Insider Threats: It's the *HUMAN*, Stupid!

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

ABSTRACT

Insider threats refer to threats posed by individuals who intentionally or unintentionally destroy, exfiltrate, or leak sensitive information, or expose their organization to outside attacks. Surveys of organizations in government and industry consistently show that threats posed by insiders rival those posed by hackers, and that insider attacks are even more costly. Emerging U.S. government guidelines and policies for establishing insider threat programs tend to specify only *minimum* standards for insider threat monitoring, analysis, and mitigation programs. Arguably, one of the most serious challenges is to identify and integrate behavioral (sociotechnical) indicators of insider threat risk in addition to cyber/technical indicators. That is, in focusing on data that are most readily obtained, insider threat programs most often miss the *human* side of the problem. This talk briefly describes research aiming to catalog human as well as technical factors associated with insider threat risk and summarizes several recent studies that seek to inform the development of more comprehensive, proactive approaches to insider threat assessment.

KEYWORDS

Insider Threat, Behavioral Monitoring, Insider Threat Ontology, Sociotechnical Indicators, Information Assurance, Counterintelligence

1 Introduction

Insider threats refer to threats posed by individuals who intentionally or unintentionally destroy, exfiltrate, or leak sensitive information, or expose their organization to outside

attacks (e.g., [1][2][3]). Surveys of organizations in government and industry¹ consistently show that threats posed by insiders rival those posed by hackers, and that insider attacks are even more costly (e.g., [4]). Emerging U.S. government guidelines and policies for establishing insider threat programs tend to specify only *minimum* standards for insider threat monitoring, analysis, and mitigation programs [5] [6]. Arguably, one of the most serious challenges is to identify and integrate behavioral (sociotechnical) indicators of insider threat risk in addition to cyber/technical indicators. That is, in focusing on data that are most readily obtained, insider threat programs most often miss the *human* side of the problem. This talk briefly describes research aiming to catalog human as well as technical factors associated with insider threat risk and summarizes several recent studies that seek to inform the development of more comprehensive, proactive approaches to insider threat assessment.

2 Background

Figure 1 (after [7]) provides an insider-threat centric representation of information security *failures*, which includes a broad array of threats ranging from intentional attacks to accidents and acts of nature. While this figure provides for outside cyberattacks, a more complete accounting of those threats would double the size and detail of the figure; instead, here the focus is on insider threats, providing only minimal coverage of other types of failures (outside cyberattacks, design/engineering failures, or acts of nature) that serve to position insider threats within the overall context of information security failures. This “big picture” concept for information security failures was influenced by discussions with Carnegie Melon’s CERT team (while I collaborated with them on unintentional insider threat research [8], [9]) and with collaborators on “cyber friendly fire” research performed by the DOE’s Pacific Northwest National Laboratory ([10], [7])—nevertheless the concepts and relationships depicted in this figure represent my views alone, and any potential misrepresentations are solely my responsibility.

While cyber friendly fire is not the focus of this paper, for completeness its definition is as follows [7]: Cyber friendly fire is intentional offensive or defensive cyber/electronic actions

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the Owner/Author.

NCS '19, April 8–10, 2019, Richland, WA, USA
© 2019 Copyright is held by the owner/author(s).
ACM ISBN 978-1-4503-6614-4/19/04.
<https://doi.org/10.1145/3332448.3332458>

¹ <https://resources.sei.cmu.edu/library/asset-view.cfm?assetID=484019>

Figure 1. “Big Picture” of Information Security Failures – Focus on Insider Threat

intended to protect cyber systems against enemy forces or to attack enemy cyber systems, but which unintentionally harms the mission effectiveness of friendly or neutral forces. As discussed in [7], cyber friendly fire falls within the general class of unintentional insider threats. As in all cases of unintentional insider threats, actors who commit cyber friendly fire do not do so with malicious intent, but rather through lack of knowledge, or due to errors in judgment or other such contributing factors. These unintentional insider threats are largely influenced by the same set of human factors (and organizational factors) that impact human errors.

There are five main types of insider threats [2]:

- **Insider Sabotage.** An act by an insider to direct specific harm toward an organization or its assets.
- **Insider Data Theft/Exfiltration.** Theft of an organization’s intellectual property by an insider.
- **Insider Fraud.** Modification, addition, deletion, or theft, of an organization’s data for personal gain, leading to an identity crime (e.g., identity theft, credit card fraud).
- **Unintentional Insider Threat.** An act or failure to act by an insider, without malicious intent, that causes harm or substantially increases the probability of future harm to an organization or its assets.
- **Workplace Violence.** Any act or threat of physical violence, harassment, intimidation, or other threatening disruptive behavior that occurs at the work site.

Data theft/exfiltration and sabotage may also be further differentiated to the extent that they include espionage (i.e., the perpetrator is working with or for a foreign adversary). Here we do not address such distinctions, but the interested reader may refer to [1] for more details about threat characteristics and mitigation strategies.

The typical monitoring approaches to insider threats are heavily influenced by defensive measures to thwart *outside* attacks—e.g., monitoring of host/network logs, intrusion detection systems, and data loss prevention products. While these are useful and *necessary* components of an organization’s insider threat program, they are not *sufficient*: Insiders have a key advantage over outside attackers because insiders are trusted individuals with access to information systems and networks that contain private data and intellectual property that is to be protected. In this sense, insiders can be more dangerous than outside threats. Therefore, a more comprehensive monitoring approach is needed that includes behavioral and psychosocial factors as well as the more typical technical factors [11].

2 Problem

Unfortunately, even after more than ten years of research that has emphasized the need for including more behavioral factors in a comprehensive insider threat program (e.g., [12-19]), most organizations have not implemented such programs. In fact, according to a recent survey of 150 information security professionals [20], about five percent of organizations have no

insider threat programs, and nearly a quarter of those surveyed have no processes in place to identify potential insider threats—they only react to and investigate insider crimes after they occur. A large percentage (48%) of respondents reported only limited insider threat programs that monitor activities of higher-risk individuals to detect potential malicious threats. Only sixteen percent of respondents reported that their organizations monitor all employees for potential malicious insider threats, and only about 7% of organizations exemplify best practices through monitoring programs that attempt to identify potential unintentional insider threats as well as malicious threats, before they occur. These most mature insider threat programs use multidisciplinary insider threat mitigation teams that review technical and behavioral indicators of insider threats.

The general lack of mature insider threat programs—and the implementation of approaches that largely reflect defensive tools aimed at outside attacks—reminds us of a proverbial comic strip sometimes called “The drunkard’s search” or the “Streetlight Effect”² that tells of a drunkard searching under a street lamp for keys he has lost because the light there is better than where he thinks he lost them. A policeman comes by to ask why he is searching under the lamp post even though the item was lost a block away, and the drunk replies, “this is where the light is.” The streetlight effect refers to the propensity for people who are looking to solve a problem to take an easier path instead of those that are the most likely to yield results. This is a cognitive bias that manifests in research and investigative methods as well as everyday life. The lesson, or challenge, that we might take from this parable in our insider threat programs is: “Do not ignore critical information!”

Arguably, one of the most serious challenges for insider threat research and operations is to identify and integrate behavioral/sociotechnical indicators of insider threat risk, in addition to cyber/technical indicators. By focusing on data that are most readily obtained, insider threat programs risk missing the *human* side of the problem, which includes those factors that are most effective in achieving a more proactive program.

3 Illuminating the Path Toward a More Comprehensive, Proactive Approach

As noted above, we’ve known for a while that there is a need to incorporate organizational and behavioral indicators into the insider threat analysis process. There are numerous possible reasons why this is not typically done, among which are lack of management resolve, lack of information, technical barriers to developing a comprehensive approach that integrates technical and “human” factors, inadequate funding, and combinations of these. Two other reasons are lack of studies to scientifically validate the proposed associations between behavioral indicators and insider threat risk; and privacy/ethical concerns about

monitoring behavioral or psychological risk factors. I shall return to these topics in Section 6.

Despite the lack of progress in adopting more comprehensive insider threat programs, there has been progress in identifying potentially relevant factors to consider when setting up or updating insider threat monitoring programs. A productive approach to defining relevant factors is the development of a formal knowledge representation, or ontology, to capture the concepts. An ontology is a formal description of concepts within a domain that defines concepts and relationships that may be applied consistently. Ontologies provide computational properties that support inferences from asserted facts, and therefore they are useful in model development. CMU CERT developed an Insider Threat Indicator Ontology [21] that focuses on technical indicators; Greitzer et al. [22] [23] developed a comprehensive insider threat ontology (called, Sociotechnical and Organizational Factors for Insider Threat, SOFIT) that includes behavioral and organizational indicators as well as technical indicators. A comparison of these two ontologies is provided in [22].

A high-level picture of the ontology is shown in Figure 2. The development of the SOFIT ontology has been described in [22], [23], and studies that have been conducted to verify the concepts based on expert judgments are described in [24][25]. SOFIT was initially developed as a taxonomy, and then implemented as an ontology with over 320 constructs (factors), including a:

- Individual (human) factor branch that contains more than 270 technical and behavioral indicators
- Organizational factor branch that includes roughly 50 contributing factors associated with organizational constructs.

Figure 2. High Level Structure of SOFIT Ontology

Individual factors reflect behaviors, attitudes, personal issues, sociocultural or ideological factors, and various biographical (life narrative) factors that may indicate increased risk. This branch of the ontology, shown in Figure 3, reflects a substantial body of work by a diverse set of researchers and practitioners focusing on concerning behaviors, sociocultural factors, and psychological factors underlying insider threats. Organizational factors (shown in Figure 4) reflect practices that may contribute to or trigger an insider threat (malicious or unintentional).

² <http://www.rightattitudes.com/2016/02/26/drun kard-search-streetlight-effect/>

Figure 3. SOFIT Individual Factors Branch

Figure 4. SOFIT Organizational Factors Branch

Organizational factors have been identified in a large body of research from industrial/organizational psychology and human factors research on human errors, including communication issues, work setting/management systems, work planning and control, and security practices. More detailed discussion of individual and organizational factors in the SOFIT ontology may be found in [25].

SOFIT was developed with an aim to provide guidance on the types of indicators that an organization might employ in a comprehensive insider threat program. Not all indicators are necessarily adopted or available, but generally speaking, the more that are used, the more comprehensive the monitoring program. A significant barrier to utilizing this and other similar knowledge bases is a lack of studies that demonstrate the validity of the concepts, i.e., statistically reliable relationships between observed indicators and documented insider threat exploits. Organizations do not typically share detailed data on such crimes, so availability of test data is lacking. Even in cases where monitored data are available to test against documented insider threats, the scope of information available is often inadequate (because of the sparse use of behavioral data). Because of these limitations, we have attempted to provide some insight about the relevance of the various SOFIT indicators by using expert judgments as proxies for operational test data.

3 Expert Judgments of Indicator “Strength”

To test the concepts defined in SOFIT, we conducted several expert knowledge elicitation studies. More detailed description of experimental methods and results are provided in [24] and [25]. Here I provide some main findings and implications.

Experts were recruited from the research and operational communities to complete email-based and online surveys. The experts had at least five years of experience in their respective domains, which included cybersecurity, psychology, and human resources, and related fields. Our surveys gathered expert judgments on:

- Threat ratings of *single* indicators
- Ratings for cases comprising *multiple* indicators
- Temporal relationships
- Comparing “concern” with “likelihood” threat ratings

3.1 Comparing Individual Indicators

One of the first studies simply sought to verify our expectations that different indicators vary in their “severity” or “strength” as indicators of potential insider threat risk. To obtain expert judgments, we initially asked eight experts (pilot study) to rate individual indicators using a five point “Level of Concern” scale (1 = NO CONCERN, 2 = SOMEWHAT CONCERNED, 3 = CONCERNED, 4 = VERY CONCERNED, 5 = EXTREMELY CONCERNED). In later studies (tapping expertise of 13-20 experts with comparable background/experience in cybersecurity/information security) that replicated and expanded

this pilot study, we used a 0-100 point scale reflecting the Level of Concern dimension, and obtained ratings of over 200 individual factors. In each of these studies we found strong support for the claim that indicators have varying effects on the judged “strength” (level of concern) as potential insider threat indicators. Figure 5 shows results obtained in the initial (pilot) study that used the smaller set of 24 indicators that were spread across different indicator classes. Results obtained in the subsequent studies with over 200 individual indicators, were very similar [24][25].

In the next stage of these studies, we asked the expert participants to rank-order different sets of “cases” comprising multiple indicators. We compiled sets of cases by combining two to five indicators into a list of observed indicators for hypothetical cases. In one study a total of 45 cases were constructed, and in another study a total of 57 cases were constructed. The judges were asked to rank order the cases on the basis of their “level of concern” as insider threats.

The results were comparable in these ranking studies [24][25], and this enabled us to compare the ability of different threat models to predict the rankings. For example, a very simple model (Counting Model) assumes that all indicators are equally concerning, so that one merely counts the number of indicators observed to arrive at a threat judgment for a given case. A more sophisticated model (Sum-of-Risk Model) assumes that the indicators differ in strength, and the threat level of a case is obtained by summing the threat values of the individual indicators comprising the case. The Counting model yields a poor fit to the obtained ranking data (accounting for only one-fourth of the variance in predicting expert judgments), while the Sum-of-Risk model accounts for about half of the variance. These results confirm that (a) different indicators contribute differentially to judged insider threat risk (and therefore, a method that simply counts observed indicators is inadequate); and (b) a model that takes judged risk ratings of individual indicators into account can perform adequately in accounting for expert judgments of risk for insider threat cases comprising multiple indicators.

Figure 5. Pilot Study Showing Varied Levels of Concern Across 24 Indicators

In another study, judgments of “level of concern” were compared with judgments of likelihood (experts were asked to provide a relative likelihood ratio comparing pairs of indicators as to the strength of their association with insider threats. Because human generated probability estimates are extremely difficult and unreliable, it was hoped that this alternative methodology would provide another means of obtaining probability values, which would be more useful in developing and applying probabilistic threat models (as opposed to the simple threat models described above). The notion is, if there is a high correlation between the “Level of Concern” ratings and the derived likelihood ratings, then this relationship could be used to derive probability values for the SOFIT indicators that may be incorporated into a Bayesian network model. In fact, the findings of the Likelihood Ratio study did reveal a high correlation between the threat values obtained in our earlier studies with the likelihood ratings obtained in this study ($R\text{-squared} = 0.762$), which then enabled us to develop a more formal Bayesian model that could be used to predict the threat rankings of cases based on the set of observed individual indicators. At the time of writing this paper, the results indicate that the Bayesian model performs slightly better than the Sum-of-Risk model in predicting the case rankings.

3.2 Preliminary Findings on Temporal Effects

We structured a survey to determine if certain kinds of indicators have more persistent effects on threat/risk judgments than other factors. For example, it seemed reasonable to expect that certain indicators would be more transient, such as cyber/technical indicators like authentication errors, or printing large files. If we observe a large number of such authentication or printing events at one point in time, but none in the several subsequent months, these indicators may not have much influence on our threat judgment. Other types of indicators—particularly psychological factors such as personality traits, which are more static—should have a more lasting impact on risk judgments. Results were still being analyzed as of this writing, but preliminary results suggest that there is a slight decrease in impact over time. All types of indicators (cyber/technical as well as behavioral/psychosocial) appear to decay over time in their influence on insider threat judgments, but it appears that certain factors (personal predispositions) decay at a slower rate than other indicators.

4 Applications

There are at least three ways that a knowledge base such as SOFIT may be applied and benefit the research and operational communities. First, the ontology informs the research and operational communities about concepts relevant to establishing or advancing insider threat programs. By interacting with the material (which is implemented in the web ontology language, .owl, and which may be viewed with an application such as Protégé), users may explore the knowledge base, its class structure, and reference source material associated with specified indicators.

A second application of the knowledge base is to provide a foundation for development of tools to assess an organization’s insider threat monitoring program. The National Insider Threat Task Force (NITTF), formed as a result of the U.S. National Insider Threat Policy [5][6], provides a set of broad-based, but minimum standards for U.S. government agencies, especially agencies and departments that use classified material. The SOFIT knowledge base has both a narrower focus (insider threat indicators) and a far more in-depth treatment of potential insider threat indicators. Therefore, it can inform a far more detailed assessment of an organization’s insider threat indicator monitoring approach. Comparing the indicators detectable by the organization’s system with the set of indicators identified in SOFIT can help identify possible improvements.

A third application is the use of the knowledge base to inform qualitative and quantitative insider threat assessment. In Section 3, a brief review of ongoing research was described that is aimed toward developing and refining predictive threat models.

5 Limitations and Future Research

The studies reported here were conducted using expert judgments as proxies for actual insider threat outcomes. The results are encouraging as they appear to match expert judgements, but a more rigorous validation is needed using real data to validate the proposed approach and threat models. Further research is needed to solidify the foundation provided in the SOFIT ontology class structure and to refine and validate predictive models. This will then enable development of operationally oriented tools for assessing insider threat/risk for individuals.

The role of organizational factors has largely been unexplored. While the ontology has provided a foundation for considering organizational factors as contributors to insider threats, much more research is needed to refine the knowledge base by estimating the contributions of specified organizational factors to increased organizational vulnerability.

This research was informative in specifying predictive models that account for roughly half the variance in predicting expert judgments. Nevertheless, there is still room for improvement. Perhaps model predictions can be improved by obtaining more accurate estimates of indicator threat values, or perhaps deficiencies may be addressed by dealing with inadequacies in the models. The models examined so far all assume that the impact of any indicator is independent from any other; more sophisticated models may be considered that take more complex relationships (e.g., interactions) among indicators into account.

Finally, the research discussed here was conducted without attempting to differentiate among different types of insider threats. Experts were not asked to assume one particular type of threat when asked to provide judgments of insider threat “level of concern” for particular indicators, or for cases comprising multiple indicators. Perhaps future research can consider more refined models that are keyed to specific types of insider threats.

6 Where to Go From Here

This talk provided an overview of an approach to insider threat mitigation that emphasizes the need for more comprehensive, proactive insider threat programs that include sociotechnical (behavioral, organizational) factors in addition to the more traditional technical factors. As has been noted in this research (e.g., [24]), monitoring sociotechnical factors may facilitate getting “left of boom,” where boom refers to the kinetic action (in this case, the insider attack). That is, if a program only tracks cyber/technical indicators, then it will most likely be caught in a “reactive” mode that will only allow cleaning up after the exploit, applying forensic tools to investigate what happened. Proactive monitoring of behavioral factors provides an opportunity to address motivations and contributing factors that are reflected in concerning behaviors, which may be observed weeks or months prior to the insider exploit/event.

6.1 Need for Validation

These comments reflect the *opinions* of several researchers in the field of cybersecurity and insider threat. The relationship between behavioral factors and insider threat risk remains largely untested (except for limited studies involving simulated data and/or expert judgments as proxies for real data and outcomes): When subjected to more rigorous evaluation studies, some proposed psychosocial factors may turn out to have minimal predictive power while others may be found to be highly informative about an individual’s predilection or potential to become an insider threat. The point is that we cannot resolve the question without validation studies, and we cannot perform rigorous validation studies unless we are able to acquire operational data that represent a broad range of psychosocial, behavioral, and organizational indicators as well as technical (cyber/network) indicators. This need is paramount to the advancement of this field of research and to the establishment of more effective insider mitigation strategies.

6.2 Overcoming Barriers to Adoption

There are many reasons why organizations have not adopted the more mature, comprehensive insider threat approaches that include monitoring and analysis of both behavioral and technical indicators. Because of the publication of industry-wide surveys, scientific works, and policies and standards set for government agencies, lack of knowledge is less a factor than lack of resolve. The SOFIT knowledge base provides a “roadmap” for identifying possible behavioral and technical factors to include in a comprehensive monitoring program. Nevertheless, there is a reluctance to adopt such approaches because of concern over privacy and ethical issues.

Any discussion of potential methodologies, tools, or organizational policies aimed at mitigating insider abuses and potential insider threats should include consideration of privacy and ethical issues. Many organizations are reluctant to consider behavioral monitoring because of feared repercussions in the form of employee morale or public relations impact. The treatment of privacy and ethical issues in the insider threat research literature

has been relatively sparse. The issues, which are described in more detail in [17] and [26], may be summarized as follows:

Employers have the right to monitor employee cyber activities and to internally use personal employee data—for example, it is common for employers to monitor employee email and Internet usage on office computers: Alder and Ambrose [27] reported that at least 40 million U.S. workers are subject to electronic monitoring. From the employer’s perspective, monitoring promotes productivity and affords better control over counterproductive employees; it is justified because employers “own” or pay for employee time and resources such as computer equipment and network connections.

However, employee monitoring (user activity or behavioral monitoring) can have severe ramifications in terms of employee job satisfaction and public relations. There is a fine line between what the organization “needs to know” and what is firmly in the realm of the employee’s expectation of privacy. Since a small percentage of employees engage in malicious insider threat activities, the vast majority comprises honest, hard-working staff who may be highly offended to learn they are being monitored.

To the employer, the cost and damage of even one instance of sabotage or espionage may exceed the cost of establishing a comprehensive insider threat program to proactively mitigate the risk. Further, risk to the company is accompanied by risk to its employees: Since insider threats can jeopardize the survival of the company, catastrophic insider crimes can destroy jobs along with the company. In this view, employees should be more accepting of monitoring as long as they can trust management to conduct it responsibly and ethically.

Privacy rights advocates argue that employees may be harmed when their companies compile such information. Critics also note that monitoring can increase employee stress, reduce commitment, and lower productivity [28]. Monitoring perceived as invasive with an implied lack of trust may contribute to employee job dissatisfaction, and management intervention on suspected employee disgruntlement issues may actually increase an employee’s frustration level [14]. Moreover, false accusations can have severe negative effects on careers of the accused.

All this consternation surrounding the issue might be more effectively addressed if the discussion about insider threat were changed from a largely “punitive” mindset to a “supportive” one. If an insider *risk* (vice “threat”) program is considered in light of how the organization can better identify ways to help troubled employees and improve management practices or policies that negatively impact morale and productivity, then the entire enterprise might be seen as providing more enlightened organizational management practices and policies. If behavioral monitoring can identify individuals who appear to be “at risk” (i.e., on the critical pathway [19] leading to insider crimes), then a more enlightened and supportive management team can provide support aimed to facilitate the troubled employee’s exit from the critical pathway—through employee assistance, management oversight/mentoring, counseling, job transfer to find a better fit to

the individual's abilities and interests, examination of compensation/benefits, etc.—and prevent the “boom.”

In this enlightened approach, employers should communicate to employees the rationale and elements of their insider threat programs and the approach for balancing security and privacy [26] to benefit all. Disclosure of monitoring policies also removes the expectation of privacy (from a legal perspective). If the process is disclosed, explained, and managed equitably across employees, it may not be considered unfair by employees, and the mutual trust relationship required for a healthy organization may remain intact [17]. In this more enlightened view, adopting an open policy of continuous assessment of employee performance and “wellbeing”—as well as identifying and correcting organizational policies and practices that negatively impact morale—can produce benefits to both sides: It need not be considered a zero-sum game.

ACKNOWLEDGMENTS

This research was partially supported under IARPA contract 2016-16031400006. The content is solely the responsibility of the author and does not necessarily represent the official views of the U.S. Government. Special thanks are due to many research collaborators, including SOFIT team members J Purl, YM Leong, PJ Sticha, DE Becker, and J Lee; former colleagues CF Noonan, TE Carroll, DA Frincke; and many others who influenced my research on insider threat over the last two decades.

REFERENCES

- [1] Dawn N. Cappelli, Andrew P. Moore, & Randy F. Trzeciak. (2012). *The CERT guide to insider threats: How to prevent, detect, and respond to information technology crimes (theft, sabotage, fraud)*. Addison-Wesley.
- [2] Frank L. Greitzer, Justin Purl, Yung-Mei Leong, & Paul J. Sticha. (2019). Positioning your organization to respond to insider threats. *IEEE Engineering Management Review*, 47(2), 75-83.
- [3] David Luckey, David Stebbins, Rebeca Orrie, Erin Rebhan, Sunny D. Bhatt, & Sina Beaghley. (2019). *Assessing continuous evaluation approaches for insider threats*. Santa Monica: RAND Corporation. ISBN: 978-1-9774-0194-6.
- [4] Ponemon Institute. (2018). *2018 Cost of Insider Threats: Global*. Ponemon Institute Research Report, April, 2018.
- [5] The White House. (2012). *Presidential Memorandum—National Insider Threat Policy and Minimum Standards for Executive Branch Insider Threat Programs*. November 21, 2012.
- [6] National Insider Threat Task Force (NITTF). (2017). *Insider Threat Guide: A Compendium of Best Practices to Accompany the National Insider Threat Minimum Standards*. Available online: <https://www.dni.gov/files/NCSC/documents/nitff/NITTF-Insider-Threat-Guide-2017.pdf>
- [7] Thomas E. Carroll, Frank L. Greitzer, and Adam D. Roberts. (2014). Security informatics research challenges for mitigating cyber friendly fire. *Security Informatics*, 3(13) (25 September 2014). <http://www.security-informatics.com/content/3/1/13>
- [8] Frank L. Greitzer, Jeremy Strozer, Sholom Cohen, Andrew Moore, David Mundie, and Jennifer Cowley. (2014). Analysis of unintentional insider threats deriving from social engineering exploits. *IEEE Security and Privacy Workshop on Research for Insider Threat (WRIT)*, San Jose, CA, May 17-18, 2014. <http://www.ieee-security.org/TC/SPW2014/papers/5103a236.PDF>
- [9] Frank L. Greitzer, Jeremy Strozer, Sholom Cohen, John Bergey, Jennifer Cowley, Andrew Moore, and David Mundie. (2014). Unintentional insider threat: contributing factors, observables, and mitigation strategies. *47th Hawaii International Conference on Systems Sciences (HICSS-47)*, Big Island, Hawaii. http://ieeexplore.ieee.org/xpl/login.jsp?tp=&arnumber=6758854&url=http%3A%2F%2Fieeexplore.ieee.org%2Fxppls%2Fabs_all.jsp%3Farnumber%3D6758854
- [10] Frank L. Greitzer, Thomas E. Carroll, & Adam D. Roberts. (2014). Cyber friendly fire: Research challenges for security informatics.” *IEEE International Conference on Intelligence and Security Informatics*. Seattle, WA, June 4-7, 2013. https://cybersecurity.pnnl.gov/documents/Cyber_Friendly_Fire_Research_Challenges_for_Security_Informatics.pdf
- [11] Frank L. Greitzer & Deborah A Frincke. (2010). “Combining Traditional Cyber Security Audit Data with Psychosocial Data: Towards Predictive Modeling for Insider Threat Mitigation.” In *Insider Threats in Cyber Security*, ed. CW Probst, J Hunter, D Gollmann & M Bishop, pp. 85-113. Springer, New York. http://dx.doi.org/10.1007/978-1-4419-7133-3_5
- [12] Eric D. Shaw, Jerrold M. Post, & Keven G. Ruby. (1999). Inside the mind of the insider. *Security Management*, 43 (12), 34-42.
- [13] Michael Gelles (2005). Exploring the mind of the spy. In Online Employees' Guide to Security Responsibilities: Treason 101. Retrieved from Texas A&M University Research Foundation website: <http://www.dss.mil/searchdir/training/csg/security/Treason/Mind.htm>
- [14] Eric Shaw & Lynn F. Fischer. (2005). *Ten Tales of Betrayal: The Threat to Corporate Infrastructures by Information Technology Insiders. Report 1—Overview and General Observations*. Technical Report 05-04, April 2005. Monterey, CA: Defense Personnel Security Research Center.
- [15] Eric Cole & Sandra Ring. (2006). *Insider Threat: Protecting the Enterprise from Sabotage, Spying, and Theft*. Rockland, MA: Syngress Publishing.
- [16] Denise Bulling, Mario Scalora, Randy Borum, Jill Panuzio, & Andrew Donica. (2008). *Behavioral science guidelines for assessing insider threats*. Publications of the University of Nebraska Public Policy Center. Paper 37.
- [17] Frank L. Greitzer, Deborah A. Frincke, & Mariah M. Zabriskie. (2011). Social/ethical issues in predictive insider threat monitoring. In: MJ Dark (Ed.), *Information Assurance and Security Ethics in Complex Systems: Interdisciplinary Perspectives*. Hershey, Pennsylvania: IGI Global. Chapter 7, pp.132-161.
- [18] Frank L. Greitzer, Lars. J. Kangas, Christine F. Noonan, Christopher R. Brown, & Thomas Ferryman. (2013). Psychosocial modeling of insider threat risk based on behavioral and word use analysis. *e-Service Journal*, 9(1), 106-138.
- [19] Eric Shaw & Laura Sellers. (2015). Application of the critical-path method to evaluate insider risks. *Studies in Intelligence*, 59(2), 41-48.
- [20] Jim Henderson & Nick Cavalancia. (2019). *2019 Insider Threat Program Maturity Model Report*. Techvangelism & Insider Threat Defense. Retrieved from: <https://cdn2.hubspot.net/hubfs/5260286/PDFs/Whitepapers/insider-threat-maturity-report-2019.pdf>
- [21] Daniel L. Costa, Mathew L. Collins, Samuel J. Perl, Michael J. Albrethsen, George J. Silowash, & Derrick L. Spooner. (2014). An Ontology for Insider Threat Indicators. In K. B. Laskey, I. Emmons and P C.G. Costa (Eds.), *Proceedings of the Ninth Conference on Semantic Technologies for Intelligence, Defense, and Security (STIDS 2014)*, 2014, 48-53.
- [22] Frank L. Greitzer, Muhammad Imran, Justin Purl, Elise T Axelrad, Yung-Mei Leong, D.E. (Sunny) Becker, Kathryn B Laskey, and Paul J. Sticha. (2016). Developing an ontology for individual and organizational sociotechnical indicators of insider threat risk. *The Eleventh International Conference on Semantic Technology for Intelligence, Defense, and Security (STIDS 2016)*, Fairfax, VA, November 15-16, 2016. Paper available at: http://ceur-ws.org/Vol-1788/STIDS_2016_T03_Greitzer_etal.pdf
- [23] Frank L. Greitzer, James Lee, Justin Purl, & Abbas K Zaidi. (2019). Design and implementation of a comprehensive insider threat ontology. 17th Annual Conference on Systems Engineering Research (CSER). Washington, DC, April 2019. <https://doi.org/10.1016/j.procs.2019.05.090>. Published in *Procedia Computer Science*, 2019, 153, 36-369. <https://www.sciencedirect.com/science/article/pii/S1877050919307495>
- [24] Frank L. Greitzer, Justin Purl, Yung-Mei Leong, & D.E. (Sunny) Becker. (2018). SOFIT: Sociotechnical and Organizational Factors for Insider Threat. IEEE S&P Symposium, Workshop on Research for Insider Threat (WRIT). San Francisco, CA, May 24, 2018. <http://conferences.computer.org/sp/2018/Resources/spw/2018/SOFITSociotechnicalandOrganizationalFact.pdf>
- [25] Frank L. Greitzer, Justin Purl, D.E. (Sunny) Becker, Paul J. Sticha, & Yung-Mei Leong. (2019). Modeling expert judgments of insider threat using ontology structure: Effects of individual indicator threat value and class membership. 52nd Hawaii International Conference on Systems Sciences (HICSS-52), Maui, Hawaii, January 2019. <https://scholarspace.manoa.hawaii.edu/bitstream/10125/59756/0316.pdf>
- [26] Angela I. T. Kiser, Timothy Porter, & David Vequist. (2010). “Employee monitoring and ethics: Can they co-exist?” *International Journal of Digital Literacy and Digital Competence*, 1(3), 30-45.
- [27] G. Stoney Alder, & Maureen L. Ambrose. (2005). An examination of the effect of computerized performance monitoring feedback on monitoring fairness, performance, and satisfaction. *Organizational Behavior and Human Decision Processes*, 97, 161-177. DOI:10.1016/j.obhdp.2005.03.003
- [28] William S. Brown. (1996). Technology, workplace privacy, and personhood. *Journal of Business Ethics*, 15, 1237-1248.