Unintentional Insider Threats: A Review of Phishing and Malware Incidents by Economic Sector

Abstract

The research documented in this paper seeks to advance the understanding of the unintentional insider threat (UIT) that results from phishing and other social engineering cases, specifically those involving malicious software (malware). The research team collected and analyzed publicly reported phishing cases and performed an initial analysis of the industry sectors impacted by this type of incident. This paper provides that analysis as well as case examples and potential recommendations for mitigating UITs stemming from phishing and other social engineering incidents. The paper also compares security offices' current practice of UIT monitoring in the current manufacturing and healthcare industries' practice of tracking near misses of adverse events.

1. Introduction

Within government agencies, the second highest threat perceived by IT professionals is that of careless and untrained insiders [1]. In a survey published by cybersecurity consultancy Solar Winds, 42% of respondents noted that insiders may inadvertently pose nearly as many risks to their agency as deliberate, malicious hackers [1]. This research supports the conclusions in the *2013 Verizon Data Breach Report* that 47% of malware was downloaded through email attachments, 48% of hacking took place with stolen credentials, 77% of social engineering takes places through phishing, and 79% of social engineering takes place through emails [2]. Symantec's 2014 *Internet Security Threat Report* revealed that the trends in this space are working against organizations: there has been a 91% increase in targeted email attacks from 2012 to 2013, attackers are conducting longer campaigns,[1] and they are sometimes even following up these phishing attacks with phone calls to convince the unintentional insider to take immediate action without the opportunity to contemplate the risk involved [3]. The same report notes that the government (all levels) is the sector of the economy most frequently attacked by spear phishing; at 16% of all attacks are in this space.

Publicly released reports already provide an initial examination of the problem of unintentional insider threat [4] as well as how this problem is influenced by social engineering [5].

The first of those reports characterized the UIT by developing an operational definition, reviewing relevant research to gain a better understanding of possible causes and contributing factors,[2] and providing examples of UIT cases and the frequencies of UIT occurrences across several categories. The second report, building off of the first, looked into social engineering and how the unintentional insider can be manipulated into acting against the organization's interests. These reports also documented an initial design of a UIT feature model, suggested mitigation strategies, and outlined incident paths for UIT-HACK incidents.

One challenge in researching the UIT problem and developing effective mitigation strategies is that the UIT topic and its related incidents have gone mostly unreported. In particular, incident reports typically lack sufficient detail to inform analyses of potential contributing factors. Our team intended for our initial work on UIT cases to inform government and industry stakeholders about the problem and its potential causes and to guide research and development investments toward the highest priorities for countering UIT [4]. Our second report published on this topic sought to advance our understanding of UIT contributing factors by focusing on a major type of UIT incident, social engineering [5].

---

1   "An attack campaign is defined as a series of emails that: A.) Show clear evidence that the subject and target has been deliberately selected. B.) Contain at least 3 or 4 strong correlations to other emails such as the topic, sender address, recipient domain, source IP C.) Are sent on the same day or across multiple days" [4].

---

2   A *factor* is a situational element or feature that may or may not be related to the existence of the incident. A *contributing factor* is a factor that has been demonstrated to be associated as a causal factor of an incident. Because our research generally has not shown causal relationships, our usage of the term *contributing factor* should be interpreted as *potential contributing factor*.

The goals of this research project were to collect additional UIT incident data, build them into a set of social engineering cases, and add that set to our database (also referred to as the insider threat database). Another goal was to analyze UIT cases to identify possible behavioral and technical patterns and precursors, with a particular focus on UIT cases that involve social engineering, to inform future research and development of UIT mitigation strategies.

This paper focuses on the newly designated PHISHING/SOCIAL threat vector and its subvectors, Malware and Credentials. These threat vectors are influenced by the previous study on social engineering [5]. The intent of this paper is to identify the frequency of incident types (single-stage or multi-stage) that occur in different economic sectors within the United States. This research is based on the sample of incident cases the team has been able to collect from publicly available sources, which is often limited because organizations are reluctant to report incidents related to UIT. Due to limited amounts of information, we are unable to show a statistically significant difference between the types of incidents across industry sectors.

2. Defining and Characterizing the PHISHING/SOCIAL Threat Vector

Our initial research produced a working definition of an unintentional insider threat:

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,3 (4) causes harm or substantially increases the probability of future serious harm to the confidentiality, integrity, or availability of the organization's information or information systems [4].

In our second round of research, on social engineering, we recognized a need to modify the original definition slightly [5]. One change was to emphasize that the unintentional insider's actions occur largely without the insider's knowledge or understanding of their impact, so we added the term

"unwittingly" to the fourth part of the definition. (This definition uses the perspective of the unintentional insider, which differs from the broader definition of social engineering acts that includes the [malicious and/or intentional] perpetrator's perspective.)

A second change was to modify the description of the target of the incident to include assets other than the organization's information system, such as financial systems. The report on the second round of research revised the definition as follows:

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.

2.1. PHISHING/SOCIAL Threat Vector

In our previous work, we defined the UIT-HACK threat vector as

An outsider's electronic entry acquired through social engineering (e.g., phishing email incident, planted or unauthorized USB drive) and carried out via software, such as malware and spyware.

Through further research, we have determined that many incidents initiated through phishing and other social engineering are not carried out by using software, but by acquiring and misusing the victim's credentials to secured systems. Because the common elements of the two types of attacks are phishing and other social engineering, we created a new, larger category of threat vector, PHISHING/SOCIAL that subsumes UIT-HACK, renames it as the subvector Malware, and adds the new subvector of Credentials:

Malware (formerly UIT-HACK)—An outsider's electronic entry acquired through social engineering (e.g., phishing email incident, planted or unauthorized USB drive) and carried out via software, such as malware and spyware.
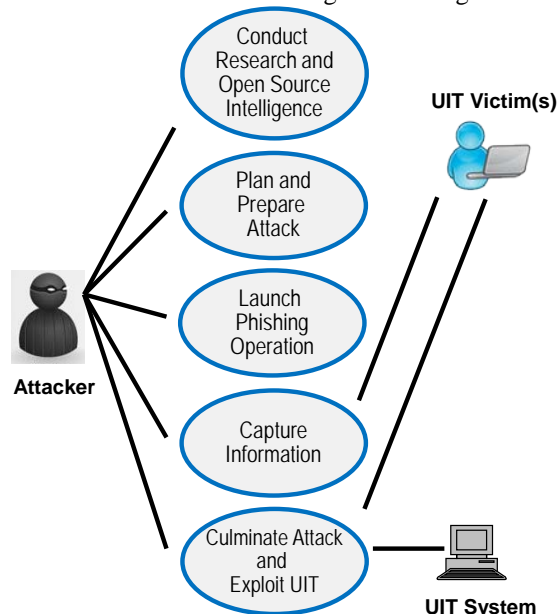
Credentials—An outsider's electronic entry acquired through social engineering (e.g., phishing email incident) and carried out through compromised credentials, including passwords and other identifying information.

The identification of the new threat subvector, Credentials, and the refinement of UIT-HACK allows researchers and those in operations to quickly

---

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

differentiate the two types of incidents and take the most appropriate mitigation actions.
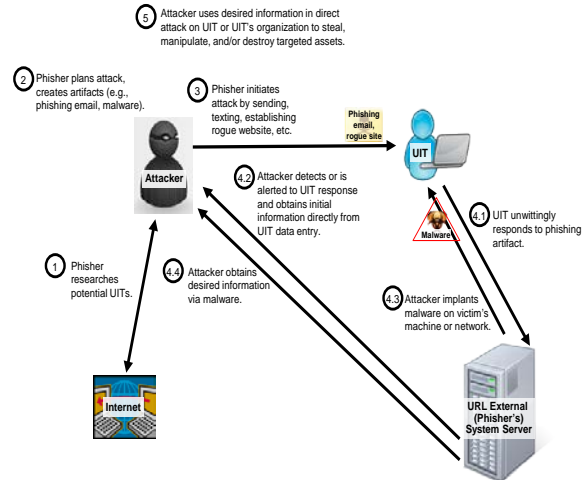
2.1.1. Single-Stage Incident. The incident progression for a single-stage incident generally comprises five steps, as shown in Figure 1. We used this variation of the kill-chain model, well known across the computer security industry, as a foundation and customized the delivery, exploitation, and command-and-control steps to accommodate the specifics of social engineering. The steps shown in Figure 1 represent the general building blocks on which more complicated incidents (multi-stage) may be based. Each phase of the incident has different objectives that can change opportunistically depending on what information is captured during the social engineering operation. The general workflow pattern of typical actions taken by an outsider allows for this flexibility.

In the first phase, the attacker researches possible targets. Based on the information gathered, the second phase, *Planning and Preparation*, progresses by the attacker preparing phishing artifacts. The attacker then executes the phishing operation by sending phishing emails to recipients in the target organization. Many recipients do not respond, but those who do respond may become a UIT. In the *Response and Information Capture* phase, the UIT unwittingly sends account information to the attacker's system. After the information is received, the attacker conducts the final phase of the incident by using the account credentials to gain access to the unwitting individual's machine and plant malware or take other measures directed against the organization.

The interaction view (Figure 2) shows each interaction and the exchanges that occur to carry out an incident. This view illustrates the collaborations of each element of the swim-lane view for a single-phase incident. The sequence of interactions shows the information exchanges during each phase of the incident.

Figure 2. Interaction View Showing Object Collaboration in a Single-Stage Social Engineering Incident

2.1.2 Multiple-Stage Incident. The multiple-stage incident follows a similar pattern as the single-stage incident, but once the attacker has 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 gather intelligence about the compromised systems or networks, and use the information to cause harm or develop subsequent spear phishing messages. The workflow diagram in Figure 3 shows the multi-stage incident chain. This diagram identifies the ordering and decision processes involved in each phase of the exploit.
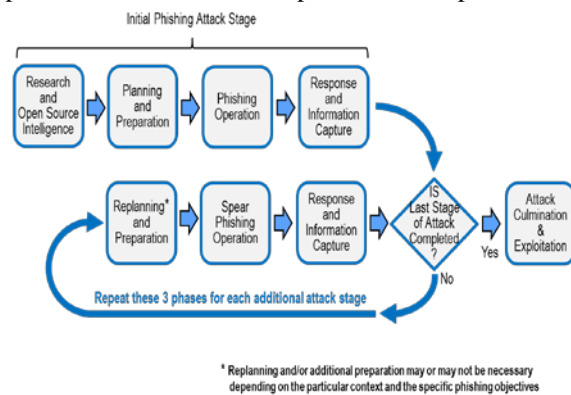
**Figure 3. Workflow Diagram Incident Chain for Multiple-Stage Phishing Exploit**

Figure 1. Use-Case Model for a Single-Stage Social Engineering Incident

3. Summary of Collected Cases

Case study research is not a valid method for making generalizable inferences. Yet without case studies, researchers are left to infer what factors and parameters are important. Collecting and analyzing PHISHING/SOCIAL and Malware 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.

3.1. Cases by Economic Sector

This section describes the process we used to categorize the incidents and identify the victim organization's sector, and it describes patterns of the UIT incidents specific to economic sectors.

3.1.1. PHISHING/SOCIAL Threat Vector by Economic Sector. Within our 110-case UIT corpus, the PHISHING/SOCIAL threat vector accounts for 44 cases. These cases break down by economic sector of the United States as shown in Figure 4:
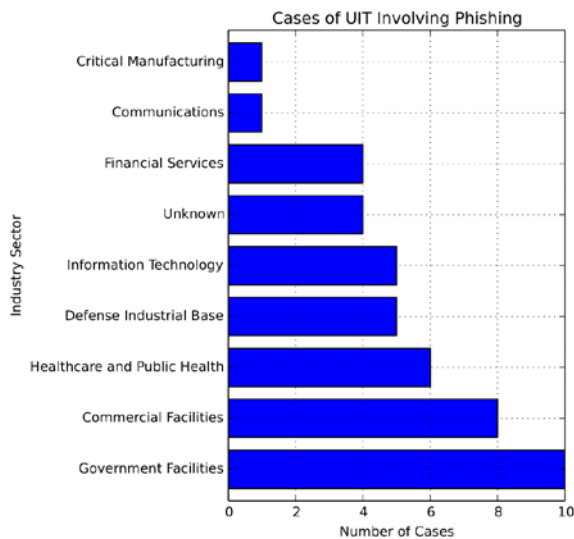

Figure 4. Cases of UIT Involving Phishing by Economic Sector

3.1.2. Malware Threat Subvector by Economic Sector. The Malware subvector offers a compelling area of research because of the high stakes this kind of incident entails. Organizations experiencing this kind of incident will mitigate it differently than they would if only the loss of credentials was at stake, as experienced in the Credentials threat subvector.

Malware incidents constitute 52% of PHISHING/SOCIAL cases in our database, or 23 out of 44. These cases break down by economic sector of the United States as shown in Figure 5:
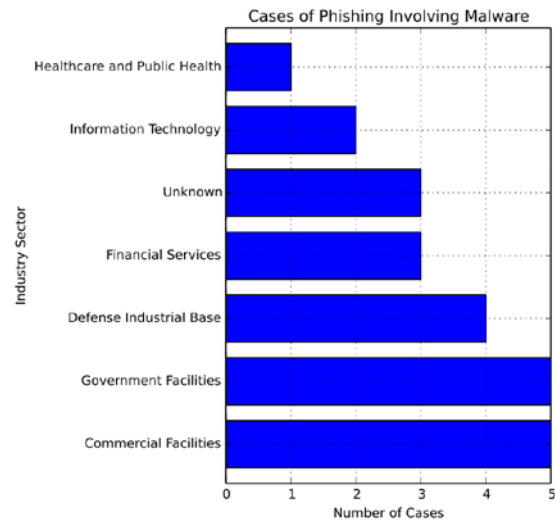
Figure 5. Cases of UIT Involving Malware by Economic Sector

3.2. Patterns of Threat Vector by Economic Sector

Our initial findings suggest that most Malware cases are multi-staged, while most Credentials cases are single-stage. Of the 23 Malware cases, 19 were multi-stage. Of all remaining PHISHING/SOCIAL cases, 15 were single-stage.

4. Initial Findings

Our initial findings suggest that multi-stage and single-stage incidents are almost evenly spread across all sectors of the U.S. economy. The only outliers of this trend are in the commercial and government facility sectors, which showed increased occurrences of single-stage PHISHING/SOCIAL incidents. However, this could be due to higher levels of reporting requirements in these sectors.

Inferring from the results cited, the team concludes that there is a possible trend in single-stage incidents leading toward outside malicious actor ownership of an organization's social media presence, while multi-stage incidents are geared more toward the introduction of malware.

These findings are not definitive, nor should they be taken as a call to action. They are based on a very limited sample set within a community that lacks

proper reporting mechanisms, requirements to report, or places in which to share this information.

Nothing in our research signaled that the mitigation strategies already outlined in previous research needs to be updated [4]. Reviewing the recommendations provided in the UIT foundational study [4], we have confirmed that these mitigation strategies apply to Malware incidents. Our research confirms that the strategies presented in the foundational study may have been useful, in each of the cases we analyzed, for preventing, detecting, or responding to such events.

As noted in the foundational study, a proactive approach to creating a healthy and productive work environment is an essential first step in managing UITs. These steps involve practices that aim to avoid workload pressure and overworked staff, thereby decreasing the propensity for staff to fall prey to phishing emails or other socially engineered attacks.

Furthermore, it is essential to provide adequate education to staff regarding methods used by adversaries to use insiders to unwittingly harm the organization (from the foundational study, Best Practice [BP] 3, BP 18). This education enables staff to recognize social media and phishing threats appropriately. Approximately 40% of the UIT cases we analyzed deal with phishing attacks and 52% of those progressed to Malware cases. Employment of proper mitigation strategies might have prevented these cases.

The protection of the network is also a key general protection strategy. Employees who are educated on the threat vectors may be less likely to click on a potentially malicious link or attachment in an email that may be malware-laced, thus protecting the network. Companies are encouraged to put in place anti-malware and anti-phishing software to proactively address potential threats (BP 19). The use of firewalls is also strongly encouraged. Maintaining data encryption (BP 13, 19) on storage devices as well as password protection for individual users (BP 7, 19) could mitigate phishing attacks leading to hacks that involve data exfiltration.

4.1. Near Misses, Understanding, and the Potential Impact of an Insider Attack

The appropriate value to invest in insider threat mitigation is unique to all organizations. The expected cost of an insider threat attack to an organization can be viewed as the estimated cost of an insider attack multiplied by the probability that an attack will occur. To accurately identify the potential cost of an attack, organizations must prioritize their critical assets and know the exposure of their critical

assets to insiders. The exposure of critical assets to insiders represents the potential for an insider attack. The higher the potential for an attack, the greater the probability of an attack.

The organization can calculate their risk exposure from insider threats by multiplying the probability of an insider attack by the expected cost of an insider attack. An organization's risk exposure can then be used to determine the appropriate amount to spend on insider threat mitigation.

Recent research related to catastrophic events, or "low-probability, high-consequence occurrences" [6], in the process industry [7] and financial services industry [8] may provide strategies for insider threat mitigation. An organization can better understand the impact of an insider attack by understanding its warning signs. These warning signs often come in the form of a near miss: "an event, observation, or situation that possesses the potential for improving a system's safety and/or operability by reducing the risk of upsets, some of which may eventually cause serious damage" [9]. UITs can often be considered near misses in that they expose potential attack vectors and methods for exfiltration.

Consider the following incident. An employee at a Fortune 500 organization is working late into the night to finish her portion of the organization's quarterly forecast. She is working on the spreadsheet containing all of the organization's forecasts for production, shipping, and supplier pricing. Tired after a long day at work, the employee writes an email to her boss, attaches the spreadsheet, and goes to sleep. The next morning, the employee opens her email to find a response to her message: "Thanks for this data; it looks like you put a lot of effort into it."

The message was written by a writer for the largest trade magazine in the industry. The employee immediately reports the mistake to her boss, who notifies the legal office and executives of the mistake. The organization contacts the trade magazine and demands that the data be destroyed, and nothing is published. After the fact, the employee realizes that her email client had auto-completed the writer's address.

This actual incident, along with many of the UIT incidents analyzed for this report, represent the opportunity for the organization to analyze a near miss. By considering the near misses, organizations can better understand the probability of insider harm, the scale, and the consequences for the organization.

Oktem developed an eight-step process for addressing near misses in organizations [9]. The steps, adapted for an insider threat team, appear below:

Identification—Employees in the insider threat detection team must know the definition of a near miss and be trained to identify a near miss.

Disclosure (Reporting)—All near misses must be tracked by the organization. A database designed to record realized UITs and near misses may provide a valuable resource for understanding the potential impact of an incident.

Prioritization—Near-miss UIT cases must be ranked according to their potential impact to the organization.

Distribution—The highest impact near misses should be reported to those who have the potential to create a similar incident.

Identification of Causes (Causal Analysis)—The organization should determine the root cause or causes of the near-miss incidents and determine if the near miss reveals a vulnerability that could be exploited by malicious attackers. In the case of a phishing attack, the cause might be a lack of training and awareness.

Solution Identification—Once the root cause is understood, the organization should develop a solution or solutions to address the vulnerability. In the case of phishing, the solution may be to provide training to increase awareness.

Dissemination—In addition to the unintentional insider in the case, solutions should be communicated to those in the organization who were impacted or have the potential to be impacted by the near miss. In the case of a phishing attack, the organization might provide training to those with access to the organization's critical assets.

Resolution (Tracking)—An organization should track solutions as well as identify related, future near misses.

Oktem's work further describes the important aspects of a "Near-Miss Management System." Implementing such a system in the context of an information security team would enable a better understanding of the probability of a UIT incident as well as a better security posture.

5. Recommendations

The research results to this point are limited by the small sample size of available cases. As our collection of sample cases increases, we can better determine the validity of the initial findings of the research. To advance the current practice and state of the art in computer and network defense, especially safeguards against phishing and other social engineering attacks, organizations should prepare and test their ability to prevent, detect, and respond to the incidents covered in this report by following the best practices in the *Common Sense Guide to Mitigating Insider Threats, 4th Edition* [10], as well as the mitigation strategies the Insider Threat team outlined in the *UIT: Foundational Study* [4].

To help the research community determine the level and means of prevention and mitigation for this kind of threat, the following research needs should be addressed across the cyber community:

- Develop an extensive, confidential, self-reporting UIT database. The database should track the security equivalent of near-miss incidents tracked by many healthcare and manufacturing organizations. A UIT database could be used to track the quality of security at an organization over time and to better assess the potential impact of a malicious insider attack.
- Perform a more detailed analysis of UIT credentials and malware incidents and near incidents to inform the development of more effective mitigation approaches and tools.

5.1. Research Needs

We conclude that many of the research needs identified in the *Unintentional Insider Threats: Social Engineering* report [5] are still valid and would be useful for addressing the UIT malware cases identified here. The earlier report's recommendations are summarized in Sections 5.2 and 5.3.

5.2. Development of an Extensive UIT Database

A major roadblock to advancing our understanding of UIT social engineering exploits, such as PHISHING/SOCIAL incidents, and our ability to counter them is a dearth of data from actual incidents. By conducting public searches, we have collected a small number of cases that contain limited details, but we expect that far more case information could be obtained directly from affected organizations. In addition, a self-reporting mechanism is needed to collect and analyze incidents of social engineering. We recommend that a feasibility analysis be conducted to assess whether organizations could be motivated to self-report incidents, how the data may be collected anonymously and nonpunitively, and how the database can collect sensitive information from organizations across the spectrum of the economy [5].

5.3. Detailed Analysis of UIT Incidents

Further research is needed to examine UIT incidents across a broad spectrum of participants in a comprehensive range of industries representing the full breadth of the economy. This research should focus on what factors are present in UIT incidents, how the affected organizations have handled these incidents, and the motivation of those conducting the PHISHING/SOCIAL or Malware exploits. Our current efforts were hampered by having access only to court transcripts and other third-party accounts of the incidents because organizations do not tend to make this information publicly available, even to research institutions. Only by collecting more detailed data and applying analysis and conceptual modeling approaches, such as those described in our previous report [5], will we, as a community, be able to advance our understanding of UIT social engineering.

6. Required Disclosures

7. References

[1] SolarWinds. SolarWinds Federal Cybersecurity Survey Summary Report. SolarWinds, 2014.

[2] Verizon. 2013 Data Breach Investigations Report. Verizon, 2013.

[3] Symantec. "Internet Security Threat Report 2014." 2013 Trends 19 (April 2014). Symantec Corporation.

[4] CERT Insider Threat Center. Unintentional Insider Threats: A Foundational Study (CMU/SEI-2013-TN-022). Software Engineering Institute, Carnegie Mellon University, 2013. http://resources.sei.cmu.edu/library/asset-view.cfm?AssetID=58744

[5] CERT Insider Threat Center. Unintentional Insider Threats: Social Engineering (CMU/SEI-2013-TN-024). Software Engineering Institute, Carnegie Mellon University, 2014. http://resources.sei.cmu.edu/library/asset-view.cfm?AssetID=77455

[6] Kleindorfer, P.; Oktem, U. G.; Pariyani, A.; & Seider, W. D. "Assessment of Catastrophe Risk and Potential Losses in Industry." Computers & Chemical Engineering 47, 20 (December 2012): 85-96.

[7] Kleindorfer, P. R.; Lowe, R. A.; Rosenthal, I.; Rongwei, F.; & Belke, J. C. *Accident Epidemiology and the RMP Rule: Learning from a Decade of Accident History Data for the U.S. Chemical Industry Final Report for Cooperative Agreement R-83033301*. The Wharton School of the University of Pennsylvania, 2007.

[8] Muermann, A. & Oktem, U. "The Near-Miss Management of Operational Risk." Journal of Risk Finance 4 (Fall 2002): 25-36.

[9] U. G. Oktem. "Near-Miss: A Tool for Integrated Safety, Health, Environmental and Security Management." 37th Annual AIChE Loss Prevention Symposium "Integration of Safety and Environmental Concepts." New Orleans, LA, March, 2003.

[10] CERT Insider Threat Center. Common Sense Guide to Mitigating Insider Threats. 4[th] Edition (CMU/SEI-2012-TR-012). Software Engineering Insitute. Carnegie Mellon University, 2012. http://resources.sei.cmu.edu/library/asset-view.cfm?assetid=34017