

The Influences of Social Networks on Phishing Vulnerability

Kathryn Coronges Ronald Dodge Cort Mukina Zachary Radwick Joseph Shevchik Ericka Rovira
United States Military Academy
first.last@usma.edu

Abstract

Phishing is a form of electronic deception in which an attacker tries to cause the recipient to do something or disclose data that they likely would not normally do by mimicking a trustworthy entity. These attacks have been increasing at an alarming rate and can cause damages in the form of identity theft, financial losses, and compromised security for organizations and governmental institutions. Additionally, phishing attacks have become very sophisticated and even more successful because of the lack of vigilance by computer users. Successful phishes have particularly strong implications for military populations, and have the potential to threaten national security. In an attempt to reduce the overall success rate of a phishing attack, this paper applies the foundations of social network analysis to identify how social network structures among a military company of future US Army officers are most influential in reducing the spread of a phish. This experimental study collected empirical and survey data in an effort to analyze the flow of information and influence of people in phishing awareness within an organization.

1. Introduction

Over the course of history, technology has vastly improved the ability of the human user. Specifically in the 1990s, internet and computer technology advanced at an extremely rapid rate, allowing humans virtually endless capabilities from the comfort of their own home. In 2000, the US Census reported that nearly 60% of American households owned a personal computer and utilized it on a daily basis (Monitor, 2002). This increased technological capacity allowed Americans the ability to communicate with someone through e-mail and perform tasks like paying the bills and shopping. Unfortunately, all of the increased benefits associated with computer technology came at the cost of personal security [3].

As users continue to utilize computers and access the internet daily, they become more susceptible to a fraudulent scheme known as “phishing.” Phishing is a form of electronic deception in which an attacker tries to cause the recipient to do something or disclose data

that they likely would not normally do by mimicking a trustworthy entity [4]. Phishers try to lure victims to falsified websites, usually through spoofed emails, by “employing both social engineering incentives and technical subterfuge to steal consumers’ personal identity data and financial account credentials” [6].

Names, social security numbers, financial account passwords, credit card numbers, and bank account information are what phishers desire most, all of which they try to capture for their own personal gain [6]. Billions of dollars are lost each year from financial institutions and millions are left with compromised identities and destroyed credit [3]. Also, falling victim to a phishing attack can also be seen as a security risk where attackers have a higher probability of accessing secure networks like governmental agencies [6]. Stealing identities and classified information through Military officers have risk to compromise security efforts.

A recent study conducted by the Anti-Phishing Working Group in 2009 reported a high of 56,362 unique phishing attempts occurring in a single month. The numbers of phishing attacks have and are expected to increase and become even more sophisticated as technology continues to advance. Because of the increased sophistication as well as the number of attacks, it has become vitally important to teach and train employees and computer users alike in an effort to reduce the success rates of phishing attempts, particularly populations whose compromised identities could threaten national security [7].

Although there has been an increased effort to educate and train computer users against phishing attacks in both military and civilian sectors, little research has been conducted on the effectiveness of training modules since they are still in their infancy. Phishing training typically includes increasing awareness about how others can access private computer files through websites, and other internet interfaces. One aspect of phishing awareness that has been given little attention so far involves increasing security resilience by taking advantage of social interconnectedness. Perhaps training that includes a component on network communication would encourage individuals to warn others when a threat is perceived. In particular, we consider whether using

one's social connections to warn others about a phishing attack would reduce the number of phishing victims, and whether the reduction would be associated with structural aspects of the network. For example, companies, which often have local area networks connecting the employees with one another, would benefit by warning their superiors, subordinates, their personal contacts of the phishing attempt.

This current study attempts to determine which warning source, friend or superior, is more likely to report the phish, and which source is more influential in reducing incidence of a successful phish. In order to explore these research aims, an experiment was developed where phishing emails were created and sent to a controlled user population. The participants in this study were provided one of two types of training techniques to evaluate their effectiveness. In addition, victims and 'warners' of the attacks were studied to see whether these behaviors could be characterized by social structural metrics.

In this paper we leverage existing efforts at the United States Military Academy to attempt to use social network analysis to gain additional insight into the impact on security awareness in the context of phishing emails. The remaining sections of the paper are organized as follows: section two presents a brief literature review of other and supporting work in the security field, section three describes the previous work in support of this study, section four describes the study process for this effort, section five discusses the results, section six is a discussion and seven concludes.

2. Related Work

In the development of the US Army's wide scale security efforts, it became imperative to first understand which populations are most susceptible to phishing. In a study conducted by Carnegie Mellon University, titled, Who Falls for Phish? A Demographic Analysis of Phishing Susceptibility and Effectiveness of Interventions, 1001 online surveys were administered to college students to study the relationship between demographics and phishing susceptibility. These surveys asked participants questions to determine their background and assess their knowledge about phishing. Participants completed a scenario in an effort to assess their behavioral susceptibility to phishing [8].

The results of this study suggested that women are more likely than men to fall victim to a phishing attack. In addition, a younger population (18-25) was suggested to be the most susceptible to online deception. This specific age's vulnerability was a direct result of their increased likelihood to engage in

risky behaviors when compared to an older population [8]. This study provided evidence that a younger user population (18-25 years old) should be tested during the execution of the experiment.

To evaluate the effectiveness of phishing attempt warnings, basic research about obedience principles can be applied. It is no surprise that people oftentimes listen to the advice given by friends and superiors, even if there are different motivations for doing so. One study, involving students' ranking of musical performances, demonstrates the influence of authoritative presence. The music was ranked significantly higher when it was labeled as professional (coming from an authoritative source). Thus, participants were more likely to change their behavior according to their perceived professional rank of the information source [9]. Additionally, numerous studies have demonstrated the link between a person's friends and their own behavior [10]. For example, individuals often accept their peers' attitudes and beliefs as reality, and these perceptions often direct their consumer decisions. In other words, people value the beliefs of their peers and will often emulate them [11].

Previous research suggests that a person will perceive a threat in a situation where a warning is provided by either a professional superior or a personal contact. Although plenty of research in these separate aspects of obedience has been conducted, little has been done to evaluate the two areas (friends vs. authority figures) with respect to one another. It is desirable to determine which will have a greater influence in the realm of phishing vigilance.

3. Experiment Background

The authors have conducted many studies in the effectiveness of training and education in stemming a person's susceptibility to phishing. The most recent effort provides the data for the social network analysis reported on in this paper. While only a subset of the population was used for the social network analysis, the entire data collection effort is described below for completeness.

The phishing emails sent to all subjects included an embedded URL that when clicked takes users to a web site where they are asked to enter sensitive information (their network credentials). In all cases the email 'bait' leveraged knowledge of the cadet population (spear phishing), where some sort of free or discounted service appealing to the target population was used.

The target population for all groups was 892 Academy college students. The population was broken down into three notification conditions. Notification condition was randomized by organizational unit (i.e.,

cadet company consisting of about 140 students). There were three notification conditions:

Group 1 (No Notification): received the phishing email, however after the user entered data into the website and clicked submit, the page returned a server error and no additional information was provided to the user.

Group 2 (Notification): received the phishing email, after the user entered data into the website and clicked submit, the page returned a notice that they fell victim to a phishing attack and provided details as to what the user should have identified in the email.

Group 3 (Training): received the phishing email, after the user entered data into the website and clicked submit, the page returned a notice that they fell victim to a phishing attack and directed the user to take the institutions phishing awareness training.

Two cadet companies were assigned to each notification condition. There were 287 members in the group one, 298 in group two, and 307 in group three. For the two companies within each notification group, the phishing email was fabricated in one of two ways. Cadet companies either received the phish from the West Point IT department, or received the phish from a cadet in a leadership position.

Emails were sent from a third party service provider outside the institution's boundary. The service selected was phishme.com [12]. The emails themselves were constructed to provide several clues designed to alert end users. By default, emails are displayed in plain text mode; users have the ability to choose to view in HTML after the email has been displayed in plain text. Figures 1 and 2 show both presentations of the email.

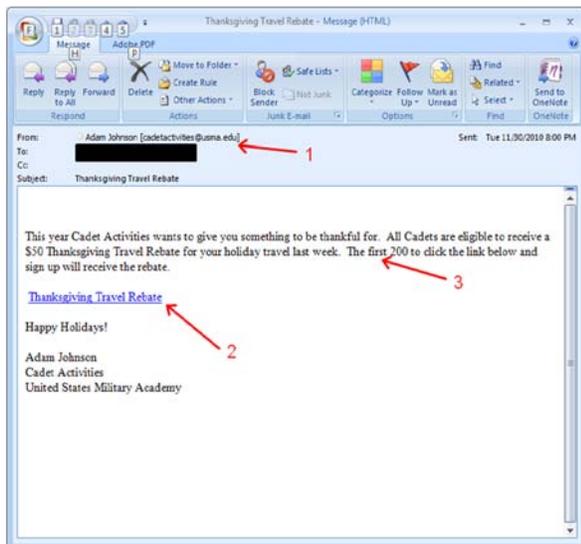


Figure 1: Sample email: html view

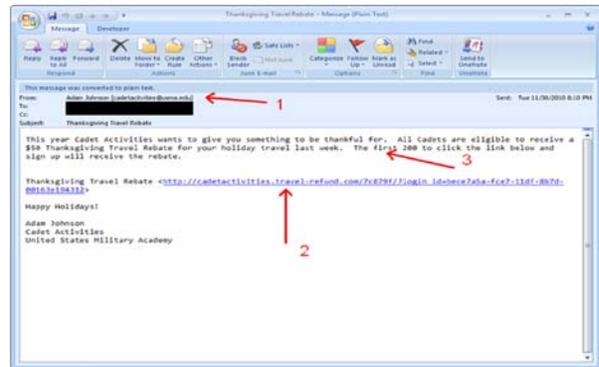


Figure 2: Sample email: plain text view

The following items details are items that are part of the annual training our users receive and should have alerted them that the email was likely not legitimate.

1. Source email address (Adam Johnson [cadetactivities@usma.edu]): The email is from a person; however the source address is an institutional address. Further the user does not exist in our global list (accessible by all authorized users).

2. The URL in the email: In figure 1, the URL is shown in the middle using html formatting for email. This is not how email is delivered by default to our users. Instead the 'presentation text' is listed as well as the actual URL string, which is how it is first displayed to users (as shown in Figure 2).

3. The email urgency: While not always a valid sign of phishing email, when an urgent email is received from a source outside the organization, it is highly suspect.

4. Finally, if they looking in the full mail headers (which is an option if they are concerned about the validity of the email); they would have seen that the originating email server is highly suspect. This is shown below using bolded and underline text.

```
... Received: from localhost.localdomain
(localhost.localdomain [127.0.0.1]) by
mail.phishme.managedmachine.com (Postfix)
with ESMTP id 5DC6B10A43 for
<REDACTED@usma.edu>; Wed, 1 Dec 2010
01:05:14 +0000 (UTC)
Date: Wed, 1 Dec 2010 01:05:11 +0000
From: AdamJohnson <cadetactivities@usma.edu>
To: <REDACTED @usma.edu >
Subject: Thanksgiving Travel Rebate
MIME-Version: 1.0
Content-Type: text/html; charset="utf-8"
X-Priority: 3
X-Pmsid:775892d4-fce0-11df-97b7-0163e4638cc
```

Message-ID:

<20101201010514.5DC6B10A43@mail.phishme.magedmachine.com>

The social network analysis was conducted on a subset of the larger population. A single organizational unit (cadet company) was asked to complete a follow-up electronic survey to obtain information on their personal social network structure. The organizational unit consisted of 128 users, and has both a well defined formal leadership structure (subordinate and superior) and an informal social network (friends). The single organization analyzed in this study participated in the larger study in the exact same manner as the other users. This particular organizational unit was assigned to group 2 condition, where they were simply informed that they were phished and notified when they fell victim to the phish. The phish appeared to come from their company commander, the highest ranking cadet in the organizational unit.

The age of the participants ranged from 18-26 years. All participants within this experiment were treated in accordance with the American Psychological Association rules and regulations.

4. Experiment Process

This network study was designed to explore which networks (formal or informal) and which structural positions within those networks have the most influence on stopping a phishing attack. All 128 participants received the same phishing email at the same time. The content of the email warned each participant that they were expected of suspicious internet activity.

The phish included a message saying that their internet activity needed to be verified. They are instructed to open a link, which when opened asks them to enter their network username and password. Participants could ignore the link, select the link (but not enter information), or select the link and enter their network username and password. The individual was identified as a victim to the phishing attack if they clicked the link embedded in the body of the email or went further to enter their personal identification numbers. Data was recorded on www.phishme.com as to which participants clicked the link and/or entered their username and password.

After the conclusion of the phishing attack, a survey was sent to the entire company. This was an electronic survey that asked participants about their demographics, social networks and about their response to the recent phishing attack. Demographics included gender, ethnicity, age, and class. Phishing items included questions about whether the participant

fell for the phish, and whether they warned anyone of the attempt. Social network items asked participants to list their friends in the organization. Immediate supervisory chain was constructed using a publicly available command roster. Also, a series of technological questions were asked to determine how many hours a day a participant uses their personal computer, what kinds of activities are performed on their computer (gaming, homework, multi-media), and their choice of internet browser.

5. Results

For this study, phishing victims are defined as users who clicked on the embedded link and those who entered their network credentials. Of the 128 users in the social network study group, 48% (62 participants) clicked on the embedded link within the email, 30% (39) of those users also entered their username and password to their network accounts, and 21% (27) successfully avoided the attach. Of 128 students phished, only 7 users in this organization warned others about the potentially malicious email.

Friendship and command network data were entered into network analytic software, Organizational Analyzer (ORA) in order to derive several indicators and visualize the network [12]. ORA software uses relational data input from the study to visualize the network, and compute network-level indicators. Network level variables include link count, density, distance, and various centralization measures.

Network density is the ratio of the number of links to the total number of possible links. Degree centrality refers to the extent to which one or a few individuals in the network have a larger number of links than the others. In-degree includes incoming links whereas out-degree accounts for outgoing links. Betweenness centrality reflects the number of times a node lies between the shortest path that exists between all other pairs of nodes. Closeness centrality is the average distance required for each node to reach all other nodes in the network. See Wasserman and Faust, 1994 for mathematical descriptions of network measures [14].

Figures 3 through 6 illustrate the structure of the two networks where circles represent individuals and links between them represent either friendship or a supervisory relationship. Table 1 compares network-level measures for the two organizational networks. Links are directed, which in the case of the friendship network means that the arrow points from the person who named the other as a friend, and for the formal network, the arrow points from the supervisor to the subordinate. Figures 3 and 5 show results of the phishing attack. Those who fell victim to the phishing

attack are represented in red, while those who resisted the attack are shown in green. Figures 4 and 6 show which individuals warned others about the attack. Blue nodes warned others of the attack while black colored nodes did not.

Supervisory Network. Of the two of highest ranking individuals, one fell for the phish. Of the 126 other members, 13% of the second highest ranking positions, 27% of the third rank, and 31% of the lowest ranking individuals failed the security test. With the exception of the highest level of leadership, position of power corresponds to about a 10% increase in security awareness. Surprisingly, those who warned the organization of the attack were more likely to be the lowest ranking members within the company. Therefore, leadership position correlates with reduced phishing vulnerability but also to less agency to help protect the rest of the company against the phish.

Network relationships allow us to understand behaviors as a consequence of direct social ties. We can see that security failures tend to cluster together. For example, of the four 3rd ranking individuals who were phished, 35% of their subordinates were also phished as compared to only 29% with supervisors who resisted the phish. As can be seen in Figure 4, warning behavior had no effect on security failures. In fact, in the case of the 3rd ranking individual who notified his subordinates, 50% of them fell for the phish – the highest failure rate in the entire company. Further, the highest ranking individual (whose name was listed as the sender of the phishing email) warned the company with apparently no increase in security awareness.

Friendship Networks. In Figure 5, friend network phishing patterns are illustrated. Individuals were 1.8 times more likely to fall for the phish if their friend was phished. There was a moderate tendency for those high in betweenness centrality to be victims. Otherwise, those who failed the security test did not have particular structural characteristics.

Figure 6 shows which individuals warned others about the phish. There were no network structural patterns found among informers – those who warned others were no more likely to be central or peripheral. Like warnings from formal leaders, warnings that came from friends (regardless of their overall popularity) had no bearing on who fell for the phish.

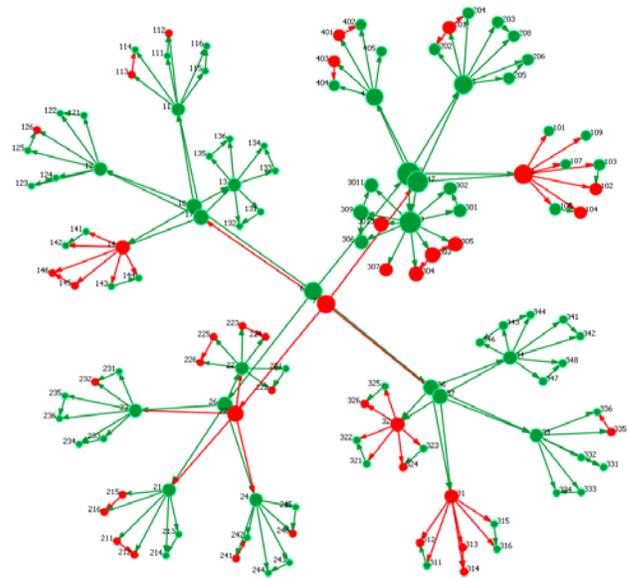


Figure 3: Supervisory Network: Phish Victims Sized by Centrality

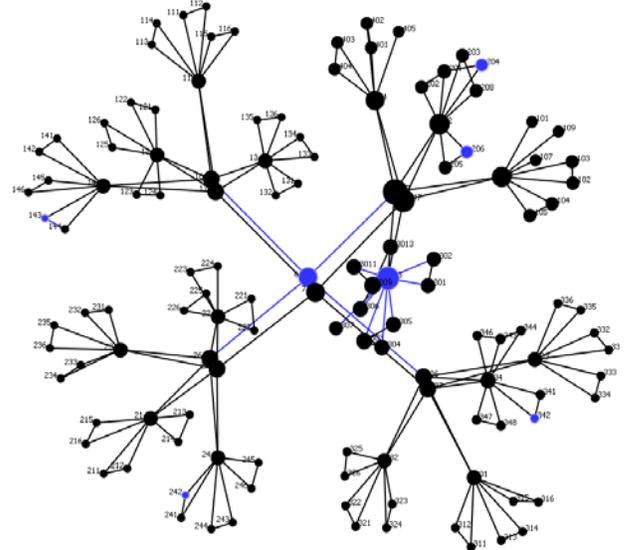


Figure 4: Supervisory, Warned Others sized by Centrality

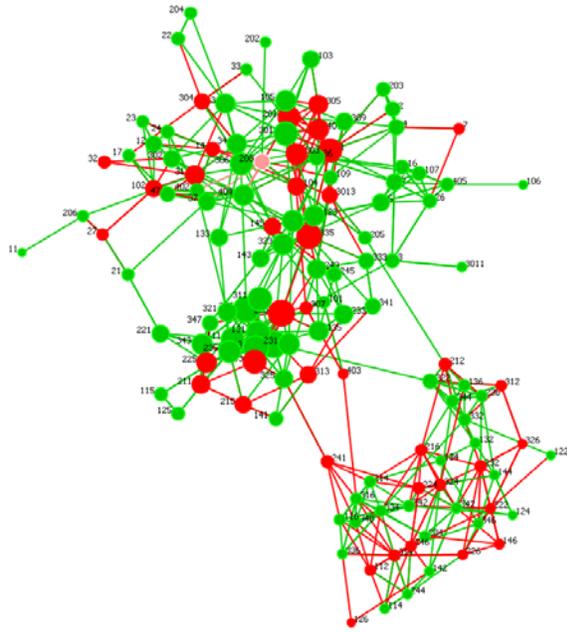


Figure 5: Friend Network: Phish Victims Sized by Centrality

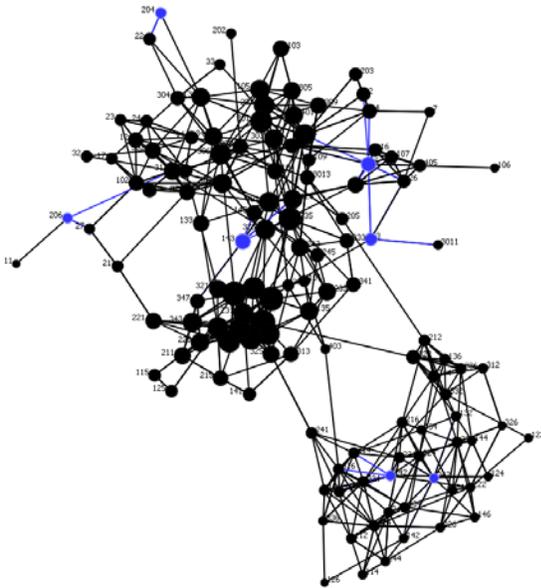


Figure 6: Friend Network: Warned Others Sized by Centrality

Table 1 compares the network metrics for the two types of networks. The friend network has significantly more links (link count) and higher density compared to the supervisory network. In addition, friendship structure has high betweenness and high in-degree

centralization compared to the supervisory network. High betweenness centralization indicates that there are key individuals who act as brokers between unconnected sections of the network.

As compared to the supervisory network, diffusion of information through the friendship network is more clustered and dependent on a few key individuals. Thus, information can spread rapidly around popular individuals but those in peripheral positions are relatively isolated from the rest of the social network.

Differences in closeness centralization show that unlike the friendship network, information can spread from any part of the formal network to any other section of this network extremely effectively. An individual only has to go through an average of two people to reach the entire unit through the supervisory network, while they would have to go through an average of five classmates in the friend network. Further, if we consider outdegree centralization, information can be transferred from a supervisor to their subordinates (“out-going links”) 37% faster in the supervisor network as compared the friend network.

Table 1: Network-Level Measures

	Friends	Supervisor Chain	Difference
Link Count	600	198	-67%
Density	0.036	0.012	-66.48%
Average Distance	5.020	2.009	-59.97%
Network Centralization, Betweenness	0.259	0.002	-99.15%
Network Centralization, Closeness	0.042	0.708	+1573.24%
Network Centralization, In Degree	0.066	0.004	-94.53%
Network Centralization, Out Degree	0.050	0.067	+34.20%
Network Centralization, Total Degree	0.058	0.036	-38.92%

6. Discussion

Network perspectives enable researchers to consider the role of structure in the dissemination or disruption of information through a community or organization. Network indicators show that friendship or informal networks are clustered, and are highly centralized, where a few individuals have key roles in spreading information. In this study, highly central individuals did not make any attempt to warn others about the phishing attack and therefore, network capabilities were not mobilized through informal connections. Supervisory networks can be much more efficient than informal networks where all individuals can be reached with few number of steps (2 steps, on average). Therefore, any attempts to warn other members of the attack should have traveled through the supervisory network relatively quickly. However, the few warnings that occurred did not spread successfully – which speaks not to network structure but perhaps the perceived trust or value around phishing information.

Results reveal several important aspects about the role of social factors in security awareness. First, leaders phishing failures correspond to their subordinates' phishing vulnerability. Second, low ranking organization members appear to play a unique role – where they are both the most susceptible but also relatively more active in warning others of the phishing attempt. Third, there was a surprisingly low level of information sharing (both in warning attempts, and warning compliance) through both formal and informal social connections.

Local Leadership. The pattern of victims in the both the supervisory and friendship networks suggests that the local level leadership has the biggest impact on susceptibility to phishing attacks. Those with supervisors or friends who were phishing victims were more likely themselves to also be victims. One interesting observation is that 2 influential people (the senior supervisor and the deputy) in the supervisory network attempted to warn people of the phishing attack, yet these warnings had little or no effect on people's behaviors. Further, direct subordinates of informants were no more likely to avoid the attack than others in the organization. These results suggest that local level leaders may have an effect on phishing vulnerabilities but apparently no effect on security resilience. In other words, superior's weaknesses were more likely to correspond to their subordinates' behaviors than were their strengths.

New User Vulnerability. The majority of the victims came from the newest users (14 victims). This group was also relatively active in warning others about the phishing, containing 4 of the 7 informants. The sphere of influence of a new user on other groups within the supervisory network is extremely limited,

and it is even more limited in the friend network. This group is the youngest, has the least experience in the organization and holds the lowest ranking leadership positions.

Interestingly, this group is the most segregated within the friendship network, indicating perhaps that they rely on their professional superiors to be role models and provide them with trusted information. Maintaining close personal relationship between new users and senior users is not common which creates a highly segregated portion on the network. The large separate cluster in the bottom right of figure 5 represents the new user class in the friend network. It is not surprising then that they fell victim to the phishing which was ostensibly sent by their superior. What is surprising though is that the upper leadership were not more active in warning others about the phishing, suggesting that security is not a major concern for this population.

Missing Informants. This organization had only 7 people warn others about the phishing attack. The few warnings that were conveyed from friends and leaders had no effect on limiting spread of phishing victims. The user breakdown for the 7 informants is 3 senior users and 4 new users. Of the senior level users, two held the most highest level of the leadership. One of them used a verbal alert and the other emailed the entire organization. Despite this, there was apparently no increase in awareness.

There are many factors that may have played a role in why so few individuals warned others about the attack, and why the few warning attempts made were unsuccessful. Determining these factors should be considered for future research because uncovering these factors could lead to better IA training and reduce susceptibility to phishing attacks from a network perspective.

7. Conclusions

This study attempted to explore the effects of social networks on the dissemination of information about a well known, yet highly successful social engineering attack – phishing emails. Local leadership appeared to influence security vulnerability, but had no influence on security resilience. Superior's phishing failures were more likely to correspond to their subordinates' behaviors than were their phishing successes. This study is one part of a large effort to understand the many influencing factors in why phishing is effective. As a form of social engineering, it is not a surprise that social relationships play a large role in effectively mitigating the impact of phishing.

In future work, the authors will further explore the efficacy of various training programs and the frequency of their delivery.

8. References

- [1] Downs, J., Holbrook, M., & Lorrie, C. (2006). Decision Strategies and Susceptibility to Phishing. Symposium on Usable Privacy and Security.
- [2] Heim, S. G. The Resonant Interface. Boston, MA: Pearson Education-Addison Wesley, 2008.
- [3] Hicks, D. (2005). Phishing and Pharming: Helping Consumers Avoid Internet Fraud. Communities and Banking , 29-31.
- [4] Jagatic, T. N., Johnson, N. A., Jakobsson, M., & Menczer, F. (2007). Social Phishing. Communications of the ACM , 94-100.
- [5] Monitor, C. S. (2002). Poverty now comes with a color TV. Retrieved November 29, 2010, from MSN: <http://webcache.googleusercontent.com/search?q=cach e:RQYryc3YGV MJ:articles.moneycentral.msn.com/Investing/Extra/PovertyNowComesWithAColorTV.aspx>
- [6] Phifer, L. (2010, April 12). Top Ten Phishing Facts. Retrieved November 29, 2010, from eSecurity Planet: <http://www.esecurityplanet.com/views/article.php/3875866/Top-Ten-Phishing-Facts.htm>
- [7] (2009). Phishing Activity Trends Report. Anti-Phishing Working Group .
- [8] Sheng, S., Holbrook, M., & Kumaraguru, P. (2010). Who falls for a phish? A demographic analysis of phishing susceptibility and effectiveness of interventions. 28th International Conference on Human Factors in Computing Systems . Atlanta: CHI.
- [9] Radocy, R. E. (1976). Effects of Authority Figure Biases on changing judgments of musical events. Journal of Research in Music Education. Vol. 24, No. 3. MENC, 119-128.
- [10] Valente, T. W. (1995). Network modeling of diffusion of innovations. Cresskill, NJ: Hampton Press.
- [11] Mangleburg T. F., Doney, P. M., Bristol T. (2004). Shopping with friends and teens' susceptibility to peer influence. Journal of Retailing. Vol. 80, No. 2, 101-116.
- [12] www.pishme.com, accessed 15 February 2011
- [13] Carley, Kathleen M. and Jeff Reminga, (2004). ORA: Organization Risk Analyzer. Carnegie Mellon University, School of Computer Science, Institute for Software Research International, Technical Report CMU-ISRI-04-101.
- [14] Wasserman, S., & Faust, K. (1994). Social network Analysis: Methods and applications, New York: Cambridge University Press.