Insider Threat Defined: Discovering the Prototypical Case

David A. Mundie, Samuel J. Perl*, and Carly L. Huth, J. D.
Software Engineering Institute, CERT Division
Pittsburgh, Pennsylvania, USA
{dmundie, sjperl, clhuth}@cert.org

Abstract

In a continued effort to better define the field of insider threat research, this study presents a survey of 30 cybersecurity experts' opinions on the attributes of a prototypical insider and insider threat case. The survey is based on the attributes in the Entity-Relationship Model developed in a previous study of 42 different definitions of insider and insider threat. To develop clearer consensus and uniformity in the field, we discuss the attributes, which, in this small exploratory study, experts saw as typical or atypical components of an insider threat case.

Keywords: insider threat, taxonomy, ontology, attributes

1 Introduction

As we noted in a previous study, for most insider threat researchers, the concepts of insider and insider threat were not Aristotelian categories, meaning there is no mutually agreed-upon set of attributes that are both necessary and sufficient for identifying an individual as an insider [1]. Instead we suspected that insider and insider threat may be "natural" categories as explicated in Roschian prototype theory [2]. This would imply that researchers have a mental construct of a prototypical insider and insider threat case, using specific attributes for categorizing an individual or incident as a more or less prototypical case.

Prototype theory has a long history in philosophy and linguistics. Wittgenstein's interest in the notion of a "family resemblance," where family members resemble each other even though there is no set of characteristics that they all share in common, is often taken as the starting point. In psychology, Eleanor Rosch's studies [2] remain the most fully developed expression of the theory. She showed that "natural" categories have members that are better and worse examples of the category, so that a robin for example is a better example of a bird than an albatross or a kiwi. She demonstrated that prototypical examples yield faster response times in queries about class membership: "Is a robin a bird?" elicits faster responses than "Is an albatross a bird?" She further showed that when asked to name a few examples of a category, prototypical examples were produced more frequently than less prototypical examples. George Lakoff [3] and Donald Langacker [4] have convincingly applied prototype theory to linguistics.

What all this means is that in prototype theory, "instances of a natural concept are defined by their resemblance to a prototype that is a best or most typical example of the concept, sharing the maximum number of features or attributes with other instances and a minimum number with instances of other concepts." [5]

Our motivation for this work is to improve our ability to analyze existing and new cases of insider threat. Since 2001, the CERT® Insider Threat Center, part of Carnegie Mellon University's Software Engineering Institute, has researched the insider threat problem, collected a large database of insider threat cases, and derived a substantial body of assessment tools, best practices, and mitigation techniques

Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications, volume: 5, number: 2, pp. 7-23
*Corresponding author: 4500 Forbes Avenue Pittsburgh, PA, 15213, Tel: +1-412-268-4112

to help organizations confront the threat [1]. As our collection of observed insider threat cases continues to grow we are faced with multiple classification challenges. When presented with a new case, our first challenge is to make a determination if the case is indeed an insider threat case according to our definition as stated in our previous work [1]:

> Current or former employee, contractor, or other business partner who has or had authorized access to an organization's network, system, or data and intentionally exceeded or misused that access in a manner that negatively affected the confidentiality, integrity, or availability of the organization's information or information systems.

If a case fits our definition, it is placed into our queue for case coding and analysis. Our coding process involves reviewing the case details and placing relevant and important case elements into a structured format allowing for analysis against other case samples. This process occasionally results in differences of opinion between researchers as to what the important elements of a specific case are and how strongly the elements must exist in the case before a case is labeled an insider threat case according to our definition.

Our research on the prototypical elements of an insider threat case is an attempt to improve the collective understanding of the most useful elements in a case and to open a broader dialog with other researchers of insider threat to identify the elements and the "correct operational measures" [6] for each studied case. In our quest to understand the concepts of insider and insider threat, we set out to test for the features which may be shared among experts, to determine if there were in fact attributes in a case that were necessary for the case to then be considered a better, more prototypical example of an insider or an insider threat. This survey is an exploratory attempt to collect data towards creating a formal shared Insider Threat investigator test for construct validity [6].

2 Methodology

To test this hypothesis, we developed a survey asking experts to determine the more prototypical example of insider threats, based on 13 specific attributes. We derived the attributes from those specified in our Entity-Relationship Model (Figure 1), which itself was developed from the examination of 42 different definitions of insider threat. The Entity-Relationship Model describes the defined relationships among the insider, the organization and the assets [1].

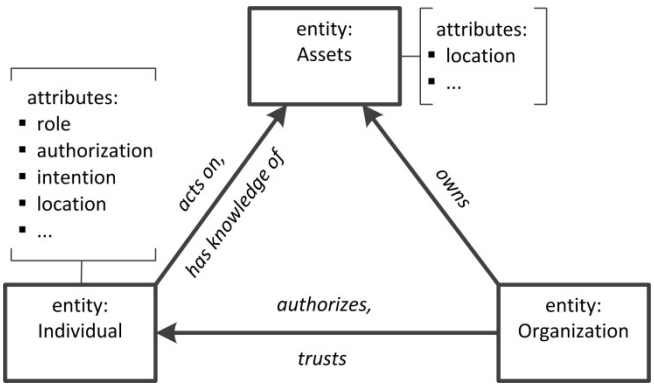

Figure 1: Entity relationship model

We tested 13 proposed attributes:

- asset importance: criticality of the asset
- action: action taken by insider
- status: status of the insider within the organization
- intention: intent of the insider
- authorization: authorization given by the organization to the insider
- organization: type of organization
- relationship: organization's relationship to the assets
- policy violation: violation of the organization's policies
- prosecutability: prosecutability of the attack
- harm: harm or impact of the attack
- individual: individual's relationship to the organization
- location: location of the insider at the time of the attack
- materiality: "'significant,' in other words 'important enough to merit attention.'" [7]

We chose to include both asset importance and materiality as separate attributes. The outcome of an asset importance test is often a range of values such as high, medium, low. Materiality in the Legal Profession remains a debated term for which multiple meanings have been assigned [7]. Materiality in the accounting profession frequently refers to the threshold for determining potential financial significance for transactions and events [8]. The outcome of a materiality test is either material or not-material.

In the survey, we set each attribute in contrasting pairs, asking the respondents which choice was a more prototypical example of an insider and an insider threat case. Both choices were the same except for the value of the chosen attribute. The following is a sample question about asset importance:

Which of the two choices is a better (more prototypical) example of the insider threat?

(a) *An employee emails a competitor a highly sensitive design for a new product.*

(b) *An employee emails a competitor a design for an old product, which the company is phasing out.*

(c) *No difference*

Approximately 30 cybersecurity experts within our organization responded to the anonymous survey, which, in addition to the contrasting pairs, collected a proposed example of insider threat in free text. During the time of the survey, all participants were individuals within our organization and working within the cyber security field. The years of experience working on cyber security ranged from 3 years to over 25 years. We did not provide training to our group of participants and intended to collect their unguided opinions. Our survey was designed to focus upon insider threat cases that included the use or abuse of Information Technology. The full survey given to our participants is included in Appendix A.

3 Results

Four attributes appeared to stand apart from the rest: asset importance, intention, action, and policy violation. The majority of respondents believed that the value of these four attributes made a case more or less prototypical.

The majority of respondents believed that a critical asset being stolen was a more prototypical insider threat case than a trivial asset being stolen, as illustrated in Figure 2.

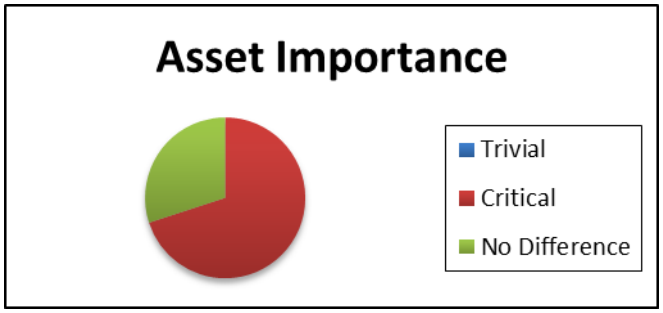

Figure 2: Asset Importance responses

In addition, the majority of respondents believed that a case in which the insider was malicious was more prototypical than a case in which with an unintentional insider, as illustrated in Figure 3.

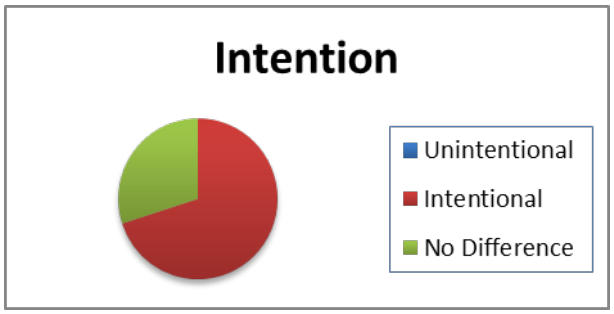

Figure 3: Intention responses

As illustrated in Figure 4, the majority of respondents also thought that a case in which the insider takes an affirmative action (rather than failing to respond to a known violation) was a more prototypical insider threat case.

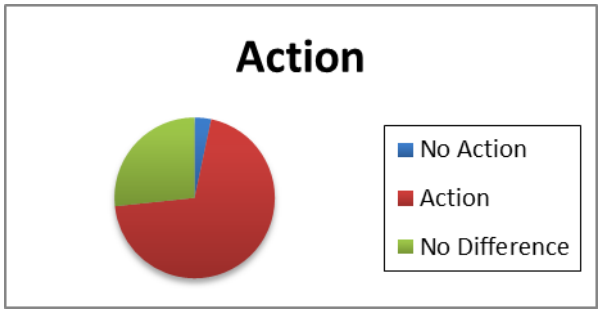

Figure 4: Action responses

Finally, Figure 5 illustrates that most respondents believed that an incident that violated organizational policy was more prototypical than a case in which the action in question did not violate organizational policy.

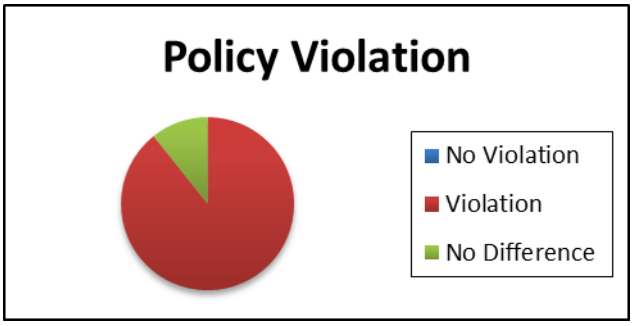

Figure 5: Policy violation responses

Most respondents felt that the rest of the attributes had little effect on the prototypicality of cases, based on the values presented. As indicated in Figure 6, 25 or more of the respondents felt that the values of the attributes did not affect a prototypical case; specifically it did not matter if

- the organization was public or private (organization) the location of the attack was in the workplace or remote (location)
- the insider compromised a relatively small number of records or a larger number (materiality)
- the employee had been at the organization a short or long while (status)

More than 19 respondents believed that the values of the following attributes did not affect a prototypical case; specifically it did not matter if

- the insider was a contractor or former employee (individual)
- the insider had a great deal of authorization on the organization's network or very little (authorization)
- the organization owned the assets or was holding the assets for a third party (relationship)

More than 14 respondents believed that the values of the following attributes did not affect a prototypical case; specifically it did not matter if

- the organization suffered a great deal of harm or a small amount of harm (harm)
- the insider's attack amounted to a prosecutable crime (prosecutability)

As noted in Section 2, the survey also had a free text question, in which experts were asked for an example of insider threat. While we had less of a response to this question, free-text analysis found several commonalities in the responses. Specifically, many responses mentioned the words "employee" (9 responses), "information" (9 responses), "employer" (4 responses), and "exfiltrates" (7 responses). While we have less confidence in this response, these words strongly correlate with the entities and relationships in our models, specifically "individual" (employee), "asset" (information), "organization" (employer), and "action" (exfiltrates).

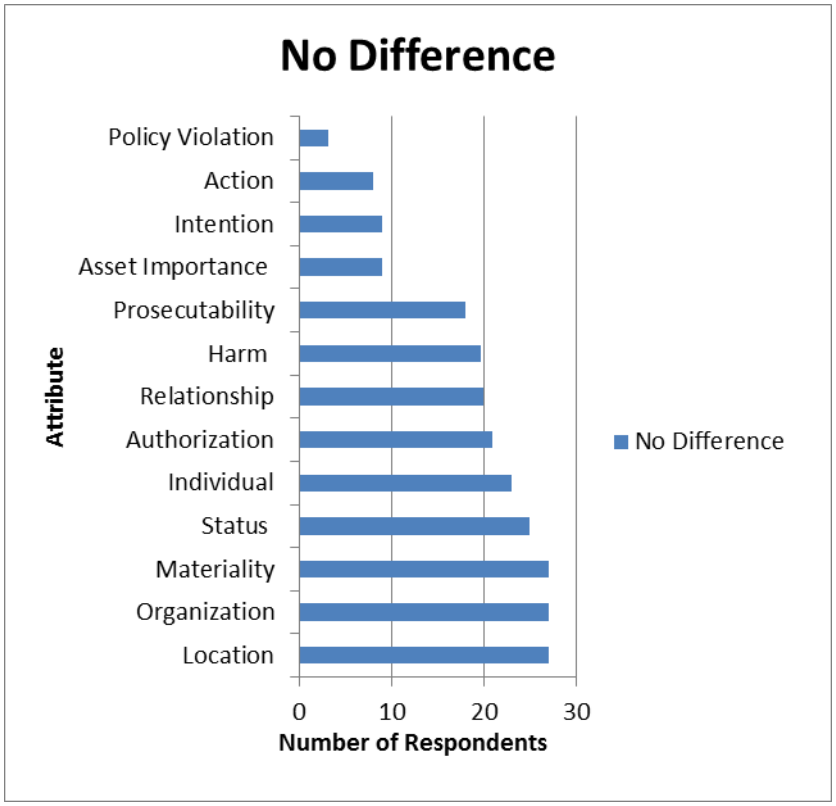

Figure 6: Number of "No Difference" responses per attribute

4 Discussion

This study, while preliminary, yielded several surprises to the researchers. First, we had hypothesized that the value of most of the attributes would more strongly affect the prototypicality of the case. However, for 9 out of the 13 attributes, there was "no difference" between the values. This could be due to a variety of factors, including that the examples presented were not sufficient to polarize the concept, or that the respondent group was in an academic rather than operational setting.

For those attributes that were found to more strongly affect prototypicality, organizational policy was the most surprising attribute, considering how little it was discussed in the insider threat definitions [1]. Intention is an attribute that could have been expected to be of particular importance to our respondents, given the amount of research into motivations and indicators of insider threat [9]. However, as non-malicious insiders become an increasing area of concern, the importance of this attribute may change. The same could perhaps be said of the action attribute. Finally, asset importance has long been considered a key attribute of insider threat; there are best practices about the prioritization of assets for protection [9].

5 Comparison with Existing Work

As we have outlined previously, many different taxonomies of insider threat have been published. While our work is exploratory, we found it interesting to view others' frameworks in light of our preliminary results. For example, several taxonomies have discussed access as a primary attribute [10], [11], [12]. However, our respondents noted there was no difference in the prototypicality of the case of a janitor

with limited access and a system administrator (e.g., authorization attribute). Knowledge is another attribute discussed in taxonomies [10], [11], [12]. However, our respondents noted that a case was not more or less prototypical if the insider had been in the organization for six months or ten years (e.g., status attribute). Interestingly, trust was noted in several taxonomies but not studied here, presenting a possibility for future work [10], [11].

With respect to the free-text question of a proposed insider threat example, Predd et al.'s holistic, four-dimensional approach to understanding insider threat by focusing on the organization, individual, system, and environment likely comes closest to addressing the responses [13]. However, it is notable that the respondents believed that a criminal incident was not any more or less prototypical than a non-criminal incident. Such responses may lead to questions about the role of the environment in a prototypical case.

6 Conclusion and Future Work

In conclusion, while this study is preliminary, it is interesting to note that the values of four attributes appear to affect whether a case is considered prototypical or not and thus may argue for the inclusion of these attributes within a definition of insider threat.

Limitations of the study include the sample size and the fact that all the respondents were individuals within our organization, though they are professionals in the cybersecurity field. In addition, two questions had 29 and 28 responses, respectively, instead of 30; however, this was accounted for in the analysis. In addition, the survey included a comment field (whose entries are not included here) for all but two questions, where it was erroneously excluded. Hypothetically, this could have altered the choices of the respondents, who could have been more likely to choose a different response if allowed to supply a comment. While the questions were workshopped with experts in survey methodology after a pilot test, other confounding factors could have influenced the answers. Due to these limitations, this study can best be characterized as an exploratory proof of concept.

Future work may include applying a similar survey to a larger and more varied population and using this work to improve ontology proposed in our previous work [1]. We began this work with the hypothesis that each of these elements could be classified as a continuum, with a prototypical insider and case of insider threat falling at a particular point along the continuum. As our contrasting pairs could only delineate two of the many points on that continuum, future work may also include fleshing out the proposed spectrum for each attribute, particularly for the four attributes our respondents identified as meaningful to the prototypicality of a case. Future work may include conducting a new survey that incorporates multiple attributes into its questions and solicits additional attributes that would define an insider threat from survey participants. Future surveys could also be conducted on populations outside of the information security field to determine if those participants had different beliefs on the attributes of a prototypical insider threat case. If larger and more diverse populations are surveyed, statistical analysis techniques could be employed. We decided to initially limit the size of our survey to see if we could find any evidence of differences in opinion as to the elements in a prototypical insider threat. Now that we have identified that differences do exist, we can employ surveys of greater size and with better external validity. As part of our original survey, we included a space for participants to provide a single prototypical example but did not analyze them for additional attributes. Future analysis work may also include looking for additional attributes in unstructured survey responses or alternatively, in the source material for documented insider threat cases [14].

7 Acknowledgments

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.

DM-0000124

References

[1] D. A. Mundie, S. Perl, and C. L. Huth, "Toward an ontology for insider threat research: Varieties of insider threat definitions," in *Proc. of the 3rd Workshop on Socio-Technical Aspects in Security and Trust (STAST'13), New Orleans, Los Angeles, USA*. IEEE, June 2013, pp. 26–36.

[2] E. Rosch and B. Lloyd, *Cognition and Categorization*. Hillsdale, N.J. New York: L. Erlbaum Associates, distributed by Halsted Press, 1978.

[3] G. Lakoff, *Women, Fire, and Dangerous Things: What Categories Reveal About the Mind*. Chicago: University of Chicago Press, 1987.

[4] R. W. Langacker, *Foundations of Cognitive Grammar: Theoretical Prerequisites*. Stanford, California: Stanford University Press, 1987.

[5] A. M. Colman, *A Dictionary of Psychology*. New York: Oxford University Press, 2009.

[6] R. K. Yin, "Applications of case study research, 4th ed," in *Applied Social Research Methods Series*. Sage Publications, 2008, vol. 5.

[7] K. Adams, "Contract drafting: Revisiting materiality," *New York Law Journal*, August 2007.

[8] J. Vorhies, "The new importance of materiality," *Journal of Accountancy*, May 2005.

[9] G. Silowash, D. Cappelli, A. Moore, R. Trzeciak, T. Shimeall, and L. Flynn, "Best practice 6: Common sense guide to mitigating insider threats, 4th edition," Software Engineering Institute, Carnegie Mellon University, Tech. Rep., December 2012.

[10] M. Bishop, S. Engle, D. Frincke, C. Gates, F. Greitzer, S. Peisert, and S. Whalen, "A risk management approach to the 'insider threat'," in *Insider Threats in Cyber Security, Advances in Information Security series*. Berlin: Springer, 2010, vol. 49, pp. 115–137.

[11] J. Hunker and C. Probst, "Insiders and insider threats: An overview of definitions and mitigation techniques," *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications*, vol. 2, no. 1, pp. 4–27, March 2011.

[12] B. J. Wood, "An insider threat model for adversary simulation," SRI International, Cyber Defense Research Center, System Design Laboratory, Tech. Rep., 2000.

[13] J. Predd, S. L. Pfleeger, J. Hunker, and C. Bulford, "Insiders behaving badly," *IEEE Security and Privacy*, vol. 6, no. 4, pp. 66–70, July/August 2008.

[14] D. M. Cappelli, A. P. Moore, and R. F. Trzeciak, *The CERT Guide to Insider Threats: How to Prevent, Detect, and Respond to Information Technology Crimes (Theft, Sabotage, Fraud)*. Addison-Wesley, 2012.

[15] M. Rouse, "Insider threat," http://searchsecurity.techtarget.com/definition/insider-threat, September 2010.

[16] E. E. Schultz, "A framework for understanding and predicting insider attacks," *Computers and Security*, vol. 21, no. 6, pp. 526–531, October 2002.

[17] P. Ning and K. Sun, "How to misuse AODV: A case study of insider attacks, against mobile ad-hoc routing protocols," *Ad Hoc Networks*, vol. 3, no. 6, pp. 795–819, November 2005.

[18] R. P. Brackney and R. H. Anderson, "Understanding the insider threat," in *Proc. of a March 2004 Workshop*. Santa Monica, CA: Rand Corporation, March 2004.

[19] J. Patzakis, "New incident response best practices: Patch and proceed is no longer acceptable incident response," Guidance Software, Pasadena, California, Tech. Rep., September 2003.

[20] S. Matthew, M. Petropoulos, H. Mgo, and S. Upadhyaya, "A data-centric approach to insider attack detection in database systems," in *Proc. of the 13th International Symposium on Recent Advances in Intrusion Detection (RAID'10), Ottawa, Ontario, Canada, LNCS*, vol. 6307. Springer-Verlag, September 2010, pp. 382–401.

[21] D. I. N. 5240.26, "Countering espionage, international terrorism, and the counterintelligence insider threat," May 2012.

[22] R. Chinchani, A. Iyer, H. Ngo, and S. Upadhyaya, "Towards a theory of insider threat assessment," in *Proc. of the International Conference on Dependable Systems and Networks (DSN'05), Yokohama, Japan*. IEEE, June/July 2005, pp. 108–117.

[23] S. Pfleeger and S. Stolfo, "Addressing the insider threat," *IEEE Security and Privacy*, vol. 7, pp. 10–13, November/December 2009.

[24] E. Cole and S. Ring, *Insider Threat: Protecting the Enterprise from Sabotage, Spying, and Theft*. Syngress Publishing, 2006.

[25] S. Brown, "Uncloaking the insider threat," http://security.ittoolbox.com/pub/SB052002.pdf, May 2002.

[26] B. Gabrielson, "Solving the insider threat problem," October 2006, presented at the University of Louisville Cyber Securitys Day.

[27] B. Bowen, S. Hershkop, A. Keromytis, and S. Stolfo, "Baiting inside attackers using decoy document," in *5th International ICST Conference on Security and Privacy in Communication Networks (SecureComm'09), Athens, Greece, Revised Selected Papers, Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, vol. 19, September 2009.

[28] G. Jabbour and D. Menasce, "Stopping the insider threat: The case for implementing integrated defense mechanisms in computing systems," in *Proc. of the International Conference on Security and Privacy (ISP'09), Orando, Florida, USA*, July 2009.

[29] B. Aleman-Meza, P. Burns, M. Eavenson, D. Palaniswami, and A. Sheth, "An ontological approach to the document access problems of insider threat," in *Proc. of the IEEE International Conference on Intelligence and Security Informatics (ISI'05), Atlanta, Georgia, USA*. IEEE, May 2005, pp. 486–491.

[30] C. Blackwell, "A security architecture to protect against the insider threat from damage, fraud, and theft," in *Proc. of the 5th Annual Workshop on Cyber Security and Information Intelligence Research (CSIIRW'09), Oak Ridge, Tennessee, USA*. ACM, April 2009.

[31] B. Bowen, M. Salem, A. Keromytis, and S. Stolofo, *Monitoring Technologies for Mitigating Insider Threats*. Springer, 2009.

[32] D. Lieberman, "Defining the insider threat," http://www.infosecisland.com/blogview/12824-Defining-the-Insider-Threat.html, 2011.

[33] J. Butts, R. Mills, and R. Baldwin, "Developing an insider threat model using functional decomposition," in *Proc. of the 3rd International Workshop on Mathematical Methods, Models, and Architectures for Computer Network Security (MMM-ACNS'05), St. Petersburg, Russia, LNCS*, vol. 3685, September 2005, pp. 412–417.

[34] M. Maybury, P. Chase, B. Cheikes, D. Brackney, S. Matzner, T. Hetheringston, B. Wood, C. Sibley, J. Martin, T. Longstaff, L. Spitzner, J. Haile, J. Copeland, and S. Lewandowski, "Analysis and detection of malicious insiders," in *Proc. of the International Conference on Intelligence Analysis, McLean, Virginia, USA*, May 2005.

[35] "Security taxonomy and glossary: Insider threat," http://www.garlic.com/~lynn/secure.htm.

[36] P. Magazine, "Encyclopedia: Insider threat," http://www.pcmag.com/encyclopedia/term/45031/insider-threat.

[37] "Insider threat law and legal definition," http://definitions.uslegal.com/I/INSIDER-THREAT/.

[38] M. Kandias, A. Mylonas, N. Virvilis, M. Theoharidou, and D. Gritzalis, "An insider threat prediction model," in *Proc. of the 7th International Conference on Trust and Privacy in Digital Business (TrustBus'10), Bilbao, Spain, LNCS*, vol. 6264. Springer-Verlag, August 2010, pp. 26–37.

[39] R. Bejtlich, "Insider threat study," http://taosecurity.blogspot.com/2006/09/insider-threat-study.html, September 2006.

[40] B. Bowen, M. Salem, S. Hershkop, A. Keromytis, and S. Stolfo, "Designing host and network sensors to mitigate the insider threat," *IEEE Security and Privacy Magazine: Special Issue on Insider Threat*, vol. 7, no. 6, pp. 22–29, 2009.

[41] M. Maybury, "How to protect digital assets from malicious insiders," http://www.thei3p.org/research/mitremi.html.

[42] A. Parker, "Meaning of insider trading," http://www.ehow.com/info_7894263_meaning-insider-trading.html.

[43] R. Kissel, Ed., *Glossary of Key Information Security Terms, NIST IR 7298*. National Institute of Standards and Technology, June 2013.

[44] N. Einwechter, "Preventing and detecting insider attacks using ids," http://www.symantec.com/connect/articles/preventing-and-detecting-insider-attacks-using-ids, March 2002.

[45] G. Magklaras, S. Furnell, and P. Brooke, "Towards an insider threat prediction specification language," *Information Management and Computer Security*, vol. 14, no. 4, pp. 361–381, 2006.

[46] G. Doss and G. Tejay, "Developing insider attack detection model: A grounded approach," in *Proc. of the International Conference on Intelligence and Security Informatics (ISI'09), Dallas, Texas, USA*. IEEE, June 2009, pp. 107–112.

Author Biography

David A. Mundie is a member of the CSIRT Development Team within the CERT® Program at the Software Engineering Institute (SEI), a unit of Carnegie Mellon University in Pittsburgh, PA. He has been at CERT since 2000 and has worked in a variety of areas including insider threat, malware analysis, and incident management capability metrics. From 2006 to 2009, he was a member of the Q-CERT project, which established a national information security team for the country of Qatar. David's current research interests include formal ontologies for information security, insider threat patterns, and models of incident information sharing. Prior to joining CERT, he worked at Texas Instruments and Western Digital on compiler development, test engineering, and process improvement. Since the time of drafting, David has retired from the SEI.

Samuel J. Perl is a member of the CSIRT development team within the CERT® Program at the Software Engineering Institute (SEI), a unit of Carnegie Mellon University in Pittsburgh, PA. He has been at CERT since 2011 and has worked in a variety of areas including insider threat, vulnerability assessment, security incident data analysis, and incident management team development. Prior to CERT, Perl gained over 10 years of industry experience working with client organizations to manage their most challenging IT security risk issues. Perl holds a M.S. in Information Security Management from Carnegie Mellon University and a B.S in Information Systems from Carnegie Mellon University.

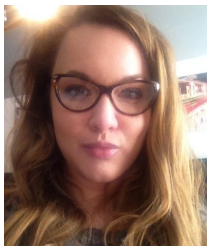

Carly L. Huth is Global Privacy Law and Data Protection Counsel for the Coca-Cola Company, where she focuses on protecting both consumer and employee data and is the co-chair of the Global Privacy Law and Data Protection Practice Group. Prior to joining Coca-Cola, Huth was an insider threat researcher in the SEI's CERT Program. Huth's areas of research included the intersections of privacy and technology. Huth has experience in both international and academic arenas. Prior to joining the SEI, Huth worked with the Intellectual Property Team in the United Nations Conference on Trade and Development and with the University of Pittsburgh's Office of Technology Management. Huth is a licensed patent attorney and a Certified Information Privacy Professional in Information Technology. She holds a Juris Doctor from the University of Pittsburgh with a Certificate in Intellectual Property and Technology Law. Huth also holds a B.S. from Carnegie Mellon University.

Appendix

A Survey Questions

1. Which of the two choices is a better (more prototypical) example of the insider threat?
   (a) An employee emails a competitor a highly sensitive design for a new product.
   (b) An employee emails a competitor a design for an old product, which the company is phasing out.
   (c) No difference.

   Primary attribute tested: asset importance. Criticality of the asset

2. Which of the two choices is a better (more prototypical) example of the insider threat?
   (a) The IT manager disables the anti-virus.
   (b) The IT manager notices that the anti-virus isn't working and doesn't report it.
   (c) No difference.

   Primary attribute tested: action. Action taken by insider

3. Which of the two choices is a better (more prototypical) example of the insider threat?
   (a) An employee of six months places a logic bomb in the organization's system.
   (b) An employee who has been with the organization for ten years places a logic bomb in the organization's system.
   (c) No difference.

   Primary attribute tested: status. Status of the insider within the organization

4. Which of the two choices is a better (more prototypical) example of the insider threat?
   (a) An employee accidentally clicks on a link, which contains malware.
   (b) An employee knowingly downloads malware.

(c) No difference.

Primary attribute tested: intention. Intent of the insider

5. Which of the two choices is a better (more prototypical) example of the insider threat?

   (a) A janitor, who is allowed access to the system only to check his email, maliciously installs a keylogger.

   (b) A system administrator with full access to the system maliciously installs a keylogger.

   (c) No difference.

Primary attribute tested: authorization. Authorization given by the organization to the insider

6. Which of the two choices is a better (more prototypical) example of the insider threat?

   (a) A book keeper commits fraud against a multi-national corporation.

   (b) A book keeper commits fraud against a government agency.

   (c) No difference.

Primary attribute tested: organization. Type of organization

7. Which of the two choices is a better (more prototypical) example of the insider threat?

   (a) A malicious insider deletes files, causing the organization about an hour of lost sales.

   (b) A malicious insider deletes files, the organization, unable to recover, closes.

   (c) No difference.

Primary attribute tested: materiality. "'Significant,' in other words 'important enough to merit attention'"

8. Which of the two choices is a better (more prototypical) example of the insider threat?

   (a) An employee exfiltrates proprietary information which is stored on his employer's system, but owned by a third party.

   (b) An Employee exfiltrates proprietary information owned by his employer.

   (c) No difference.

Primary attribute tested: relationship. The organization's relationship to the assets

9. Which of the two choices is a better (more prototypical) example of the insider threat?

   (a) A contractor uses another employee's credentials which is in violation of company policies.

   (b) A contractor uses another employee's credentials which is not in violation of company policies.

   (c) No difference.

Primary attribute tested: policy violation. Violation of the organization's policies

10. Which of the two choices is a better (more prototypical) example of the insider threat?

(a) An employee posts company information on a website, which is in violation of company policy but is not a crime.

(b) An employee posts company information on a website, which is in violation of company policy and is a crime.

(c) No difference.

Primary attribute tested: prosecutability. Prosecutability of the attack

11. Which of the two choices is a better (more prototypical) example of the insider threat?

(a) An employee accidentally emails 5,000 customer records.

(b) An employee accidentally emails 50,000 customer records.

(c) No difference.

Primary attribute tested: harm. Harm or impact of the attack

12. Which of the two choices is a better (more prototypical) example of the insider threat?

(a) Contractor stealing software design documents.

(b) Former employee stealing software design documents.

(c) No difference.

Primary attribute tested: individual. The individual's relationship to the organization

13. Which of the two choices is a better (more prototypical) example of the insider threat?

(a) While at work a contractor exfiltrates 30 proprietary documents.

(b) While at home, a contractor exfiltrates 30 proprietary documents.

(c) No difference.

Primary attribute tested: location. The location of the insider at the time of the attack

14. Please provide an example you believe to be a prototypical example of insider threat:

B Definitions Studied [1]

- An insider threat is a malicious hacker (also called a cracker or a black hat) who is an employee or officer of a business, institution institution, or agency. [15]
- The term can also apply to an outside person who poses as an employee or officer by obtaining false credentials. The cracker obtains access to the computer systems or networks of the enterprise, and then conducts activities intended to cause harm to the enterprise. [15]
- An insider is someone who is authorized to use computers and networks. [16]
- An insider has access to the keying materials or full control of some nodes. [17]
- An insider has "access, privilege, or knowledge of information systems and services." ([18], p. 10)

- An insider is anyone who operated inside the security perimeter.[19]
- An insider is a database subject who has personal knowledge of information stored in one or more fields marked confidential. [20]
- CI insider threat. A person, known or suspected, who uses their authorized access to DoD facilities, personnel, systems, equipment, information, or infrastructure to damage and disrupt operations, compromise DoD information, or commit espionage on behalf of an FIE. [21]
- Defining the term "insider" in an airtight manner is hard because the boundary between insiders and outsiders is fuzzy. We assume that every legitimate user is an insider. Note that the term "insider" can have both physical and logical connotation. Physical outsiders can be logical insiders and vice versa. For example, an authorized user who may be physically far away from an organization but has wireless or VPN connectivity. Similarly, users may be physically inside an organization but have no authorized access to use the computation infrastructure. Insiders are in a unique position with the privileges entrusted to them and the knowledge about their computational environment, and this already translates directly to a certain amount of capability. Insider abuse can occur within this default capability, but more dangerous scenarios occur when an insider widens his realm of capability. Since insiders have access privileges to use the computational infrastructure, it represents resources at their disposal that can be used against the parent organization, so resources for an insider attack are freely available. Unlike external attackers who use the internet as an umbrella of anonymity and can be sloppy, insiders have a strong incentive to avoid detection. They are a part of an organization and bound by the organization policy, and if caught, an organization has all the necessary information about the insider and the legal resources to prosecute him. External attackers can become insiders too by compromising an internal system and learning about the computers in the neighborhood. However, there is an inherent risk to the attacker that the compromise may be discovered and the corresponding security hole patched. [22]
- Insider: someone with legitimate access to an organization's computers and networks. Notice that we don't define what "legitimate" means and thus don't provide a single bright line distinguishing insiders from outsiders. Both legitimate access and the system's perimeter are a function not only of system-specific characteristics but also of a given organization's policies and values. For instance, an insider might be a contractor, auditor, ex-employee, partner temporary business partner, or more. Thus, the organization itself can best determine who is an insider. [13]
- The insider threat: people with legitimate access who behave in ways that put our data, our systems, our organizations, and even our businesses' viability at risk. Such behavior -malicious might not be malicious; it might be well-intended but still have unwelcome consequences. [23]
- There exist many different definitions of the terms "insider" and "insider threat". One common definition is that "an insider is defined as an individual with privileged access to an IT system"]. (this is the definition used by the Department of Homeland Security (us) research project "Human Factors, Awareness, and Insider Threats", 2007-2009.) [11]
- In 2008, a cross-disciplinary workshop on "countering insider threats"' concluded that "an insider is a person that has been legitimately empowered with the right to access, represent, or decide about one or more assets of the organization's structure.' [11]
- Numerous definitions for the term "insider threats" have been proposed. According to Cole and Ring, "an insider threat is anyone who has special access or knowledge with the intent to cause harm or danger". [24]

- According to Brown, insider threats are conducted by individuals working within a company that are granted any level of access to the organisation's network resources. These insiders know how to corrupt its valuable network resources, and access confidential information. They understand the company's security methods and how to best exploit them. [25]

- As a general DoD definition, the "insider" is anyone who is or has been authorized access to a DoD information system, whether a military member, a civilian employee, employee of another federal agency or the private sector. Some definitions, however, address the broader scope of "system components" or "computer software code" inserted inside a system and intended to carry out a malicious act. Of interest regarding the many broad descriptions of insider is that the definition proposed is often dependent on the perspective of the individual defining the problem. The real question arises, is the perpetrator simply someone exhibiting bad behavior or is this person representing a serious threat to our nation. [26]

- The definition of insider threat should encompass two main threat actor categories and five general categories of activities. The first actor category, the "true insider," is defined as any entity (person, system, or code) authorized by command and control elements to access network, system, or data. The second actor category, the "pseudo-insider," is someone who, by policy, is not authorized the accesses, roles, and/or permissions they currently have but may have gotten them inadvertently or through malicious activities. The activities of both fall into five general categories: 1) exceeds given network, system or data permissions; 2) conducts malicious activity against or across the network, system or data; 3) provided unapproved access to the network, system or data; 4) circumvents security controls or exploits security weaknesses to exceed authorized permitted activity or disguise identify; or 5) non-maliciously or unintentionally damages resources (network, system or data) by destruction, corruption, denial of access, or disclosure. [26]

- Some investigators have cited four categories of the insider problem: traitor, zealot, browser, and well intentioned. The traitor category includes persons who have a malevolent intent to damage, destroy, or sell out their organization. The zealot category involves an insider who believes strongly in the correctness of one position or feels the organization is not on the right side of a certain issue. The browser category consists of persons who are overly curious in nature (often a violation of the need-to-know principle), while the well-intentioned insider commits violations through ignorance. Downloading shareware, disabling virus protection software, using unapproved CDs can all provide the assistance a hacker needs to penetrate a system. The well- intended user can become the unwitting and unknowing associate. [26]

- We define insider threats by differentiating between masqueraders (attackers who impersonate another inside user) and traitors (an inside attacker using their own legitimate credentials). One possible solution for masquerade detection involves anomaly detection. In this approach, users actions are profiled to form a baseline of normal behavior. Subsequent monitoring for abnormal behaviors that exhibit large deviations from this baseline signal a potential insider attack. The common strategy to prevent inside attacks involves policy-based access control techniques to limit the scope of systems and information an insider is authorized to use, and hence, limit the damage the organization may incur when an insider goes awry. Prevention techniques may not always succeed, and thus, monitoring and detection techniques are needed when prevention fails. In this paper, we are focused on different techniques aimed at detecting masqueraders and traitors. [27]

- Some of the available definitions of the term "insider" denote it as any person or system that has a privileged access to the domain or system that is being protected. [28]

- Insider threat refers to the potential malevolent actions by employees within an organization, a specific type of which relates to legitimate access of documents. [29]

- We define an insider as one who has legitimate access to an organization, its systems, information or other resources. The insider threat is a risk that an insider can misuse their access or knowledge to cause harm to the organization. We also mention the insider weakness where an insider performs unsafe actions or fails to apply adequate protection that may expose the organization to accidental damage or malicious attack. We do not count outsiders that appear to be insiders because they have gained internal access by defeating system defenses. [30]

- Users with privileged knowledge about a system. [31]

- An insider can be defined with regard to two primitive actions: violation of a security policy using legitimate access, and violation of an access control policy by obtaining unauthorized access. [32]

- Any attack launched from inside the network by an employee, contractor or visitor that damages or valuable assets by exploiting means (multiple accounts) and opportunity (multiple channels). [32]

- Using this definition, we can see that trusted insider threats is a matter of asset value and threat surface – not just access control. [32]

- To ensure the model adequately addresses the insider threat, it is necessary to clearly define the aspects that are being captured. In this context, an insider is any individual who has been granted any level of trust in an information system. This description does not limit the insider to specific borders such as firewalls, routers, or a local area network. The system itself could be a conglomeration of networks. What is important is that once users have been granted any authorized explicit right to the information system, they are now considered an insider and are part of the system protection state. The malicious insider is therefore someone who violates the protection state of the system and is depicted as the root node of the tree representation. The four subordinate nodes are the specific types of actions a malicious insider may perform. [33]

- An insider as anyone in an organization with approved access, privilege, or knowledge knowledge of information systems, information services, and missions. A malicious insider (MI) is one motivated to adversely impact an organization's mission through a range of actions that compromise information confidentiality, integrity, and/or availability. This research explores three fundamental hy- potheses motivated by our study of MIs. [34]

- A disgruntled insider with knowledge of the victim's system. (see also abuse of privilege, insider attack, internal vulnerability, insider) [35]

- The potential risk that employees and officers of a company can cause more harm to the IT infrastructure o threats such as viruses and cracker attacks. Also known as an "authorized user threat," disgruntled employ especially if their feelings are not made public. [36]

- An insider threat is a malicious hacker who can be an employee or officer of a business, institution, or agency. An insider threat is also known as a cracker or black hat. It can also refer to an outside person who poses as an employee or officer by obtaining false credentials. The cracker obtains access to the computer systems or networks and then conducts activities intended to cause harm to the enterprise. [37]

- For the purposes of this paper, an insider is “a human entity that has/had access to the information system of an organization and does not comply with the security policy of the organization”. This definition does not define the type of access (logical or physical, existing or revoked). Also, it does not define the level of skill required by the insider to meet his objectives. [38]
- They’re actually talking about attacks caused by insiders, not ”insider threats.” working with their language, an insider threat would be ”those who misuse or destroy sensitive or confidential information, as well as it equipment that houses this data.” [39]
- Insider attacks are attacks by users with privileged knowledge about a system. [40]
- An insider is anyone who has approved access, privilege, or knowledge of information systems, information services and missions. A malicious insider is one motivated to adversely impact an organization’s mission by taking action that compromises information confidentiality, integrity, and/or availability. [41]
- Illegal insider trading. Most of the time when someone mentions insider trading, they are referring to the illegal kind. Basically, this happens when someone buys or sells stock based of confidential, non-public information. The purchase or sale of stock in these scenarios constitutes a violation of responsibilities or duties by the trader. The insider may be an officer or employee of the company, or it may be someone who has inside information as a result of some kind of professional or contractual relationship with the company. [42]
- Tipping refers to the act of inside information. Tipping is illegal and if the person who is tipped makes trades based on the information, it is classified as illegal insider trading as well. Examples of people likely to be involved in tipping include family members, friends and business associates of someone who has insider information on a stock. [42]
- An entity with authorized access that has the potential to harm an information system through destruction, disclosure, modification of data, and/or denial of service. An entity with authorized access (i.e., within the security domain) that has the potential to harm an information system or enterprise through destruction, disclosure, modification of data, and/or denial of service. [43]
- Einwechter defines insider threat as someone entrusted with authorised access who instead of fulfilling assigned responsibilities, manipulate access to a system to exploit it. [44]
- An ‘insider’ is a person that has been legitimately given the capability of accessing one or many components of the IT infrastructure, by interacting with one or more authentication mechanisms (plain text password, PKI, biometric or smart card token). [45]
- The term “insider” can be defined as any user that either “currently or at one time was authorized to access an organization’s information systems.” This applies to users that are physically located within an organization as well as users that are logically within an organization. The term user is extended to agents, a process created and run by a user. This takes into account scheduled jobs and processes that may be executed without the user being present or without their knowledge. External attacks can become or considered insiders through the proxy of a current insider. A current insider could perform an activity that meets the desires of an external attacker such as executing malware. The ‘insider threat’ refers to harmful acts that insiders might carry out. Harmful acts can be a number of activities including computer abuse, unauthorized access, disruption of service, loss of integrity and theft. Insiders have greater privileges and knowledge of their organization than external attackers. Insiders have the trust of their organization and normally goal oriented. [46]