Analysis of Unintentional Insider Threats Deriving from Social Engineering Exploits

Frank L. Greitzer
PsyberAnalytix
Richland, WA USA
Frank@PsyberAnalytix.com

Jeremy R. Strozer, Sholom Cohen, Andrew P. Moore, David Mundie and Jennifer Cowley
Software Engineering Institute, Carnegie Mellon University
Pittsburgh, PA USA
jrstrozer@sei.cmu.edu; sgc@sei.cmu.edu; apm@cert.org; dmundie@cert.org; jcowley@cert.org

***Abstract*—Organizations often suffer harm from individuals who bear no malice against them but whose actions unintentionally expose the organizations to risk—the unintentional insider threat (UIT). In this paper we examine UIT cases that derive from social engineering exploits. We report on our efforts to collect and analyze data from UIT social engineering incidents to identify possible behavioral and technical patterns and to inform future research and development of UIT mitigation strategies.**

Keywords— unintentional insider threat; social engineering

I. INTRODUCTION

The insider threat is recognized as a major security risk by computer and organizational security professionals, more than 40% of whom report that their greatest security concern is when employees accidentally jeopardize security through data leaks or similar errors [1]. The accidental or unintentional insider threat (UIT) problem has only recently been studied more formally as part of the insider threat problem; recent research [2] has provided an operational definition of UIT, a literature review on possible causes and contributing factors, and a tabulation of frequencies of UIT occurrences across several categories. This initial work served to inform government and industry stakeholders about the problem and its potential causes and to guide research and development (R&D) investments toward the highest priority R&D requirements for countering UIT.

The present paper reports on further UIT research that seeks to advance our understanding of contributing factors by focusing on UIT incidents involving social engineering. In particular, we review our efforts in collecting and analyzing social engineering UIT incident data to identify possible behavioral and technical patterns and precursors. The paper is organized as follows: Section 2 provides working definitions of UIT and social engineering, the latter being defined within the context of UIT. Section 3 provides an update on relevant research literature, focusing on social engineering UIT incidents. Section 4 describes the case collection requirements we developed to guide collection and reporting of UIT cases and provides examples of representative UIT cases involving social engineering exploits. Section 5 provides initial conceptual models, discussing results synthesized from our research and case study analyses to identify patterns that may be useful in designing mitigation strategies. Sections 6 and 7 provide conclusions and recommendations, respectively.

II. DEFINITIONS

Building on the original work in [2], we use a slightly updated definition of UIT:

> *An unintentional insider threat is (1) a current or former employee, contractor, or business partner (2) who has or had authorized access to an organization's network, system, or data and who, (3) through action or inaction without malicious intent, (4) unwittingly causes harm or substantially increases the probability of future serious harm to the confidentiality, integrity, or availability of the organization's resources or assets, including information, information systems, or financial systems.*

Malicious intent includes the intention to cause harm. Harm can also be caused by those who have no malicious intent (i.e., are non-malicious), either by action or inaction, even if they knowingly break a rule (i.e., the guard who fails to check all badges does not mean to allow a malicious actor into the building, but he lets someone in who sets the building on fire.).

The updated definition includes several minor changes. One change emphasizes that the unintentional insider's actions occur largely without the insider's knowledge or understanding of their impact; we added the term "unwittingly" to the fourth part of the definition. A second change is to broaden the description of the target of the attack to include assets such as personnel and financial systems. Thus an organization's assets include people, organizational information including protected personal information and intellectual property, financial data and information systems.

A UIT incident typically results from actions (or lack of action) by a nonmalicious insider (although not all such cases are characterized as completely nonmalicious, and

DOI 10.1109/SPW.2014.39

individuals involved may not always be identified). The unintentional insider's actions are often in response to an attacker's social engineering activities. We adopted the following working definition of social engineering and related exploits, in the context of UIT incidents:

> *Social engineering, in the context of information security, is manipulation of people to get them to unwittingly perform actions that cause harm (or increase the probability of causing future harm) to the confidentiality, integrity, or availability of the organization's resources or assets, including information, information systems, or financial systems.*

Social engineering represents a type of confidence scheme aimed at gathering information, committing fraud, or gaining computer system access. Social engineering, almost by definition, capitalizes on human psychology, such as cognitive limitations and biases, which attackers exploit to deceive the victim. This differs from other types of UIT incidents, such as cases in which an individual inadvertently discloses sensitive information without any interaction with an outside party (e.g., posting information on public databases or losing information by discarding it without destroying it). The adversary (or adversaries) masterminding the social engineering UIT incidents may have one or more malicious objectives that correspond to the intended impact to the organization, such as financial loss, disruption, or information compromise.

This type of exploit does not typically constitute a single attack, but rather a step that occurs within a more complex sequence of actions that compose a larger fraud scheme. We have found it useful to identify two levels of social engineering incident:

- **Single-stage attack—**As the name implies, the exploit is carried out in a single social engineering incident. The attacker obtains information as a result of the exploit and uses this information to cause further harm to the insider's organization. The attacker does not use the information to conduct further social engineering exploits.
- **Multiple-stage attack—**The attacker capitalizes on information gained from an initial exploit to execute one or more additional social engineering exploits. Some multiple-stage exploits play out over a matter of minutes or hours, while others may last for weeks or longer as the attacker applies the compromised information to cause harm.

III. Social Engineering Taxonomy

Several researchers have tried varied approaches to categorizing types of social engineering attacks. For example, Peltier breaks down social engineering into two main categories, human based and technology based [3]. Another decomposition uses the categories of close access (essentially human-to-human), online, and intelligence gathering [4]. Some combination of each of these perspectives applies: social engineering often occurs in multiple stages, so that a UIT social engineering incident may fall into multiple social engineering taxonomic categories. We adopted a simple yet comprehensive categorization that is consistent with descriptions of social engineering exploits in the scientific literature as well as real cases reported in court documents and other print media (see Figure 1). This provides a mutually exclusive, exhaustive organization of the various forms of social engineering exploits. Our research focuses on the portion of the taxonomy that applies to UIT incidents.

At the highest level of the taxonomy, we distinguish between whether or not exploits use interpersonal interaction. While social engineering is typically thought of as an interaction between people, UIT exploits commonly begin with the attacker gathering intelligence on the individual or organization being targeted for an attack. One type of intelligence gathering is referred to as *dumpster diving* or trashing [4], in which an attacker searches for sensitive information in the garbage (e.g., bank statements, pre-approved credit cards and student loan documents that are carelessly thrown away). A second type of intelligence gathering is *open source research* [4] that includes searching websites (e.g., Facebook, company websites) for information on targets that may be exploited in a second phase of a social engineering attack.

Social engineering attacks that include interpersonal interaction involve direct communication (such as in person or by telephone) or interaction that is mediated through electronic means (e.g., electronic media, email, and Internet). These attacks are characterized by exploitation of human psychology to deceive the victims and achieve some objective (financial, sabotage, etc.).

Non-electronic social engineering exploits are designed to gain physical access to computer systems or the information they contain. Social engineers use people skills such as friendliness, impersonation, conformity, decoying, and sympathy to exploit trust relationships and gain desired information [4]. One form of non-electronic social engineering is *shoulder surfing*, or stealthily looking over the shoulder of someone who enters security codes or passwords. Another broad method is *impersonation*, or creating a character and playing out a role to deceive others. Whether by telephone or in person, an attacker who uses impersonation typically pretends to be someone in a position of authority and attempts to persuade the victim to provide sensitive information. *Reverse social engineering* is a sophisticated form of non-electronic social engineering, in which the attacker creates a situation in which the unwitting victim believes that the attacker can help solve a problem. Typically the attacker poses as a technical aide to fix a problem that the attacker created or that does not exist. The attacker communicates his capability to help, such as through advertising or a phone call. Finally the victim invites the attacker to assist, which eventually allows the attacker to access to the desired information.

The methods of most concern in this study are those in the Electronic Means branch of the taxonomy. The literature describes many of these types of exploits. To simplify the discussion without loss of generality we describe the

following representative electronic social engineering exploits (as shown in Figure 1; see also Table I, which summarizes the salient characteristics of social engineering attacks, typical information sought, and possible consequences of the incident):

- **Baiting/Trojan horse**—an exploit that uses malware-infected physical media (e.g., CD-ROM, USB drive) to perpetuate an attack. Looking legitimate, the Trojan horse relies on the curiosity or greed of the victim who finds and uses the device, enabling installation of the malware on the targeted organization's internal computer network.
- **Fraudulent websites and social media**—an exploit that uses a fraudulent website (or social media site such as Facebook) to trick the victim into clicking on a link that downloads malware to the victim's computer.
- **Pretexting/reverse social engineering**—an exploit that creates and uses a real or an invented scenario (the pretext) to increase the chance that a targeted victim will divulge information or perform actions that would be unlikely in ordinary circumstances. A sophisticated example of pretexting is reverse social engineering, which was described above in the context of non-electronic social engineering scams. When applied to electronic (online) interactions, reverse social engineering has proven to be a very effective computer-based exploit.
- **Phishing/spear phishing**—an exploit generally defined as a phisher impersonating a trusted third party to gain access to private data. Typically, the phisher sends an email that appears to come from a legitimate business or individual (e.g., a bank, credit card company, or fellow employee) requesting verification of information and warning of dire consequence if it is not provided. The e-mail usually contains a link to a fraudulent web page that appears legitimate—sometimes with company logos and content—and requests private information (e.g., Social Security number, bank account number, banking PIN). Social engineering, and particularly phishing, has become more sophisticated over time: attackers learn which techniques are most effective and alter their strategies accordingly [5][6]. An example is spear phishing, in which the attacker initially gathers personal information about the target victim and uses it to tailor the phishing scheme, which increases the probability of success [7].

As seen in Table I, the information sought and potential outcomes are, not surprisingly, much the same as the targeted information and consequences in cyber-attacks

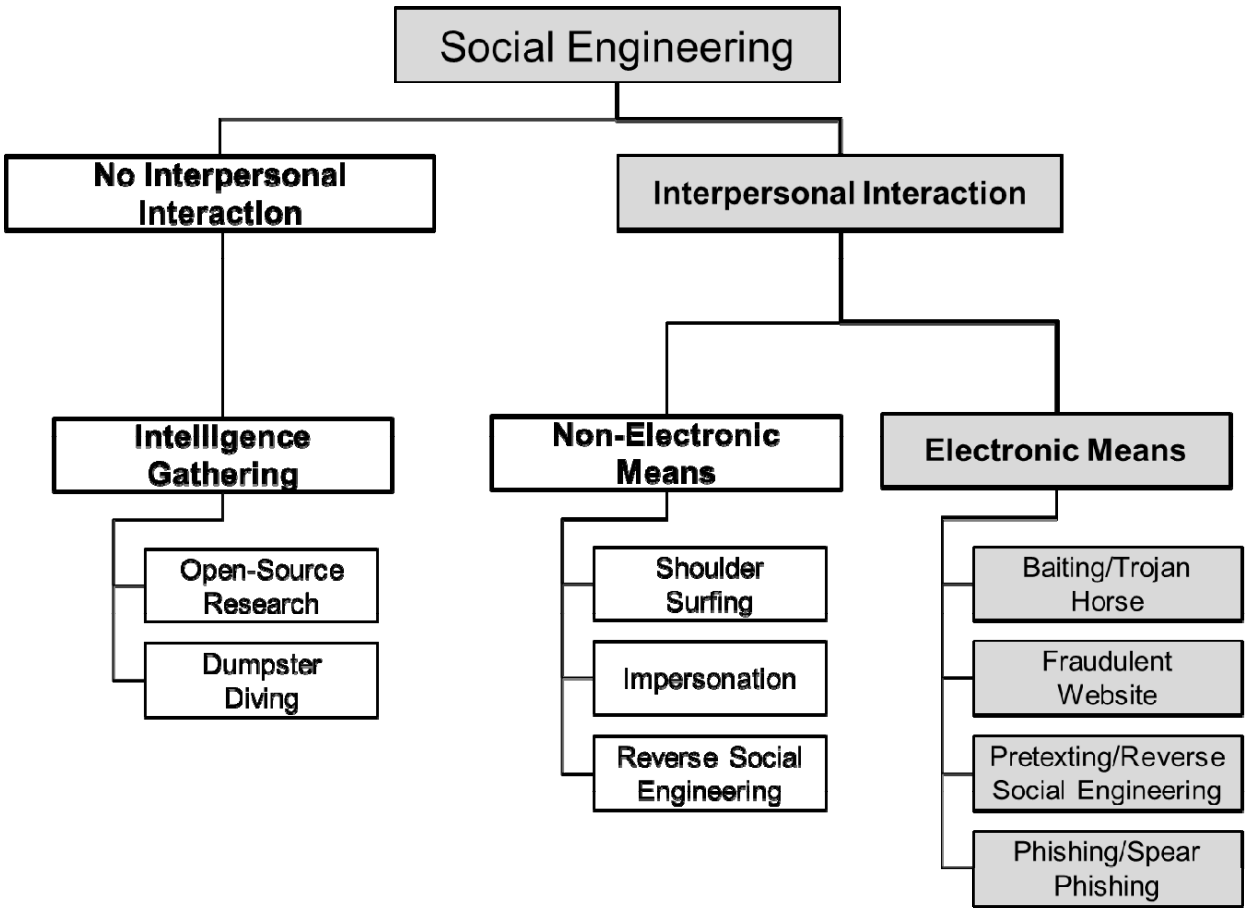

Figure 2. Social Engineering Taxonomy (highlighted branch of interest to this study)

generally, although the methods of attack differ somewhat, especially regarding salient characteristics in the first column of the table. These characteristics inform our approach to describing social engineering incidents and identifying patterns in these attacks.

TABLE 1: SOCIAL ENGINEERING CHARACTERISTICS

Salient Characteristics	Typical Information Requested	Potential Consequences/ Outcome
Appeal • usually good news or bad news • sense of urgency • sensitive or confidential matter • impersonating known sender **Desired response** • provide specific information • update personal/account information • click on link in email message • open an attachment **Suspicious indicators** • generic greetings • suspicious context • poor grammar or spelling • strange or unusual sender • incorrect information • illegitimate embedded URLs	• account information • user name • password and PIN • credit card number • Social Security number • bank account number • bank routing number • email address • telephone number • other personal information	• financial loss • identity theft • personal, confidential, or proprietary information stolen • intellectual property stolen • computer compromised, malware or virus implanted • data, software, and/or hardware assets manipulated or destroyed • personal or organizational embarrassment • political gain • denial of service

IV. RELATED RESEARCH

We organized possible contributing factors into three broad categories: demographic, organizational, and human factors. Below we briefly describe these factors and summarize findings from our literature review (see [2] for a more complete discussion).

A. Demographic Factors

Demographic factors are within-person factors that characterize who people are and their past experiences. Such factors include age, gender, personality traits, and culture.

Studies of age differences in phishing susceptibility have been inconsistent, with some findings inconclusive or indicating no age effects and some showing increased susceptibility for college-age participants. For example, a role-playing survey study conducted by Carnegie Mellon University researchers [11] examined the relationship between demographics and phishing susceptibility for 1,001 online survey respondents. They found that participants between the ages of 18 and 25 were more susceptible to phishing than other age groups (26–35, 36–45, 46–55, and older than 56) [11]. On the other hand, a phishing experiment found no evidence for age-related patterns in phishing susceptibility for students at different undergraduate levels (freshman, sophomore, junior, senior) [8]; note that the age range in this study is more restrictive than that in [11]. Another study in a university setting found no significant differences in phishing susceptibility between students, faculty and staff [6]. However, a success rate of 72% was found in a phishing experiment with 487 Indiana University students ranging from 18-24 years old [10], with a slightly higher susceptibility in students at the younger end of this scale. (This success rate is comparatively high—we may speculate that the higher rate was obtained due to the more sophisticated phishing attack that was employed). The reported age effect occurred within an age range that is not comparable to other studies: the restricted age range in this study falls entirely within the youngest age categories defined in [11].) In summary, these studies appear inconclusive, and it is possible that age differences were confounded with experience in some of these studies.

Studies of gender differences in phishing susceptibility have also been inconsistent. In the role-playing survey by Carnegie Mellon University, [11] females were more susceptible than males to phishing. Similar gender differences were reported by Halevi and colleagues [12], who suggested that women may feel more comfortable with digital communication and may be more inclined to reply to emails that advertise commercial offers or prizes. In contrast, a large-scale phishing experiment conducted with more than 10,000 human subjects in a university setting found no significant gender-related patterns in phishing susceptibility [8]. In an initial phishing attack involving spoofed email that navigates the user to a website to change a student's password, males and females were equally deceived; in a second phase of the attack that used a survey to harvest personal information, nearly 61% of the victims were male compared to only 39% females. Differences in these findings may be due to variations in methodologies used (surveys vs. experiments); it is also possible that differences may have resulted from varying levels of control for confounding variables like experience, course of study or job position.

Personality traits are stable, inherent aspects or characteristics of a person's personality (e.g., neuroticism, agreeableness, extroversion, openness, conscientiousness [13]). Differences in personality may influence the manner in which people interact with others, how they approach decisions, how they respond to job uncertainties or job pressures, or how they react to social engineering exploits. Empirical studies have reported that neuroticism was more highly correlated to responding to a phishing email scheme [12] and that openness was associated with higher social engineering susceptibility [12][14]. Examining more detailed facets of the broader personality factors, Workman [15] reported that people who are higher in normative commitment, more trusting, and more obedient to authority are more likely to succumb to social engineering. While there is limited research on personality traits, these studies

suggest that it may be possible to use personality profiles to identify individuals who are at a higher risk of falling for social engineering scams.

Cultural factors include characteristics of the individual's attitudes and ways of experiencing life that the individual adopted. We were unable to find any formal, comparative studies of possible cultural differences in susceptibility to social engineering exploits across defined cultural variables. At best, one can only draw tentative conclusions in comparing the few experiments conducted in non-Western cultures with those reported from Western countries. For example, a study conducted in Saudi Arabia with 200 students reported a 7% response rate to phishing email [14], comparable to a variety of studies in Western cultures reporting 3-11% response rates (e.g., [9], [16], [8]). These findings suggest that there is little, if any, difference in phishing susceptibility, at least between the Western and Middle Eastern populations studied.

Despite the lack of strong relationships between social engineering susceptibility and various demographic factors, it will be useful to record demographic information as case studies are collected and tabulated to help resolve inconsistencies, methodological uncertainties, or apparent contradictions in findings among published studies.

B. Organizational Factors

Organizational factors refer to management practices, policies, work environment, workload, and related aspects of the workplace that may contribute to performance deficiencies and human error, which in turn underlie certain types of UIT incidents. Direct mention of such organizational factors in published research within the cybersecurity and insider threat domain is rare, although these factors play a prominent role in the scientific literature on safety and human error (e.g., [17]).

Organizational factors can produce system vulnerabilities that adversaries may exploit, either directly or more typically indirectly by capitalizing on increased likelihood of human errors and lapses in judgment that place stressed workers at risk of being deceived by social engineering scams. Management and management systems may fail to assign sufficiently qualified personnel to tasks or to provide sufficient materials resources [18]. Increased errors or lapses in judgment may be caused by work environments or work planning and control systems that negatively impact employee satisfaction or cause stress. Security practices are often difficult and confusing for an average computer user, and usage errors caused by these difficult security systems can yield serious consequences [19]. Systems that are difficult to understand or use are negatively perceived by users and are less likely to be used [20]. Difficulty using security systems may also encourage users to employ shortcuts around these system processes, which may make them more susceptible to UIT incidents.

Organizational factors are difficult to identify as contributing factors to socially engineered exploits and are difficult to change. UIT social engineering exploits are evolving so rapidly that organizational policies and practices cannot be created quickly enough. In addition, organizational staffing involves a variety of educational backgrounds, often not from the computer sciences, which can encumber the identification of, and warning communications about, potential exploits. Organizations often must balance operational goals (e.g., short product development cycles, multiple product release dates per quarter) with security goals (e.g., protecting intellectual property and other assets from adversaries) to maintain a competitive edge in the market. Historically, many organizations value operational goals above security goals.

C. Human Factors

Despite the best organizational efforts to educate users or impose security practices and security control systems and safeguards, social engineering scams, especially phishing schemes, continue to succeed. A number of studies and research papers emphasize the need to better understand the psychological aspects of social engineering exploits—why people fall for these scams—to develop more effective security practices, systems, and training approaches to combat them. Much of the research is focused on phishing and spear phishing exploits, although the findings may be generalizable to social engineering threats. Research suggests that human factors may contribute to increasing human errors in the context of UIT incidents.

Lack of attention (preoccupation, distraction) is identified as a contributing factor in many studies. Users tend to not pay attention to the source, grammar, and spelling in a phishing email, instead focusing disproportionately on urgency cues [21]; and they may miss cues in the address and status bars of emails [9][22]. In addition, high cognitive load (high subjective mental workload) can cause narrowing of attention. Workplace stressors (e.g., organization-imposed time pressures) contributing to higher levels of subjective mental workload tend to negatively impact human performance by, for example, narrowing visual attention such that important cues attributed to malicious activity may be missed [23][24] and by reducing cognitive resources needed for effective job performance [25][26][27]. Lack of knowledge, memory failure, or faulty judgment or risk perception are potential factors in UIT risk. For example, knowledge or memory deficits may underlie the inability to recognize design inconsistencies that distinguish real and fake error messages [28][9][5]. Human factors may also account for the observed tendency for users to ignore the organization's warning notices against phishing attempts [8]. Human decisions tend to be biased and are not purely logical, and an individual may devote insufficient cognitive resources for correct reasoning and judgment [29]. An example of decision making bias occurs when individuals tend to think that threats are highly unlikely (e.g., they underestimate the abilities of social engineering attackers and overestimate the defensive capabilities of organizational security systems), and consequently ignore such threats [30]. Some users feel that use of strong security features will impede their job [22].

Annoyance with popup messages may lead users to click on fake popups [28]. Users may also lack awareness of potential risks involved in clicking fake popups [28]. Risk tolerance/poor risk perception is a factor in social engineering UIT risk. A high-risk-taking or risk-tolerant individual may exhibit risky behavior despite cybersecurity training, while a risk-averse individual may be less likely to knowingly take risky actions [11].

Stress/anxiety due to workplace conditions negatively affects employee performance [18]. Heavy or prolonged workload and constant time pressure may be correlated with higher task error rates; time pressure negatively impacts even well-trained individuals [31].

D. Summary and Implications of Research Findings

In summary, many studies emphasize the need to better understand the psychological aspects of social engineering exploits in order to develop more effective security practices, systems, and training approaches to combat social engineering. Some social engineering campaigns may be so well-crafted that individuals may still be exploited no matter what countermeasures (e.g., training, policies, etc.) are employed. For the less sophisticated campaigns that offer perceptible cues that a message is potentially exploitive, some human factors discussed above may predict the probability of being exploited. Several studies reported that users tend to ignore or do not recognize cues that a particular socially engineered message is malicious. A possible reason includes a lack of attention to these cues or a lack of knowledge about the exploitive nature of the message. In addition, the narrowing of attention can be exacerbated by high cognitive load (high subjective mental workload). There is a need for more research to identify other possible explanations for this result.

Regardless, phishers exploit these cognitive limitations of network users through visual deception to spoof legitimate email messages or websites. In addition, phishing schemes exploit a tendency for humans to focus disproportionately on urgency cues (i.e., the message urges the reader to act quickly). Susceptibility to social engineering attacks also may be traced to problems with poor judgment or cognitive biases: people sometimes underestimate the likelihood of the threats and thus ignore them. Because the vast majority of email and online experiences are scam-free, people can habituate to cues and consequently miss the phishing cues, a common phenomenon under conditions of high workload.

Risk tolerance and risk perception, as well as employee values and attitudes, represent other significant human factors to be considered in addressing social engineering threats. It is evident that organizations might decrease vulnerability to social engineering UIT by identifying more risk tolerant employees and adopting management practices and training that foster greater conformance with security policies.

The use of deception and obfuscation in social engineering UIT incidents, particularly phishing, presents special challenges for research aimed at developing effective mitigation strategies. Deceptive practices that exploit human psychological limitations and vulnerabilities are all the more challenging because the adversaries continue to change tactics. No matter how skilled, savvy, or trained an organization's employees are, there will always be a chance that a phishing campaign will succeed, especially because it takes only one individual to succumb to the scam to open new opportunities for the social engineer attacker to execute further exploits against the organization. Thus, the research community as well as responsible organizations and stakeholders are obligated to continue research and information gathering to inform the development of effective training and mitigation tools. Indeed, one implication of the increasing sophistication of social engineering attacks is the need to continue to examine these threats so that new information can be incorporated into updates of training and mitigation strategies. The next section provides a current status update on characteristics and patterns that we have observed to date, based on a small but growing collection of social engineering UIT case studies.

V. Case Representation and Sample Data

Though case studies do not constitute a valid research method for making generalizable inferences, without them researchers are left to infer what factors and parameters are important. Collecting and analyzing UIT social engineering case studies is helpful for identifying factors and relationships that may be addressed later in experimental and observational research, enabling statistical testing of hypothesized relationships (e.g., causal, correlational, moderating, mediating, predictive) between factors and incidents. By informing experimental and observational research, case study research improves the validity and generalizability of these hypothesized relationships.

For ease of presentation and to help reveal certain patterns, we categorize cases into single- or multiple-stage attacks. This categorization reflects our observation, based on examining cases collected, that many of the incidents may be decomposed into separate stages that share certain common characteristics that make up patterns or building blocks of incidents.

Although clients or users may be considered to have a business relationship with the organization, they would not necessarily be considered organizational insiders. Therefore, an argument may be made to exclude cases that take advantage of an organization's clients (e.g., banking customers). On the other hand, organizations have a vested interest in discouraging or preventing social engineering attacks aimed at their customers—these attacks can damage the organization's reputation and cause loss of customers and revenue. Thus, organizations may take steps to help prevent or combat social engineering threats to information security, such as by informing customers about these threats, how to recognize such threats, and clarification about their privacy and security policies (including identifying the kind of information that is requested from

clients via email, and information such as passwords, which are never requested). Therefore, we have included cases of this nature in our database.

We created an incident template to represent UIT social engineering incidents that we have collected from sources such as Internet searches and reports referenced in the literature. Examples are described below (the full set of cases is reported in [32]).

A. Example of Single-Stage Attack

Figure 2 shows the incident file for a single-stage attack. In this case, the targets of the exploit had all been trained in identifying and resisting phishing attempts after a previous, similar attack. However, the phisher was able to provide a very realistic email (high obfuscation) to entice potential UITs, and about five staff members succumbed. The breach involved lists of visitors and their identifying information, so this constituted a serious security threat. However, the organization was able to resist repeated attempts to access more secure types of information.

Incident ID 24

INCIDENT ID: 24

INDUSTRY: Government

STAGING: Single

INCIDENT: Employees were duped by a phishing email about HR benefits that exploited a zero-day vulnerability and downloaded malicious code. The malware masked itself on systems and was designed to erase itself if it tried to compromise a system and was unsuccessful.

BREACH: Only a few megabytes of encrypted data were stolen, but the organization failed to recognize additional dormant, malicious code.

OUTCOME: The organization was forced to disconnect Internet access after administrators discovered data being externally siphoned from a server. After initial shutdown, the organization restored external email access but prohibited attachments.

RESPONSE: This was the second widespread social engineering attack. The organization implemented extensive training after the first. The specific response to this incident is unknown.

Figure 2. Sample Incident Description (Single Stage Attack)

B. Example of Multiple-Stage Attack

Figure 3 shows the incident file for a multiple-stage attack against a bank. As is common in many of the cases we collected, information relating to possible contributing factors was difficult to obtain and was gathered by carefully examining numerous, separate information sources. This case resulted in a lawsuit with considerable numbers of court filings of documents and testimony from both the bank and the manufacturing firm. Details about this attack are available, and a thorough study of this case greatly illuminates the nature of many types of phishing exploits and insider responses.

Incident ID 5

INCIDENT ID: 5

INDUSTRY: Banking and finance, manufacturing

STAGING: Multiple

INCIDENT: The phisher impersonated the company's bank, requesting information to address security concerns. The insider clicked on a link in a phishing email and entered confidential information.

Stage 1 - phishing to multiple bank customers

Stage 2 - spear phishing to executives with likely wire transfer authority

BREACH: The disclosure included credentials and passwords that enabled outsiders to transfer funds to accounts in several countries.

OUTCOME: The bank was able to reverse 70 percent of total money lost.

RESPONSE: The company recovered the remainder in a court settlement resulting from a lawsuit brought against the bank.

Figure 3. Sample Incident Description (Multiple-Stage Attack)

VI. Analysis

In all, there are 28 cases in our UIT social engineering database. All of the cases were found online, such as through search engines. Three of the cases (10.7%) have more than one source reference. A breakdown of the sources is as follows:

- news articles: 25/28 (89.3%)
- journal publications: 1/28 (3.6%)
- blog: 1/28 (3.6%)
- other: 1/28 (3.6%)

A. Inferred Contributing Factors

Following the research described in Section IV, we examined data relevant to demographic, organizational, and human factors as possible contributing factors:

Demographic Factors

- **Gender**—The gender of victims is stated directly in some of the case study reports. In other cases, we inferred the gender of the victim based on the case description; some involved both male and female victims. While many attacks on financial institutions identified the victim as male, there is insufficient data to enable conclusions about gender effects.

- **Age**—For software development groups, we assume that victims would likely be in their 20s and 30s; the ages of midlevel financial or government victims are probably somewhat higher. However, no conclusions can be reached based on the case studies we examined.

Organizational Factors

- **Security systems, policies, and practices**—Many of the case studies reveal organizational policies and procedures. In some cases the victims violated those policies, but most incident summaries do not provide sufficient information to determine whether or not these factors are involved: Organizations generally do not have an automated means of tracking employee actions or to warn employees of possible violations.
- **Management and management systems**—Many of the cases reveal that a simple login identification and password provided the attackers with access to internal emails, company data, and even access to entire computer networks. In one case, the attacker seemed to have attained computer network access directly from the login. Organizations must regularly perform extensive security audits to determine how best to improve internal controls; they cannot rely on security established during initial installation of a system.
- **Job pressure**. Certain industry categories such as news services place a premium on obtaining and distributing information as quickly as possible. Employees of such organizations may be more prone to outside influence from social engineering due to this pressure.

Human Factors

- **Attention**—At least one case study identified fatigue as a contributing factor: the phishing message was received late at night, and the individual responded without completely analyzing the message. A phisher—in this case a spear phisher—may have information about work hours or other conditions that could affect the likelihood of an attack's success.
- **Knowledge and memory**—Many of the case studies included information about prior staff training. Organizations that provide such instruction indicate that even with training a large percentage of employees respond to phishing attacks (this is consistent with research findings). Constant refreshers or other means should be applied to maintain user knowledge.
- **Reasoning and judgment**—Some case studies indicated that an employee's safeguards were lowered, perhaps because of the realistic nature of the phishing message and pretext created through reverse social engineering (i.e., offers to assist in preventing or addressing outside attacks, solving bank account problems, or supporting system operations).
- **Stress and anxiety**—In one case (multiple-stage attack exhibited in Figure 2, above) the victim knew that the organization and its customers were receiving phishing emails. This may have increased his desire to accept an offer of mitigation that appeared legitimate, though it was actually another phishing attack.

B. Attack Progression Analysis

We examined the collected cases using a number of conceptual models or analytic methods to gain a better understanding of the problem of social engineering UITs, to highlight common features across multiple exploits, and to identify possible mitigation strategies. Here we briefly describe findings from an attack progression analysis and characterization of attack patterns; more details are discussed in [2].

Social engineering attacks leverage human psychology and how humans interact with technology. Phishing illustrates how social engineering exploits work. Most phishing attacks have three components: the hook, the lure, and the catch [33]. The hook is the seemingly legitimate email form, website, or mechanism used by the phisher. The lure is the incentive aimed to trick the potential victim into taking the bait. The catch is the information acquired in the attack. Phishing attacks use different approaches toward social engineering of potential unintentional insiders.

Phishing emails can be simple or more highly sophisticated. In the simplest cases, the attacker sends an email message offering a reward, such as gifts or free trips, or reduced insurance or utility rates. The message generally directs the reader to a URL where the user enters a system password and other login information. In more sophisticated cases, the message may have the look and feel of company letterhead. Again, the company may be a cell phone provider, a bank, or the insider's own organization. The message generally serves the same purpose as the simple email message described above.

Multiple-stage social engineering attacks are common. The first stage uses one of the above methods to obtain account privileges on the UIT's computing resources. The attacker then uses the login information to search the UIT's internal system for detailed information about employees, company policy, or privileged data. The attacker uses insider knowledge about higher-level personnel to implement spear phishing attacks. These messages, customized and targeted at individuals rather than large groups, tend to contain information specific to the addressee and to specific internal enterprise conditions. The attacker's goal is to obtain administrator privileges that may allow the attacker to access to proprietary data, interfere with internal financial operations, or cause damage to operations through a denial of service or other attacks.

It is useful to describe these attacks using a "kill chain" analysis, which originated in the military as a way of analyzing attacks in the physical world. The approach decomposes attacks into a sequence of phases, with the aim of making them easier to understand and defend against [34]. The technique was adopted for use in cybersecurity by Cloppert [35], who analyzed cyber-attacks into the now-classic six phases: Reconnaissance, Weaponization, Delivery, Exploitation, Command-and-Control, and Exfiltration. To describe the social engineering attack

Figure 4. Workflow Pattern Showing Phases of a Single-Stage Phishing Attack

progression, we use a variation of the kill-chain model, with some customizations in the delivery, exploitation, and command-and-control steps to accommodate the specifics of social engineering. A single-stage attack chain for a phishing attack is shown in Figure 4.

The steps shown represent general building blocks on which more complicated attacks may be based. Each phase of the attack has different objectives that can change opportunistically depending on what information is captured during the social engineering operation. The general workflow pattern allows for this flexibility. In the first phase, the attacker researches possible targets. Based on information gathered, the attacker prepares phishing artifacts. Following this planning and preparation phase, the attacker executes the phishing operation by sending phishing emails to recipients in the target organization. While most recipients do not respond, those who do respond become UIT victims. In the Response and Information Capture phase, the UIT unwittingly sends account information to the attacker's system. When this information is received, the attacker conducts the final phase of the attack by using the account access to plant malware or take other malicious measures.

The multiple-stage attack follows a similar pattern, but once the attacker has UIT system access, the attacker identifies other potential UITs and subsequently directs social engineering at them. The attacker may also use the access gained to probe the UIT's system to obtain various forms of internal system information. The workflow diagram in Figure 5 shows the general attack chain. This diagram identifies the ordering and decision processes involved in each phase of the exploit.

C. System Dynamics Modeling

Another way to describe and characterize social engineering exploits is to use system dynamics modeling. System dynamics modeling helps analysts model and analyze critical behavior within complex socio-technical domains as it evolves over time [36]. Here we describe a system dynamics model that captures the complex

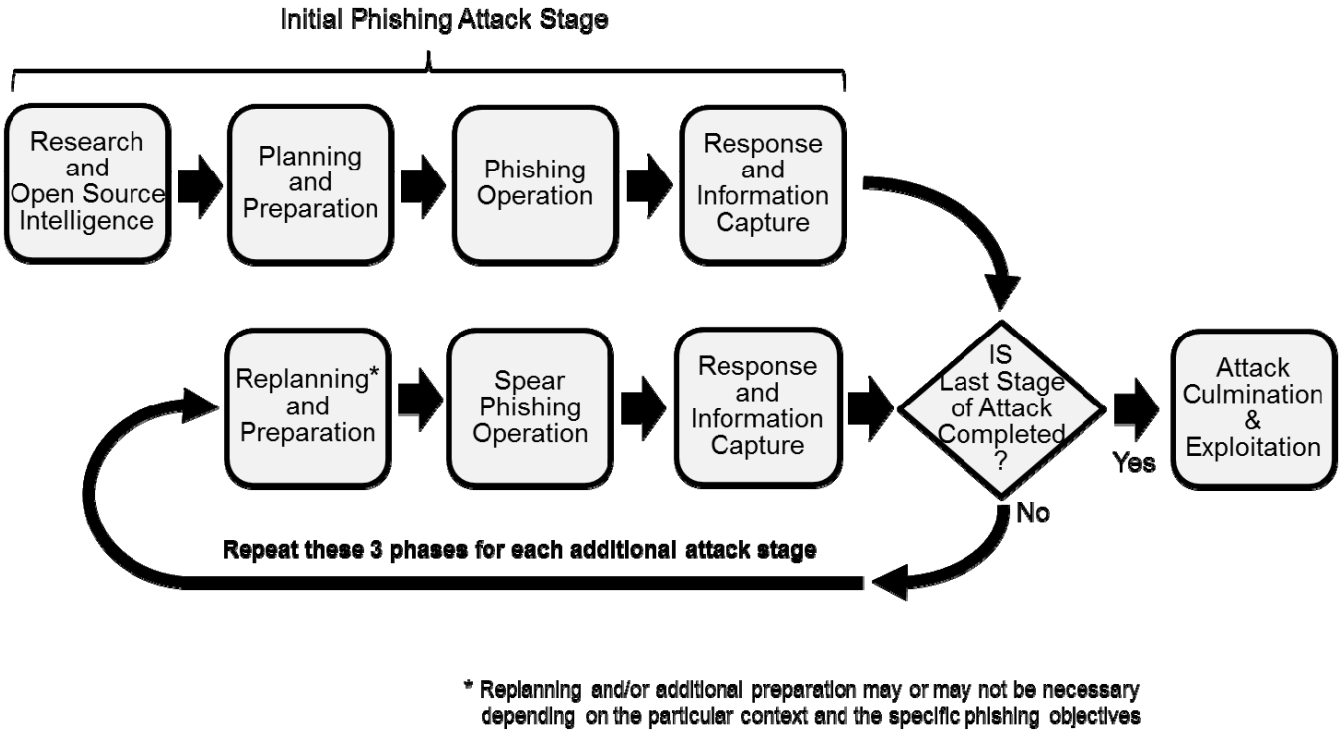

Figure 5. Workflow Diagram Attack Chain for Multiple-Stage Phishing Exploit

interactions within a social engineering scenario. We use this modeling approach to focuses on key aspects of the social engineering UIT incident that may be leveraged to identify mitigation strategies.

The modeling approach uses causal loop diagrams that show qualitatively how related variables affect each other [37]. The nodes indicate variables and the connecting arrows show the relationships between them. Arrows are labeled to indicate how the variable at the arrow's source influences the variable at the arrow's target. Basically, a positive influence indicates that the values of the variables move in the same direction and so is labeled with an "S," whereas a negative influence indicates that they move in the opposite direction, indicated by an "O" label. A connected group of variables can create a path that is referred to as a feedback loop. The type of feedback loop is determined by counting the number of negative influences along the path of the loop. An odd number of negative influences indicates a balancing loop; an even (or zero) number of negative influences indicates a reinforcing loop. Balancing loops often represent actions that an organization takes to mitigate (or control) a problem. Reinforcing loops often represent the escalation of problems but may include problem mitigation behaviors.

Our model of the social engineering UIT incident makes use of several well-defined structures in system dynamics models. A very relevant structure for social engineering is the Confirmatory Bias Loop (see Figure 6). This is a reinforcing feedback loop that reflects the tendency of decision-makers to pay attention to data that supports (or is

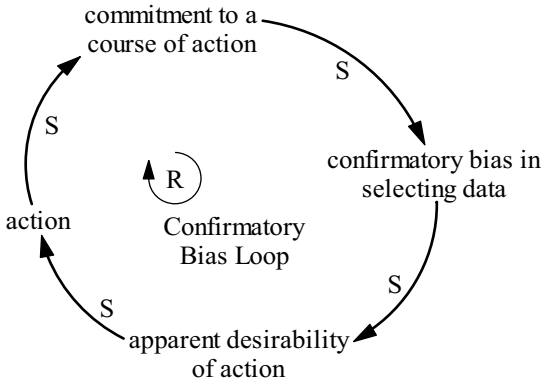

Figure 6. Confirmatory Bias

at least consistent with) their past decisions and to downplay conflicting information [38][39]. This bias can skew the basis for decision-making so that alternate decisions are overlooked in favor of the preferred decision. The associated feedback loop portrays the reinforcing nature of a decision-maker's cognitive process.

Figure 7 depicts a more complete system dynamics model that integrates and extends causal feedback loop and confirmatory bias influences described above. The outsider's planning and preparation are represented in feedback loops R1 and B1 associated with phishing and spear phishing exploits. In particular, the B1 loop is another view into the initial phishing attack stage of the multiple-stage exploit represented in Figure 5. This increases the

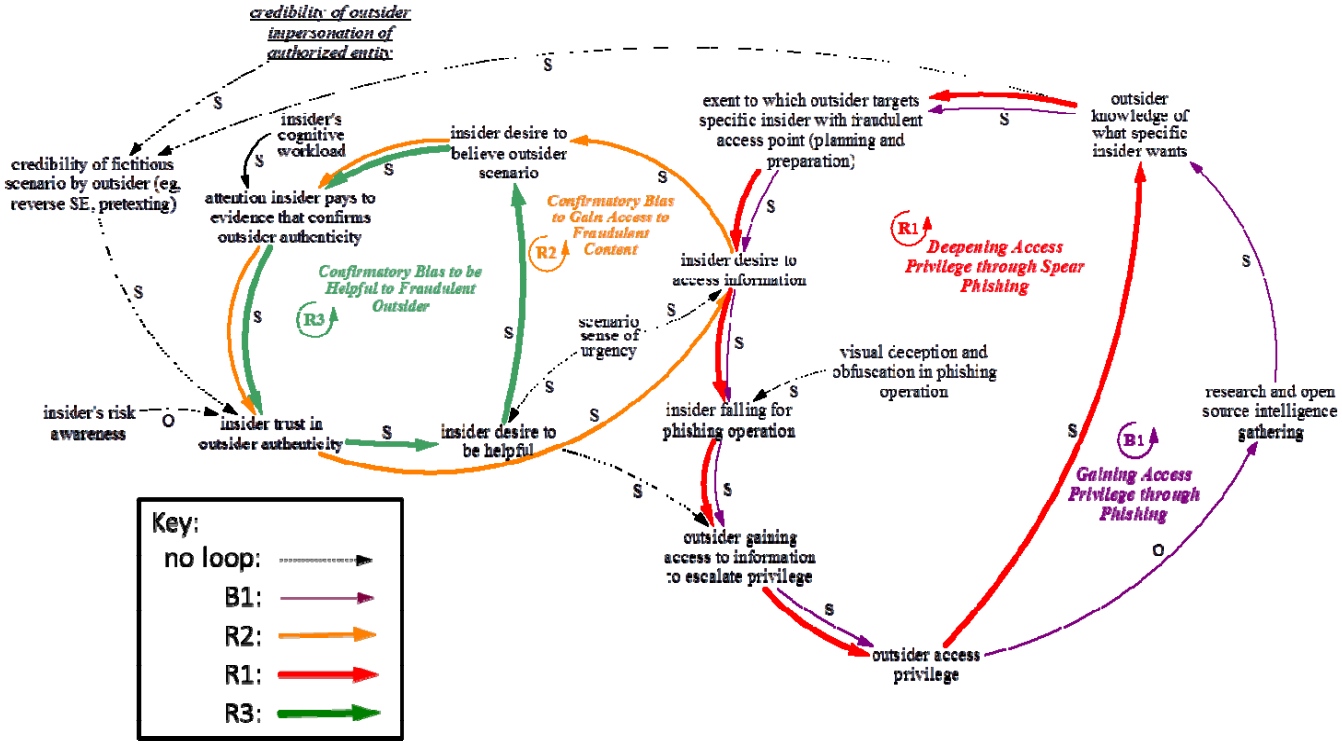

Figure 7. Causal Loop Diagram of Social Engineering of Insiders by Outsiders

outsider's access to the point that a spear phishing attack can take place in subsequent attack stages. The insider's desire to access information is reinforced through the outsider's planning and preparation from the R1 and B1 feedback loops, as well as the insider's confirmatory bias. This confirmatory bias is depicted for two situations: In the first, the insider desires access to information supplied by the outsider's created (deceptive) scenario, as depicted in the R2 feedback loop. The second is where the insider desires to be helpful to the malicious outsider in need as depicted in the R3 feedback loop. Both loops portray the reinforcing of trust in the outsider's authenticity and the subsequent desire to access information or to be helpful. The key to the confirmatory bias tendency is that the growing desire to believe the scenario put forth by the outsider leads the insider to focus on evidence that confirms the legitimacy of the scenario and ignore evidence to the contrary. The greater the insider's desire to access the information provided by the outsider (loop R2) or to help the outsider (loop R3), the more likely is the insider to provide the outsider undeserved access, resulting in the outsider attack, as shown in the bottom part of the figure. Trust in the outsider's authenticity is enhanced by a credible scenario provided by the outsider, supported by an accurate impersonation, as seen in the upper left. The model depicted in this figure reflects the impact of cognitive limitations that were discussed in Section IV, namely:

- High levels of cognitive workload can increase the chances that the insider will believe the deceptive scenario painted by the outsider and trust the outsider's authenticity.
- The insider's overall awareness of the risks of social engineering also plays a role in the trust the insider places with the outsider's scenario.
- Creating a sense of urgency increases the chances of falling for the deception, either as a need to be helpful or to access the information (phishing exploit) provided by the outsider.

VII. Implications

While it is beyond the scope of this work to examine current mitigation practices, the foregoing discussion and characterization of social engineering attacks in terms of possible contributing factors (especially organizational and human factors) and patterns help to inform a brief consideration of possible mitigation approaches and strategies. Here we briefly discuss and speculate about possible implications for mitigations suggested by systematic analyses of patterns and models of social engineering exploits discussed in previous sections.

A. Implications of Patterns and Characterizations

The kill-chain pattern indicates that the adversary must progress successfully through each phase of the chain before the desired objective can be achieved. Thus the chain and the adversary can be disrupted by a single mitigation that is applied successfully to just one phase of the chain. This observation has strong implications for concepts for and approaches to mitigation. In the present context, a sophisticated multiple-stage social engineering attack (such as one involving phishing followed by spear phishing) aims to breach successive layers of organizational defenses by progressively gaining access through social engineering methods. The attack continues iteratively, and sometimes opportunistically, to take advantage of individual or organizational responses until the final layer of defense is breached. Because the ultimate success of a multiple-stage attack depends on the success of each of the individual (i.e., iterative) stages leading up to the final attack, the kill-chain approach affords a UIT organization multiple opportunities to detect and defeat such attacks.

A systematic analysis of patterns in workflow diagrams reveals points at which opportunities for mitigation arise. Listed below (and illustrated in Figure 8) are various mitigation strategies that apply to different phases of a single-stage social engineering attack (this is generalizable to multiple-stage attacks):

- *Research and Open Source Intelligence phase*—The organization may limit the amount of information that is publically accessible. While it is not possible or desirable to eliminate this information completely, the organization may benefit from instilling controls and safeguards in its public relations and information dissemination processes to avoid excessive disclosures. Similarly, employees may be given direction or policies about avoiding certain disclosures on social media sites.
- *Planning and Preparation phase*—One possible approach is to make it difficult or expensive to copy organizational artifacts that make a spoofing email or website look legitimate. This could impair or discourage attacker's efforts to masquerade or impersonate organizational assets. Anticounterfeiting strategies such as encrypted emails are well known, but not commonly used.
- *Launch Operation phase*—Improved training and awareness are an organization's most potent mitigation tools for thwarting social engineering exploits that target human psychological characteristics and limitations. Periodic injection testing and associated training may be used to maintain staff vigilance and knowledge about the most current social engineering tactics. Organizations also should strive to maintain productive work attitudes and information security awareness through human factors and organizational practices. Effective management and resource planning can help ensure employee productivity and avoid stressful work environments that may lead to errors in judgment.
- *Information Capture and Culmination/Exploitation phases*—Organizations should enable and maintain improved tools for computer and network defense cyber monitoring to keep up with the rapidly evolving kinds of exploits that adversaries use. Cybersecurity systems that locate malware and other threats include antivirus,

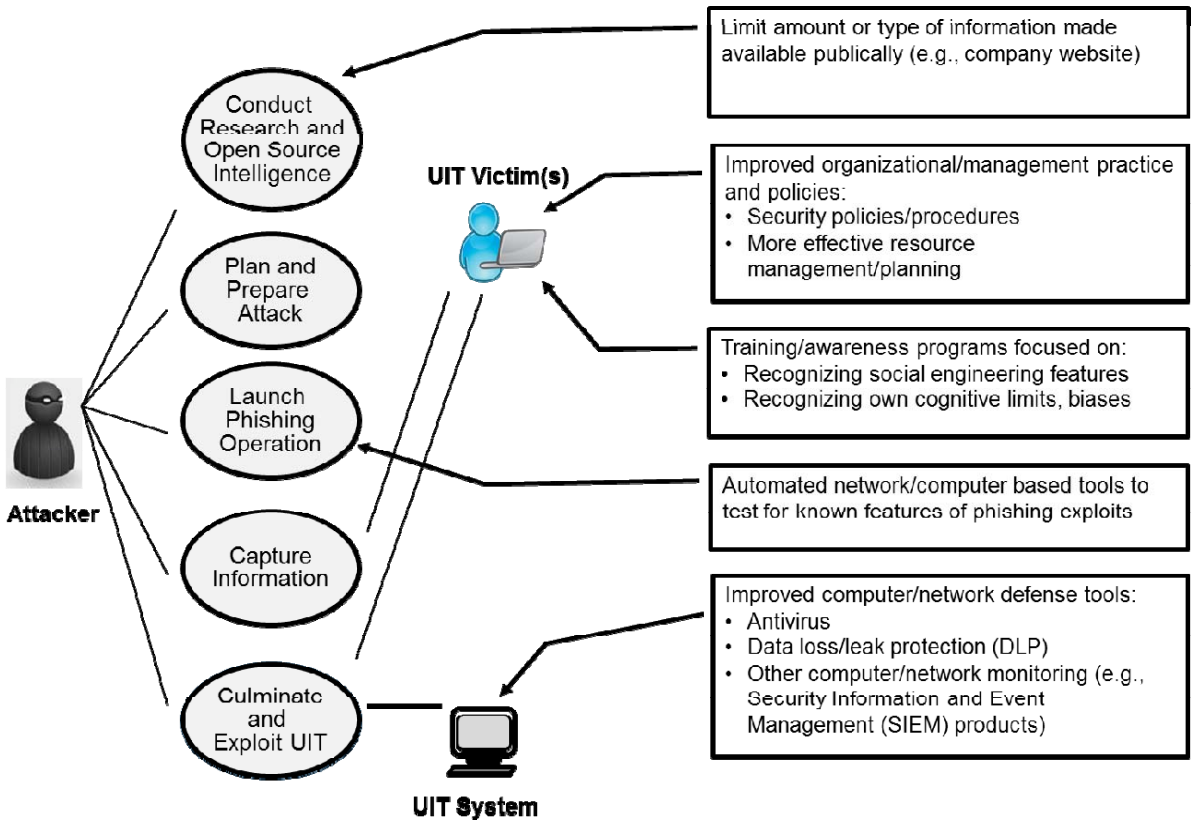

Figure 8. Mitigation Strategies that Apply to Different Phases of an Attack

data loss and leak protection (DLP) tools, and Security Information and Event Management (SIEM) products.

B. Implications of System Dynamics Model

The system dynamics model may be used to help identify possible mitigation strategies. Consider a hypothetical case depicted in Figure 9. This shows how the reinforcing feedback loops (R1, R2, and R3) involving the escalation of the phishing exploit and the cognitive limitations of the insider can be dealt with by balancing feedback loops:

- Feeback loop B2 represents organizational processes aimed at reducing the effectiveness of social engineering exploits in taking advantage of insiders. Feedback loop B2 involves the recognition of the exploitation by the organization and improved training on the nature and risks of social engineering to organizational insiders. Specifically, the organization provides more effective training and awareness about how malicious, outside attackers use obfuscation and social engineering techniques to deceive insiders. Such training may involve various topics relevant to human factors described in Section 4, aimed at raising self-awareness about cognitive limitations and biases, fostering greater security awareness and more accurate risk perception, and encouraging more diligent application of computer security policies.
- Feedback loop B3 represents organizational processes aimed at reducing the effectiveness of early-stage social engineering activities that aim to acquire intelligence about the organization that may be used in an initial phishing attack. Specifically, the mitigation approach seeks to reduce the amount of publicly available information about the organization and its employees that malicious outside social engineers can use to develop initial attack plans and associated artifacts for luring insiders into their traps.

Not shown in the figure are other possible opportunities for mitigation that would be aimed at different parts of the system dynamics model. For example, mitigation in the form of more effective firewalls or automated tools for recognizing flaws in phishing emails might be applied to balance spear phishing efforts in R1. While this example does not reflect all possible approaches that adversaries might take in executing the social engineering attack, it is instructive and representative of how the analysis can reveal opportunities for applying measures to circumvent (balance) the actions of malicious attackers implications drawn from the example are not exhaustive. And, of course, all of these mitigation approaches are hampered by time delays. The longer the delay associated with the organization's mitigation action, the less effective it will be in preventing the successful execution of social engineering exploits.

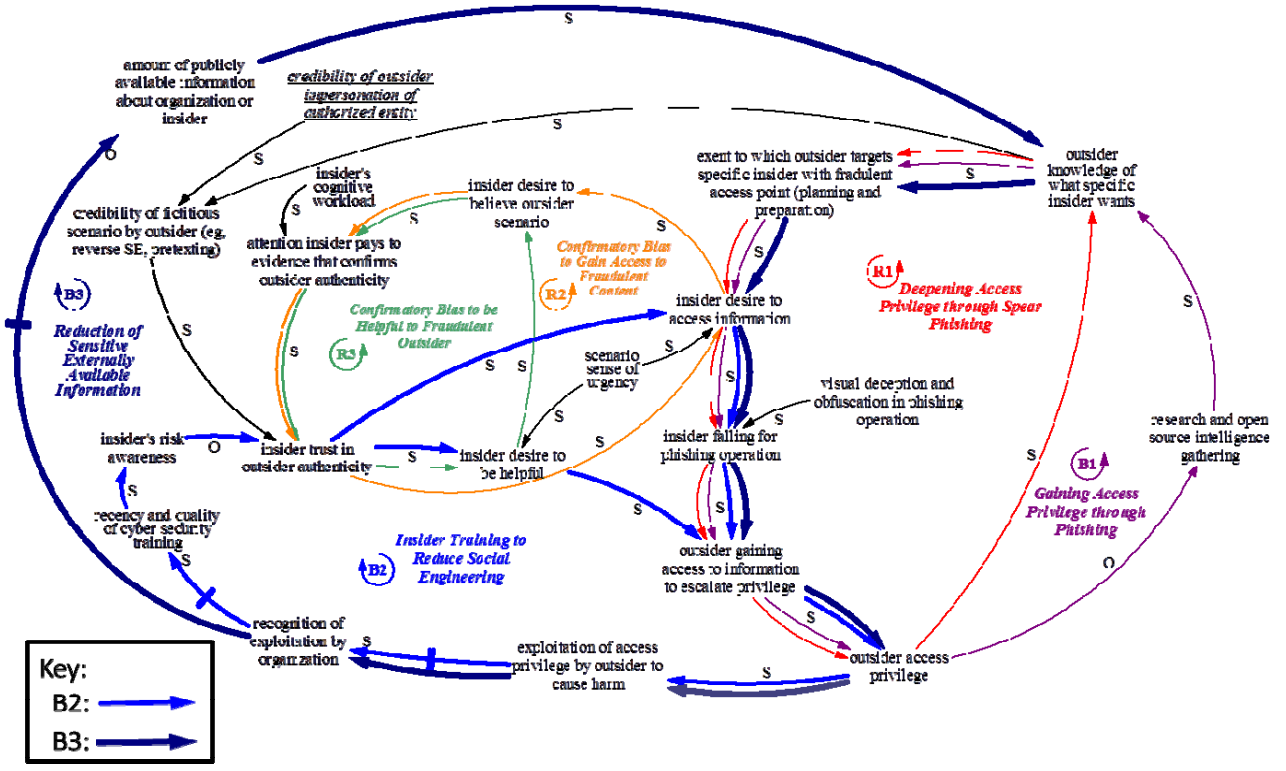

Figure 9. Causal Loop Diagram of Avenues for Social Engineering Mitigation

C. Summary and Conclusions about Mitigation

Our analysis of collected case studies reveals a number of commonalities or patterns that may be useful to take into account when developing mitigation tools or strategies. Using kill chain concepts, we recognize common features, or building blocks, in single-stage and multiple-stage social engineering attacks. Each stage contains recognizable patterns or building blocks that compose the attack. Each stage includes multiple phases. To be successful, the attack must succeed at every phase. Some phases represent actions of the attacker, while other phases represent actions of UIT victims. Mitigation strategies and tools should be crafted to target specific characteristics in each attack phase.

Our review and analysis of research and case studies suggest the following mitigation strategies to reduce the effectiveness of social engineering attacks:

1. Organizations should examine their management practices to ensure that they meet human factors standards that foster effective work environments to minimize stress (e.g., minimizing time pressure and optimizing workload) and encourage a healthy security culture. Such management practices must be applied with due consideration being given to productivity-vs-security trade-offs. Also, because employees may perceive information security compliance as interfering with job functions, it is important for organizations to allocate a certain amount of employees' time to fulfilling the compliance requirements.
2. Organizations should develop and deploy effective staff training and awareness programs aimed at educating users about social engineering scams, including learning objectives to help staff attend to phishing cues, identify deceptive practices, and recognize suspicious patterns of social engineering exploits.
3. Research is required to develop more effective network and workstation monitoring tools to recognize attributes of social engineering artifacts (e.g., emails).
4. The research and stakeholder community should develop mitigations that apply to specific attack phases, such as the following:
   - *Research and Open Source Intelligence phase*—Both the organization and individual employees may benefit from limiting the amount of information available on organizational websites or individuals' social media sites, which might be exploited by outsiders.
   - *Planning and Preparation phase*—Efforts should be made to make it difficult or expensive to copy organizational artifacts that make a spoofing email or website look legitimate. Anticounterfeiting strategies that allow encrypted emails are well known but not commonly used.

- *Launch Operation phase*—Phishing exploits target human psychological characteristics and limitations, so improved training and awareness are an organization's most potent mitigation tools. Periodic injection testing and associated training may maintain staff vigilance and knowledge about the most current social engineering tactics.
- *Information Capture and Culmination and Exploitation phases*—Organizations should enable and maintain improved tools for computer and network defense cyber monitoring to keep up with the rapidly evolving kinds of exploits that adversaries use.

VIII. Conclusions and Recommendations

A challenge in conducting research on the contributing factors and mitigating strategies of socially engineered UIT incidents is the lack of peer-reviewed academic research on the topic. Additionally, the lack of quality reporting of socially engineered UIT incidents and case studies makes it difficult to study contributing factors; this is in part due to concerns about security, proprietary business practices, and litigation as well as the immaturity of reporting processes.

The use of deception and obfuscation in socially engineered UIT incidents presents special challenges for research aimed at developing effective mitigation strategies. For example, some phishing campaigns can be so well obfuscated that they appear 100% genuine to humans, and the adversarial success rate is very high. Other, less-obfuscated messages capitalize more on human limitations (e.g., highly fatigued employees may have lower performance thresholds) to succeed. To add to the complexity, adversaries continually change their deceptive tactics. Despite these challenges, the research community as well as responsible organizations and stakeholders have an obligation to continue research and information gathering to inform the development of effective training and mitigation tools.

Countering the UIT social engineering problem poses major challenges to organizations, who must balance operational goals with security goals to maintain a competitive edge in the market. Because organizational policies and practices are resistant to change, it is a great challenge to keep up with the rapidly changing, increasingly sophisticated social engineering attacks. Some social engineering campaigns may be so well crafted that they can defeat the organization's best countermeasures (e.g., training and policies). Attackers succeed even if only one employee succumbs to an exploit, so an organization's strategy to combat UIT social engineering must be comprehensive and include cybersecurity tools in addition to up-to-date security practices and training.

Research is needed to further study possible contributing factors, particularly organizational and human factors. Additional case study data must be collected to increase understanding of characteristics of social engineering attacks. By characterizing and conceptually modeling the UIT social engineering problem, the research reported in this paper has sought to inform mitigation development efforts and identify research needs to more effectively combat UIT social engineering exploits.

IX. Acknowledgments

We wish to acknowledge and thank CMU intern Arley Schenker for contributions to this research.

This material is based upon work funded and supported by Department of Homeland Security under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center sponsored by the United States Department of Defense.

Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of Department of Homeland Security or the United States Department of Defense.

This material has been approved for public release and unlimited distribution.

CERT® is a registered mark of Carnegie Mellon University.

DM-0000592

References

[1] AlgoSec. "The State of Network Security 2013: Attitudes and Opinions." AlgoSec, Inc., 2013.

[2] CERT Insider Threat Team. *Unintentional Insider Threats: A Foundational Study.* Software Engineering Institute, May 2013.

[3] Peltier, T.R. Social Engineering: Concepts and Solutions, *Information Systems Security*,2006, 15:5, 13-21

[4] Laribee, L. *Development of methodical social engineering taxonomy*. Master's Thesis, Monterey, CA: Naval Postgraduate School. Amazon Digital Services, June 2006.

[5] Downs, JS, MB Holbrook, & LF Cranor. "Decision Strategies and Susceptibility to Phishing." *Symposium On Usable Privacy and Security (SOUPS)*, July 12-14, 2006, Pittsburgh, PA, USA.

[6] Downs, JS, M Holbrook, and LF Cranor. "Behavioral response to phishing risk." (Institute for Software Research, Paper 35), *APWG eCrime Researchers Summit*, Pittsburgh, PA, October 4-5, 2007.

[7] O'Brien, T.L. "Gone spear-phishin'." *The New York Times* (4 December 2005) http://www.nytimes.com/2005/12/04/business/yourmoney/04spear.html?pagewanted=1&ei=5088&en=2f313fc4b55b47bf&ex=1291352400&partner=rssnyt&emc=rss&_r=0

[8] Mohebzada, J. G., El Zarka, A., Bhojani, A. H., & Darwish, A. "Phishing in a University Community." *International Conference on Innovations in Information Technology (IIT),* 2012, 249-254.

[9] Dhamija, R., Tygar, J. D., and Hearst, M. "Why phishing works. In *Proceedings of the SIGCHI conference on Human*

Factors in Computing Systems, CHI '06. New York, NY: ACM, 2006, 581–590. http://dl.acm.org/citation.cfm?id=1124861

[10] Jagatic, T.N., Johnson, N.A., Jakobsson, M., & Menczer, F. "Social phishing." *Proceedings of the ACM, 50(10),* 2007, 94-100.

[11] Sheng, S, M Holbrook, P Kumaraguru, L Cranor, & J Downs. "Who Falls for Phish? A Demographic Analysis of Phishing Susceptibility and Effectiveness of Interventions." *28th ACM Conference on Human Factors in Computing Systems*, Atlanta, GA, April 10-15, 2010.

[12] Halevi, T, Lewis, J., and Memon, N. *Phishing, Personality Traits and Facebook*. Cornell University Library. 2013. http://arxiv.org/abs/1301.7643

[13] Digman, J. M. (1990). Personality Structure: Emergence of the five-factor model. *Annu. Rev. Psychol., 41*, 417-440.

[14] Alseadoon, I., Chan, T., Foo, E., & Nieto, J. G. "Who is more susceptible to phishing emails?: A Saudi Arabian study." *23rd Australasian Conference on Information Systems*, Geelong: December 3-5, 2012.

[15] Workman, M. "Wisecrackers: A Theory-Grounded Investigation of Phishing and Pretext Social Engineering Threats to Information Security." *Journal of the American Society of Information Science and Technology*, 2008, 59 (4), 662-674.

[16] Jakobsson, M., and Ratkiewicz, J. 2006. "Designing Ethical Phishing Experiments: A Study of (ROT13) Ronl Query Features," Pr*oceedings of the 15th international conference on World Wide Web*. Edinburgh, Scotland: ACM, pp 513-522.

[17] Decker, S. *The field guide to human error investigations*. Burlington, VT: Ashgate, 2002.

[18] Leka, S., Griffiths, A., & Cox, T. *Work Organization and Stress: Systematic Problem Approaches for Employers, Managers, and Trade Union Representatives. Protecting Workers Health Series, No. 3*. Geneva, Switzerland: World Health Organization, 2004.
http://www.who.int/occupational_health/publications/pwh3rev.pdf

[19] Whitten, A. and Tygar, J.D. "Why Johnny can't encrypt: A usability evaluation of PGP 5.0." *Proceedings of the 8th USENIX Security Symposium*, Washington, D.C., 1999.

[20] Venkatesh, V., Morris, M., Davis, G.B. and Davis, F.D. (2003). "User acceptance of information technology: Toward a unified view." *MIS Quarterly*, 2003, *27(3),* 425-478.

[21] Vishwanath, A., Herath, T., Chen, R., Wang, J., & Rao, H. 2011. Why do people get phished? Testing individual differences in phishing vulnerability within an integrated, information processing model. *Decision Support Systems, 51(3),* 576-586.

[22] Erkkila, J. "Why we fall for phishing." In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems CHI 2011*, Vancouver, BC, Canada: ACM, May 7–12, 2011.

[23] Houston, B. K. "Noise, Task Difficulty, and Stroop Color-Word Performance." *Journal of Experimental Psychology 82, 2* (1969): 403-404.

[24] Stokes, A. & Kite, K. *Flight Stress*. Ashgate, 1994.

[25] Davies, D. R. & Parasuraman, R. *The Psychology of Vigilance.* Academic Press, 1982.

[26] Hockey, G. R. J. "Changes in Operator Efficiency as a Function of Environmental Stress, Fatigue, and Circadian Rhythms." H*andbook of Perception and Human Performance (vol 2).* Edited by K. R. Boff, L. Kaufman, & J. P. Thomas.Wiley, 1986.

[27] Wachtel, P. L. "Anxiety, Attention and Coping with Threat." *Journal of Abnormal Psychology 73, 2* (April 1968): 137-143.

[28] Sharek, D., Swofford, C., & Wogalter, M. "Failure to Recognize Fake Internet Popup Warning Messages." *Proceedings of the Human Factors and Ergonomics Society 52nd Annual Meeting*, 2008, pp. 557-560.

[29] Kahneman, D. & Tversky, A. Prospect Theory: An Analysis of Decisions Under Risk. *Econometrica*, 1979, *47(2):* 263-291.

[30] Sandouka, H. Cullen, A.J., & Mann, I. "Social Engineering Detection Using Neural Networks." In *IEEE International Conference on CyberWorlds (CW'09)*. 2009, 273-278. http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?arnumber=5279574

[31] Lehner, P., Seyed-Solorforough, M., O'Connor, M.F., Sak, S. and Mullin, T. "Cognitive biases and time stress in team decision making." *IEEE Transactions on Systems, Man, and Cybernetics Part A: Systems & Humans*, 1997, 27, 698-703.

[32] CERT Insider Threat Team. *Unintentional Insider Threats: Social Engineering.* Software Engineering Institute, January 2014.

[33] Parrish Jr, J. L., Bailey, J. L., and Courtney, J. F. 2009. "A Personality Based Model for Determining Susceptibility to Phishing Attacks," *Decision Sciences Institute*, pp 285-296.

[34] Myers, L. "Cyber Kill Chain is a Great Idea." Infosec Institute, 2013.

[35] Cloppert, M. "Security Intelligence: Attacking the Cyber Kill Chain." *SANS Computer Forensics,* October 14, 2009. http://computer-forensics.sans.org/blog/2009/10/14/security-intelligence-attacking-the-kill-chain

[36] Sterman, J. Business Dynamics: System Thinking and Modeling for a Complex World. McGraw-Hill, 2000.

[37] Meadows, D. *Thinking in Systems: A Primer*, Chelsea Green Publishing, 2008.

[38] Sastry, A., "Archetypal Self-Reinforcing Structures in Organizations: A System Dynamics Perspective of Cognitive, Social, and Institutional Processes," *Proceedings of the International System Dynamics Society.* June 1989.

[39] Staw, B.M., and J. Ross. "Understanding behavior in Escalation Situations," *Science* 246:216-220. 1989.