unaware users who did fall for it will now be made aware and will know to contact their administrator immediately.

User training should also be an extension of the policies defined earlier. In such trainings, the associated policies that support the topic should be reviewed, making the users more aware of the firm's stance and expectations of them. Such trainings should support the policies and procedures defined, thus helping the end users to gain a greater understanding of their purpose and intent. The firm must also ensure that their reward system, or disciplinary procedures, align with the goals of policies and training. Too often an organization will encourage ideals like reporting, only to punish the individual who did so for falling victim to such a crime.

4.4.3 Technology

Technology. The poster child for all things IT/cyber related. The blinking lights and sophisticated toolsets that have been glamourized by authors and Hollywood alike are not

always the sensational products they are portrayed to be, and yet major advances in security tools are bringing us closer to a realization of that day. No longer is information security solely about firewalls and proxy servers. Today's collection of tools include dynamic training interfaces, machine learning detection algorithms, continuously monitored environments that present vulnerabilities to threats in near real time, big data techniques that pour through massive amounts of data finding useful trends for the analyst, and many more. Today's tools are faster, more robust, and produce better indicators across a much larger volume of information than has ever been accomplished in human history. With cyber-attacks continuing to rise in frequency and sophistication, we can expect that defense toolsets will continue to advance and become more efficient and accurate through time in an effort to keep pace with our adversaries.

While it has been noted that there is no magic bullet to prevent the potential for a cyber-related attack, the prevalent use of technology in a defense structure has become the primary method for mitigating such attacks. When planned and configured correctly, technological solutions offer the benefits of being scrupulous in their methods, reliable over time, and predictable in their results. To be clear, there is not a single tool in existence that has proven to be 100% successful in defending a system or network of systems over time. Even when multiple tools are deployed the success rate is high but not foolproof. What should be taken from this knowledge is that such tools, while highly successful, are still just tools and not answers. Notwithstanding its limitations, technology still provides an important and highly beneficial role in the defense of an organization's assets.

Like other aspects of information security, technological tools can, and should, be layered with depth and breadth in mind, thus providing adequate and best protection. However, we must carefully assess the value of the assets being protected in order to make an appropriate business decision on the amount of time and money invested into an acceptable risk factor. Most organizations are not in the business of security alone. Corporations and other businesses exist not to protect data and assets but rather to make money. Government entities provide services, infrastructure, and protection to its citizens, not just cyber security upon its environments. Very few cases feature information security as the primary goal of the firm. With that in mind, it would make very little sense for an organization to spend tens of thousands of dollars on defense technology if the asset itself only has an intrinsic value of \$5000. A financial analysis must accompany the security analysis when deciding on technology or any other mitigation tactic.

Similar to the need for a financial analysis accompanying the security analysis, so should the organization consider change management and the human side of implementing technological tools. Questions like “how will this impact my users?” and “will the culture of my firm, in its present state, accept such an implantation?” need to be asked and addressed before pursuing and implementing any type of technological toolset. I recently attended a forum where IT restructuring was discussed as part of a greater topic. During the course of that conversation a participant told of a recent experience she had while working at a national laboratory. Their management had decided to move from a legacy email system to the more modern Gmail platform. What seemed like a simple change was met with great resistance. Eventually the initiative failed and the lab reverted to its

former legacy system. The change failed because the company culture was well established in the ideal of an unchanging environment. Workers who had been there for 20+ years had grown accustomed to their workplace environment and toolsets, and had come to expect an unchanging atmosphere. Thus when even a simple change in their technologies was presented, they rejected it in favor of what they had come to know and expect.

The benefits provided by technology are enormous, and a firm would do well to carefully consider how to best utilize this asset in defense of their data and systems. Breadth and depth of toolsets and technologies should be considered in conjunction with business factors and complexities. When an organization does successfully plan, prepare, and implement a technological solution, in conjunction with policy and education, they can expect a more robust and complete security posture than any one of the individual approaches alone can provide.

Chapter 5: Technical Defense Zone

At the base of the Phishing Layered Defense Framework (PLDF) lies the Technical Defense Zone. It is this layer that provides the foundation of the defense strategy and front line defense to the users and data within the network. This zone represents the tool sets that are the first point of contact for the pending attack scheme. Before the end users receive the phishing email, and before they are lured into following a link to a phishing website, these tools and technologies seek to detect, warn, and mitigate the threat. These are the automated tools that once installed and configured properly, act autonomously in accordance with their configurations and rule sets. Though they may be united in their overarching goal of providing protection from a perceived threat, their approaches to so doing vary in design and implementation. What remains consistent, however, is that each tool's objective seems to fall into one of three different realms: detection, warning, and prevention.

Another constant, as researched by West (West 2008), shows that the best strategy for inserting and utilizing technological toolsets comes from using them in a manner that the end user is unaware of, or oblivious to, except where the need for warning necessitates human involvement. The two main reasons for this approach are risk and competing resources which are discussed later in this document. Suffice it to say for now that users prefer, and better results are achieved, when users do not perceive the protection that is being afforded them behind the scenes.

5.1 Detection Toolsets

Detection is the foundation for all other technical mitigation strategies. Without detection, we are hopeless in our efforts to prevent pending attacks or warn unaware users. It is commonly supposed that detection is the act of catching a probe or attack in progress as it attempts to enter the firm's network. However, I propose that detection can happen in one of three stages: pre, mid, and post attack. Thus there are three distinct opportunistic time frames by which an organization can identify the threat of an attack.

In the "pre-attack" stage there are very few indicators that a looming threat is lying in wait to attack the organization's valuable resources or assets. Occasional analytics reports may indicate that some form of reconnaissance has transpired, but those reports are hard to follow and are not always completely accurate. In this stage the security analyst must recognize that detection is not as much about detecting active attackers as it is about detecting the supposed threats that could exist, then match them against the known vulnerabilities that do exist in the firm's environment. Very similar to the way a disaster recovery policy is created and implemented, so could a phishing detection campaign begin. In disaster recovery planning a committee identifies the value of their assets, identifies potential threats to those same assets, and designs a mitigation or recovery strategy for each identified threat. Similarly, a firm could plan and prepare against phishing attacks by identifying their assets (systems, people, etc.), identifying potential threats (credential harvesting, Trojan horses, Nigerian 419 scams, etc.), and designing mitigation tactics to defend against them. (Jakobsson 2005) The idea is that

the company detects their resources and threats through reflection, intuition, and contemplation.

In the “mid-attack” phase of detection, technology can be used to great effect to discover and filter out illegitimate traffic from the legitimate. In this specific area much research is being done to more efficiently and effectively discover and report phishing traffic.

Two distinct approaches seem to be at the heart of much of the research. The first approach analyzes the links and/or the URLs that traffic is being directed to. One particular approach strips each link, or URL, out of the site or email and sends it through a battery of tests that range from a comparison of where the link says it is going and where it is actually sending the individual, to comparing it against a database of known bad URLs. (Ram Avtar 2011)

The second common approach in “mid-attack” detection looks at the actual web site the individual is sent to. Common features of sites are examined and compared to known phishing sites. With the advent of hacking tools that include automated phishing extensions, phishing sites are not only more frequently discovered but their designs tend to be very similar to one another. Another approach taken by Blum treats the different parts of the site as individual aspects to be analyzed. Attackers often use mimicry to trick their victims into a false sense of security. For example, an attacker might scrape, or copy, Google’s Gmail login page and present it to the victim as a legitimate login page. However, when the victim inserts their credentials, they are harvested by the attacker and the victim is redirected to the actual Gmail login page. Often the victim assumes that they mistyped their password and is none the wiser to the attack. In Blum’s approach each portion of the webpage is pulled out and a hash value is calculated and compared

against the legitimate webpage. A threshold is set and compared against the average matched rate to distinguish the false site from the real one. (Aaron Blum 2010) The idea is that even small changes will register in the hashed values and could be used to detect phishing attempts.

The final “post-attack” phase of detection is done primarily through forensic analysis. When a phishing attack is reported a forensic investigation can reveal how the attack happened, why it happened, and what could be done to prevent a similar attack in the future. A regular analysis of systems logs and traffic logs might also reveal a previously unknown successful phishing attack. This type of post-mortem review can seem pointless to some who view it as a moot point since the attack is in the past, but to the keen analyst it can provide great detail about the firm’s threats and vulnerabilities that will lead them to closing those doors to future attacks.

5.2 Preventative Toolsets

Preventative tools are the toolsets used by firms to eliminate or filter out phishing threats before they reach their intended target, or that render the functional aspects of the attack useless. These tools range in platform from which they are deployed as well as functional use within the organization. Common examples of these toolsets include custom written filters that use targeted algorithms to find and block phishing attempts, content filters and proxy servers, firewalls and antivirus products, phishing and spam filters inside web browsers, configuration settings including the disabling of scripting languages on systems, and many others. While the tools and approaches may vary, the goal of these toolsets are the same, to block or make nonfunctional any and all detected

phishing attempts. These toolsets provide varying degrees of success when it comes to preventing phishing attacks and further research could be done to evaluate their individual effectiveness as well as a combined success rate when different combinations are utilized.

What research has been done however, indicates that these tools provide better protection when their use is imperceptible to the end user. In an article written by Jason Hong, he seems to concur that one of the first steps to improving protection against phishing is to stop what you can from ever reaching the end user. Make it invisible. His approach to this principle is to create defenses that are imperceptible to the end user. Using filters on emails and enabling phishing filters in web browsers are two examples of ways that administrators can make protection invisible to the end user (Hong 2012)

5.3 Warning Toolsets

Where the previous categories have suggested that their technical toolsets should remain unperceivable to the end user, the category of ‘warning toolsets’ takes a hard left turn to that approach. In fact the warning aspect is all about the end user. It is about letting the user make a more informed decision as they traverse digital space. These toolsets include warning banners, generated alerts, training toolsets, and others that are designed to grab the attention of the end user and make them aware of a detected situation. Such warnings as “Possible phishing attempt detected” or “This is an untrusted connection” are common inside web browsers and can even be seen from antivirus tools and intrusion detection engines.

Such warning tools walk a very fine line of properly warning users of a potential threat and becoming a reporting nuisance that train users to ignore such signals. With today's operating systems it is commonplace to see pop-up windows that ask the user to confirm their actions or inform them of a process taking place. It is advised that security alerts should not resemble these normative alerts but instead stand out and be very clear of the presence of danger. (West 2008) The reasons for this from a human perspective will be explored later, but from a technical perspective such alerts are different in nature from a confirmation or informative window. These alerts speak to the safety of not only the system from which it is generated but to the various systems that could be potentially exposed to the same threat. Thus they should capture the attention of the user and persuade them to truly look over the data being presented. They should be bold in color, maybe even a red as red is associated with danger, and concise and clear in their message. Again this would be a good place for further research into what features users are most likely to pay attention to when important information should be delivered.

Such tools should also be properly configured and tuned to the environment in which they reside. In access control we refer to the term 'crossover error rate' (CER) when speaking of the optimal point at which the likelihood for a false positive meets the likelihood for a false negative. Move to the right or to the left of that optimal point and it is expected that either the false positives or the false negatives will increase above the other. Similar thought or concern should be given to the tuning of any device that alerts users. If too many false positives are created the users are likely to begin distrusting the system. Too many false negatives could result in a resiliency effect where the user begins tuning out the warnings, treating them as white noise.

Simply inserting a warning system is not sufficient enough when seeking to better inform users. Careful thought and consideration should be given to how it is used, when information is presented to the user, and what type of information is given. The end goal is not necessarily to persuade the user from taking an action but to make them more informed over what that action might entail. The action could be legitimate and it may be something that only a human in that situation would be able to ascertain.

Chapter 6: Human Defense Zone

The “Human Defense Zone” represents the people of the organization. From the CEO to the mailroom girl, each person in the firm acts as a resource and as a vulnerability to the organization and its assets. Each person comes with a unique perspective on risk and risk avoidance. Because life experiences have been different for each individual, their responses to a perceived threat can be vastly different and can also be situational.

Humans use rational and intellect to make decisions which makes their responses very different from that of a computer. Where a computer system will use rule sets or probabilities, humans will use past experience and logic that can sometimes be clouded with hopes and aspirations. Where computers generally look at the problem in isolation, humans will involve the environment, the timing, how the outcome will affect them personally, how the outcome will affect others, long and short term benefits, and so on.

In short, humans do what computers cannot, and conversely, computers do what humans cannot. It then stands to reason that computers are vulnerable to attack where humans are not and humans are vulnerable to attack where computers are not. Because of this, the mitigation strategies, or layers of defense, are very different for this part of the defense structure than that of the technical zone. Here we focus on psychology, awareness, and culture. We look at preventing phishing through human interaction with people and technology, and try to understand the psyche of an individual in order to understand what makes them vulnerable and how to best prevent them from becoming a victim.

While technology is used in an effort to protect or shield humans from phishing attempts, it is unlikely to ever become 100% successful. Phishing, by design, leverages human vulnerabilities in order to attack the systems and data inside the organization. Every day

many of these attacks make their way through the technologies designed to stop them and end up in the user's inbox, on the other end of the phone line, on the computer screen posing as a legitimate website, or on the victim's desk in the form of a letter. Once the attack has made it this far, the defense strategy can no longer be about technology. The safety of the data and systems of the firm is now in the hands of the user. It is up to them to make the right decision and it is up to the organization to give them the tools and perspective necessary to fight against such an attack.

6.1 Psychological Safeguards

The psychology of a phishing attack can be as much of a resource to stopping it as any log file or security analysis. Understanding both sides of the attack gives the defender a better perspective as to what makes his/her organization vulnerable, and just how far the attacker might go to exploit that vulnerability. For example, understanding that the motivating factor of one attacker might solely be glory driven, indicates that their resolve to work through layers of defense could be weak as they will likely give up quickly in order to find an easier target that validates their "skills" as a hacker. However, the same attacker might find greater resolve if the target is perceived to be of high value.

Compromising a higher value target would yield greater prestige and greater validation.

On the other side of the attack, users inside the organization typically vary in their acceptance of risk. Some users inherently refrain from risk and are thus weary of anything that could lead to a negative outcome. Others are more willing to accept risk especially if the reward is perceived to be good enough. To the later, the temptation of financial gain could be enough to lure them into falling for a typical Nigerian 419 scam.

To the former, a well-crafted email that indicates that the user's bank account could be at risk and they need to change their credentials right now, could be enough to get that user to follow a link to a fake sight where their credentials are harvested. Considering the psychological nature of the firm's employees is to consider what normal behavior is. An organization cannot rely on an idealized view of how people might react to a phishing scam. They must, instead, understand what is typical of human behavior and create a defense strategy around that. On this topic one researcher said it well when he stated "IT Security Professionals and their leadership cannot afford to think in purely strategic and tactical terms without considering what normal behavior is." (Small 2011) Surely a firm must first consider what it is before it can consider what it can become.

While it is important to consider the psychological nature of both the users inside the firm and the attacker, the organization usually has more influence on the users and can better create defenses around them. Thus the balance of this section is dedicated to understanding the psychology of the user and what can be done to mitigate their vulnerabilities.

The first thing to remember is that for most users, security is not their primary responsibility. The company accountant, who with only days left to get the company's quarterly reports filed, would likely place a precedence of financial reporting over a phishing threat warning banner. An operations manager is more likely to study the newest applications of Total Quality Management than the organization's resource manual on avoiding phishing attacks. Management and security professionals should realize that users have plenty on their plates already and don't want to be bothered with security. They want to focus on the work they were hired to do and not meddle with

topics that seem to be outside the scope of that work. To this end, a security professional might focus on making as many of the security tools as possible function in the background in a way that is unperceivable to the end user. As described earlier, when a user must be alerted it should be done in a way that is unmistakable as being important. When a user's attention is required they must know that it is important to do so and not mistake the message for just another user message.

Most users do not understand how attackers operate or the tactics they use.

Consequently, those users rely on life experiences to make decisions on how to proceed. (Lacovos Kirlappos 2014) The online world does not utilize the same boundary systems that we might expect from the real world. Social norms and acceptable behaviors can be different from their real world counterparts. Additionally, the World Wide Web is a global environment taking in many cultures and nations. Where the typical user may perceive only the local customs of that region in their office environment, their online environment is filled with a melting pot of the norms of other cultures. In some of these situations real life experiences do not translate well to the online environment and do not provide adequate understanding to enable the user to make the best decision. Without flooding them with knowledge they won't use and do not care about, organizations need to provide context and understanding of the threats that exist and some of the tactics used against targeted individuals. Awareness, covered in greater depth later, can be an outstanding mitigation tool at the psychological level as attackers often view the human mind as their greatest asset. (Small 2011)

The psychology of risk and rewards is one that has been extensively researched, and in context of IT security, I tend to agree with one researcher who believes that all users, to

some degree, want to take risks. While being risk adverse is the nature of human beings risk abstinence is not realistic. Since users are going to take risks and must assess the level of risk they are taking on a daily basis, firms should do a better job of teaching them to assess the risk vs reward behavior rather than just telling them to avoid risk altogether. (Lacovos Kirlappos 2014) They should be empowered to make decisions but given the tools necessary to make more informed decisions. Technology and training could be included into the firm's defense structure to ensure that users are armed with information and techniques needed to make better and more informed decisions. This is a realistic business approach that frees workers to be efficient and encourages productivity while helping them to act as defenders for the network and its resources.

Other psychological factors include but are not limited to: "Users do not think they are at risk," "Users aren't stupid, they are unmotivated," "Security concerns are too abstract, users are motivated by concrete factors that they can see and touch," and "Losses are perceived disproportionately to gains" (people are risk adverse). (West 2008)

Addressing these, and other conceived factors, with policy, education, and training can provide strength to an organizations greatest vulnerabilities, its people. Further research on the key psychological factors that make humans vulnerable to this type of threat could be done in an effort to strengthen this aspect of the framework. Much research has been done on psychological factors but no one has been able to narrow it down to the most crucial and key factors. This could be a valuable insight that would allow this framework, and others, to better target effective mitigation strategies.

6.2 Awareness

The United States Coast Guard defines awareness (specifically situational awareness) as: “the ability to identify, process, and comprehend the critical elements of information about what is happening to the team with regards to the mission. More simply, it’s knowing what is going on around you.” (Guard 2014) With that in mind we can extrapolate the same principles to phishing and defending against it. A good defense from the human element should begin with some element of situational awareness. System users do not necessarily need to be security professionals to combat phishing but they should be cognizant of the threats and strategies used by attackers to compromise their systems. They should be able to recognize and identify the key indicators of a phishing attack. They should also be able to identify and understand the key indicators of safe web browsing. Users should also understand why/how the indicators of trustworthiness they use in real life can fail them online. (Lacovos Kirlappos 2014)

Lacovos also proposes that awareness is delivered in two basic forms: 1) general public awareness and/or education campaigns and 2) context-specific warnings while online. (Lacovos Kirlappos 2014) General public awareness and education campaigns are what you would typically think of when you think of a firm’s training program or even a public service announcement. For instance, I have recently noticed a TV public service announcement that includes directors from the Internal Revenue Service (IRS) and the Department of Homeland Security (DHS). In this particular campaign they are warning against responding to emails portrayed to be from the federal government that ask for personal information. Their response is that the US Federal Government will never ask for this type of information in an email. Other forms of education campaigns could

include a training program put on by an organization before allowing users to utilize their systems, regular trainings as outlined by company policies, and even media displayed around the office offering reminders of security principles.

The US Federal Government on the civilian side abides by the framework defined by the Federal Information Assurance Management Act (FISMA). Included in this framework are many publications written by the National Institute of Standards and Technology (NIST). The NIST SP800-12 document is an entire section that defines and lays out the goals of a security training program as it should be implemented in the US government. What makes this publication so interesting is that it defines three specific goals for such a program as seen in figure 4. These goals include “Awareness,” “Training,” and “Education.” The document further explains that each goal is used to define a specific attribute. In the case of awareness, the goal is to define the “what” of security and is mostly informational. Training defines the “how” and adds knowledge to awareness, and education defines “why” providing insight and understanding to the training. This is important to understand because not every user needs the insights provided under the “education” umbrella. Depending on their specific role in the firm, a user might only need to focus on awareness in certain topics and not need the deep dive into training and education. Remember, most users do not have the primary goal of security in their job description and often times less is more for them. Psychology tells us that users are often unmotivated and can be overwhelmed with tasks and therefore will reject some information to make room for that which is important to them. For these users, awareness training with a simple and direct method may be sufficient. For IT

administrators and security personnel, perhaps training and education is required. In the end, the training should be tailored to the people receiving it.

Comparative Framework			
	AWARENESS	TRAINING	EDUCATION
Attribute:	"What"	"How"	"Why"
Level:	Information	Knowledge	Insight
Objective:	Recognition	Skill	Understanding
Teaching Method:	<u>Media</u> - Video - Newsletters - Posters, etc.	<u>Practical Instruction</u> - Lecture - Case study workshop - Hands-on practice	<u>Theoretical Instruction</u> - Discussion seminar - Background reading
Test Measure:	True/False Multiple Choice (identify learning)	Problem Solving (apply learning)	Eassay (interpret learning)
Impact Timeframe:	Short-term	Intermediate	Long-term

Figure 5

**Image retrieved from NIST SP800-12 ((NIST) 1995)*

Context specific warnings provide another form of awareness to the end user in an organization. These types of warnings are most often seen as banners or pop up messages in web browsers or from antivirus tools. These messages act as cautionary warnings alerting the user to a detected violation of safe browsing habits or potential risky behavior. The main problem with these types of communication is that users have become resilient to their warnings. Banners and pop ups are commonly seen as a nuisance rather than a tool. Users, in a hurry to accomplish their task, often bypass such warnings with reckless abandonment, eager to arrive at their webpage or submit the requested information and move on to their next task. Yet, "It is the lack of cognitive recognition and risk deferral that makes an individual vulnerable and susceptible to the social engineering and other attack vectors that are so successful." (Small 2011)

To make context specific warnings more effective the research indicates that they should be limited to warnings that are of the utmost importance. Additionally, they should

clearly stand out from the typical pop up message. As discussed earlier, such warnings should be clear and concise and should grab the attention of the end user. If such users feel like the warnings and banners are just another nuisance, they will likely ignore it, but if they truly stand out and look and feel important, they stand a better chance of being read. It is then up to the training received by the individual to teach them to interpret what they are seeing and how to handle such a situation.

Active and passive training techniques for phishing defense are among the most widely studied subjects I found in my research. Researchers like Kumaraguru and the folks at Carnegie Mellon University have done extensive work to find different approaches to training and educating users of the detection and dangers of phishing. In one such study they created an active training system called “Phishguru” that deliberately sends users phishing emails. If the user fails to identify the attack and instead follows the instructions in the email, they are redirected to a training page in their web browser where they are trained again on the dangers and indicators that they missed. With this approach they were able to demonstrate that users trained using the “Phishguru” system were less likely than those who were not trained to fall for a phishing scheme up to 28 days later.

(Ponnurangam Kumaraguru 2009)

Whether it is active or passive; awareness, training, or education; a program to disseminate knowledge and awareness to all users within an organization is a key factor in combating phishing attacks and all of its variants. Awareness and training programs should be well defined in company policies. Education and awareness programs should be well thought out and designed to target a specific audience whether that audience is management, security professionals, accounting staff, or other. Technology can aid in

this effort through the use of well-designed and properly placed warning banners or through sophisticated training tools like “Phishguru.” Awareness, like the other categories in the defense framework should, at its most basic level, include a defense in depth approach.

6.3 Culture and Reporting

One of the areas rarely discussed as a mitigation tool is that of a firm’s culture. Culture within an organization represents the norms and unenforced behaviors that exist among the people that make up the organization. Culture within the firm can affect the way people feel toward the organization and toward each other. It may influence desires and aspirations and can even determine the likely actions of an individual. The company culture may determine how one responds to an event or plans their normal routine. It may influence the level of loyalty and dedication one feels toward their organization and specifically toward their job.

Perhaps one of the reasons that culture is rarely discussed as a mitigation tool is because it can be a two edged sword. When an organization enjoys a strong company culture, the variables that are affected by culture act in a positive manner, but when a firm experiences a weak culture those same variables can be unpredictable and sometimes detrimental. For example, an assembly line worker in a company with a strong culture of quality and excellence may see a defect in a product and decide to remove that product and report the defect to the engineering department for analysis. Conversely, that same worker in a company with a weak culture that reflects a lackadaisical attitude may allow

the same product to continue to move down the line and out the door to market, giving little thought or care to the end customer whom would receive it.

Culture can also affect, for better or worse, the way cyber threats are handled within the organization. For example, rather than an assembly line worker let us suppose that the employee is now a marketing supervisor who receives a somewhat suspicious email from an old account representative. She notices that the email did not come from the company's email platform and that it is asking her to follow a link to a third party site in order to update information for a project that was completed last month. Upon following the link to the third party site she is presented with a verification page asking her to use her company credentials to verify her identity. She recognizes this as a phishing attempt and, in a weak company culture, thinks nothing more of it than to disregard the email.

However, in a strong culture of awareness and security she might alert the systems administrator or security staff and forward the example on to them as well. This would enable the security staff to review the threat and send out warnings to the rest of the organization, thus better defending the organization as a whole against this cyber threat.

“Security is benefited greatly when organizations, communities, or groups share “sanitized” security data. The attackers are very good at this and are experiencing positive results from their collaboration. We can expect the same types of results when security is shared among security communities.” (Small 2011)

Creating a strong organizational culture can seem like a daunting task at the onset, but much in the way of research and study has been done to help one work through the process. I will not go into detail here on how to create a strong culture within the organization but will focus on two key factors that appeared frequently in my research.

The two key factors that I kept coming across were 1) a rewards system and 2) the idea of psychological safety. In many change management frameworks the idea of rewards or small-wins are frequently utilized. Researcher John Kotter identified small-wins as one of his eight key steps to successfully implementing any change. (Kotter 2014)

Researcher Kurt Lewin described a step of refreezing, or validating the change process. (Burnes 2004) The validation process itself acts as a form of reward system. Ryan West expressed a need for rewards speculating that “increasing the immediate and tangible reward for secure actions may increase compliance.” (West 2008) While many researchers agree that a reward system of some type is recommended, at least one warned of the dangers of a misguided one.

In business strategy there is concept that ‘intended’ strategy often differs from ‘realized’ strategy. In other words, what we say we value or are going to do does not always align with our actions or what we actually do. Often, what actions or events a firm rewards is not in accordance with what they claim to value. An organization that claims to value employee safety might reward a manager for reducing company costs by reducing the number of safety officers in that organization. Another firm that promotes honesty as one of their values may reprimand an employee who refuses to endorse a project that hides funds in offshore accounts in an effort to evade taxes. These two examples may seem obvious but what about American politics. In the United States, we claim that we want honesty and well documented, well thought out, plans for leadership. However, when a candidate takes this path during an election they are often picked apart and battled against by the media and advocacy groups. Often times, it is the politician who works in generalities that is able to appease both sides enough to win the contest. (Kerr 1975) In

security, if we are to gain the trust of our users then we must ensure that our rewards coincide with our message. If we value reporting then we cannot reward users for minimal incidents reported throughout the year. We also must be careful how and when we punish those who self-report incidents they became involved in, which leads to the second point: psychological safety.

Psychological safety can be thought of as “a shared belief held by members of a team that the team is safe for interpersonal risk taking.” (Edmondson 1999) Members of the team feel comfortable in expressing their beliefs and making decisions that present a moderate degree of risk without fear of reprisal from other team members or management. They feel confident in speaking open and freely, expressing their thoughts and concerns to coworkers, team members, and/or management. In Edmondson’s research she confirmed that teams who enjoy a significant level of psychological safety are generally more efficient and effective at their jobs. (Edmondson 1999) The firm as a whole is benefited and the mission is better achieved.

The same applies to an organization that hopes to include reporting as part of their mitigation strategies against cyber threats, specifically phishing. Users should feel encouraged by the organizational culture to report such attacks to those who can evaluate the attack and propagate warnings out to the rest of the firm. They need to feel a level of comfort and safety from reprisal in so doing. Great benefits can be gained when a company and its resources, human or otherwise, are able to work together and share information in an effort to further their success.

Chapter 7: Further Research and Study

The PLDF presents a novel and unique approach to the threats of phishing to organizations of differing size and structure. This particular approach is sound as it is based off of proven strategies and techniques though it in itself is unproven. Further research might be done in the future to validate the PLDF framework and test its completeness and assumptions. Particular attention might be paid to the psychology of risk as it relates to online behavior. Further development and understanding of a proper rewards system that reinforces acceptable online behavior and a strong culture of security would also be of great value to the PLDF model. Additionally, as technologies change and methods of implementation improve providing greater insight toward phishing defense, it is hoped and expected that further study could be done to determine how such improvements might impact the PLDF.

Chapter 8: Conclusion

We are well into the age of digital wealth where the value of an organization's digital assets typically outweigh the value of its tangible assets. This new age grants companies, its employees, its customers, and even its enemies greater flexibility when accessing the firm's resources, particularly those assets that utilize a digital medium. The benefits and convenience of having greater access, and an increase in the number of avenues by which access is obtained, is vast and impressive. Yet, along with such conveniences arises an equally vast and impressive range of threats and attacks. Those who seek to gain unauthorized and/or illegal access to such resources, and whom are sufficiently motivated, will utilize any and all methods to achieve their goals. This includes more than an attack on the technologies that hold such data alone. An ever increasing number of attacks are focusing on the human element as attackers focus on social engineering and phishing style attacks. If left unchecked, it is the human, and not the technological vulnerabilities, that may be the greatest of all the threats to an organizations data and resources.

The PLDF is a model that seeks to address the phishing threat within an organization. Its concept is derived from research and consideration of defense in depth as well as defense in breadth principles. The crux of the framework lies in its duality of focus between the technological elements and the human elements of a phishing attack. The breadth of the framework is derived from the categories that are encompassed in these two main zones while the depth of the framework's defenses are a product of the Categorical Defense Structure (CDS). This structure defines the mitigation tactics that a firm might employ to defend their resources in a layered manner. Stemming from research on the

McCumber cube, the CDS integrates defenses from technology, education, and policy into a cohesive and targeted defense strategy.

When utilized as a framework, and properly employed, it is expected that the PLDF will provide a more complete and robust defense structure against the ever increasing threat of phishing against the organization. The PLDF is a toolset designed to give the defenders of a given firm's resources the upper hand in an ongoing battle involving the confidentiality, integrity, and availability of the firm's assets. It is further, a framework that should be re-assessed at regular intervals by the organization, and it is expected that as such organizations do so they will enjoy constant protection from phishing events over time.

Technologies change and their implementation in the organizational structure will continue to change, but what remains constant is that humans are an integral part of the organization. They bring with them the benefits of rationalization and judgment and the vulnerabilities of the same traits. It is, and will be, crucial for management of such organizations to provide adequate and appropriate defenses to protect the data and other resources entrusted to them. The PLDF framework is one that enables management to accomplish such a goal.

References

Special Pub 800-12 -- An Introduction to
Computer Security: The NIST Handbook.

AISeC '10 Proceedings of the 3rd ACM workshop on Artificial intelligence and security

"M O h # k Journal of
Managment Studies 41:6

) OM K O \ bbn.com.7 ydU° @--

Science Quarterly Vol 44

7 o h t y o t o h o
o Communications of the ACM vol 54

8 y o # o .o uscg.mil.U

= arxiv.org/abs/2007.04831

= U o o O # K) h h M
‡ 7 h °) h o
- @ CHI 2010 [electronic resource]; we are CHI; the 28th
Annual CHI Conference on Human Factors in Computing Systems; conference
proceedings & extende

= K u . . . Communications of the ACM K V . . .

=

U U K# k U

U ‡ h @)) 'lockheedmartin.com.'K

U ‡ h @))

@U *ibm.com*.
 U
 K U U h h *Financial Cryptography (Vol 5)*
 7 # t

Zhang, Wei. "How Could I Fall for That? Exploring Phishing Victimization with the Heuristic-Systematic Model." *System Science (HICSS), 2012 45th Hawaii International Conference on* , 2012: 2374 - 2380 .