Guest Editorial: Emerging Trends in Research for Insider Threat Detection

William R. Claycomb[1], Philip A. Legg[2], and Dieter Gollmann[3]
[1] *Software Engineering Institute, Carnegie Mellon University, USA*
claycomb@cert.org
[2] *Department of Computer Science, University of Oxford, UK*
phil.legg@cs.ox.ac.uk
[3] *Security in Distributed Applications, Technische Universität Hamburg-Harburg, Germany*
diego@tuhh.de

The insider threat is one of mankind's most enduring security challenges. For as long as people have placed trust in one other, they have faced the risk of that trust being violated. Historically, consequences of insider attacks included compromised organizational security, financial loss, and risks to human health and safety. Prior to the information age, attacks mainly targeted tangible assets, such as people or money; now insider attacks target additional assets related to information technology (IT), such as data and systems. For instance, malicious insiders may steal intellectual property, sabotage corporate IT systems, or use IT systems to commit financial fraud. Insider attacks have plagued humanity for millennia, and researchers and security professionals continue to struggle to fully understand the breadth of the problem and to propose solutions proven to have measurable effects on reducing the occurrence and impact of attacks. Even defining "insider threat" can be problematic, depending on the problem space. One definition used in the IT security arena is as follows:

> A malicious insider threat to an organization is a 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. [1]

This definition focuses on information systems; in a general sense, though, it could be argued that an "insider threat" is anyone trusted by another to safeguard something of value. Thus, anyone within an organization in a position of trust (i.e. practically everyone) could be considered an insider threat. The misuse or violation of that trust is an insider /emphattack, and occurs at the expense of the trusting entity.

We believe the challenge to detecting the malicious insider threat, either before (hopefully) or during/after an attack, includes two distinct and necessary components:

- the ability to observe malicious or potentially malicious behavior of the insider
- the ability to determine that the observed behavior presents a threat to the organization

Observing Malicious or Potentially Malicious Behavior

Many times an attack is preceded by *observable events* that compose *indicators*, or clues, about current or future malicious behavior. Here we refer to "observable events" as actions by the insider that can be observed, either by other people or by automated systems. For instance, a person may be observed entering an access-controlled area either by a security guard stationed nearby, or by an automated badging system that logs successful badge-swipes required prior to entry. Sometimes indicators are single

Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications, volume: 5, number: 2, pp. 1-6

events, such as copying the entire customer database onto removable media. Other times indicators are comprised of a combinations of events, such as logging in outside normal working hours and printing large amounts of corporate intellectual property. Indicators can take many forms, such as anomaly-based, threshold-based, rule-based, or model-based. [2]

Often, the line between observable and non-observable events is unclear. For instance, some organizations may have the ability to monitor every action a user takes on corporate IT systems, while other organizations either do not have that ability or choose not to utilize it, perhaps due to technical, practical, ethical, or legal limitations. Should we consider events that can be observed, but are not, as "observable events"? What about events that occur outside organizational boundaries or on personally owned devices? What about emotions? Some might consider emotions to be non-observable, but there are outward-facing behaviors (i.e. irritability, lack of contentiousness, etc.) that give clues about a person's underlying emotional state. Some researchers even suggest that emotional states such as stress can be detected by measuring differences in how a user interacts with IT systems, such as their typing patterns, or "keystroke biometrics." [3]

When observing human behavior, we often consider two distinct types of activities: behavioral (i.e. interpersonal; human to human contact), and technical (i.e. human interaction with IT systems). Both type of activities are important to monitor, and indicators can comprise both behavioral and technical events. We refer to these types of indicators as /emphsociotechnical.

Determining that Observed Behavior Presents a Threat to the Organization

To date, many researchers in the information security field have studied the problem of identifying real insider attacks among large sets of potential indicators returned by detection systems. Some of that work, including papers presented in this special issue, are related to anomaly detection, particularly among very large data sets. Results are promising, but two very important questions persist regarding the relationship between anomalous and malicious behavior[1]:

1. How often is malicious behavior anomalous?
2. How often is anomalous behavior malicious?

Regarding question one, we know that some - but not all - malicious behavior is anomalous. Consider an employee who often copies confidential presentations onto removable media to take to offsite meetings - behavior that is authorized by the company. However, the employee becomes disgruntled, and one day, while traveling between the office and a meeting, the employee stops at a competitor's office and allows the competitor to copy the sensitive information from the removable media. In this case, the observable behavior (copying the confidential presentation to removable media) was malicious, but not anomalous. Based on empirical evidence, in many cases of insider attacks, the perpetrator used authorized access and actions that were part of normal business processes. [1]

Similarly, we can say for question two that some - but not all - anomalous behavior is malicious. Anyone with experience configuring detection systems can confirm this; false positives often far outnumber true positives when searching for indicators of malicious behavior. Consider an indicator based on abnormal printing activity: an employee could easily trigger this indicator for a variety of reasons, such as preparing a portfolio of documents to read on an upcoming trip, printing background materials for a visiting colleague, or preparing multiple copies of a report to distribute during a meeting.

So we are left with the lingering question of where the ground truth of these relationships lies. It almost certainly depends on many variables, such as the type of organization, the type of indicator

[1] In this case, because we rely on observed events to detect malicious behavior, we must define "anomalous" as the deviation from *observed* normal behavior, rather than the presence of malicious intent in the user's actions.

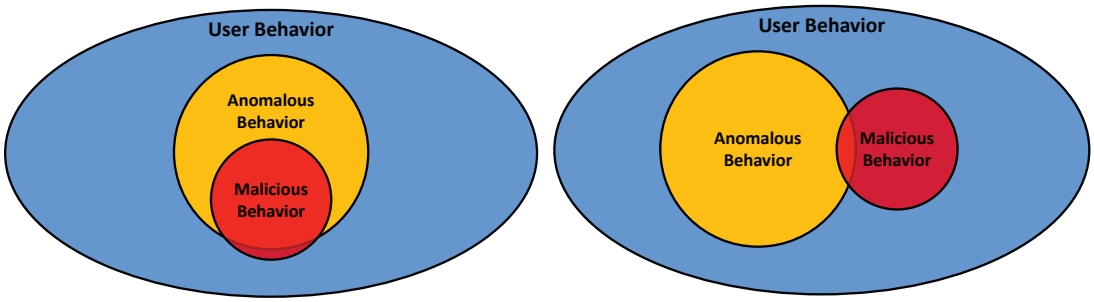

(a) Most Malicious Behavior is Anomalous (b) Most Malicious Behavior is Not Anomalous

Figure 1: Two Possible Answers to the Unanswered Question Regarding the Relationship Between Anomalous and Malicious User Behaviors

being detected, factors about the insider, etc. But even knowing where to start looking is often difficult to ascertain. Though the relationship between anomalous and malicious depends on the performance of the detector or classifier, it seems for now that most often anomalous behaviors are not malicious. Conversely, we have a smaller heuristic base to refer to when considering the question of whether or not the majority of malicious insider behaviors are anomalous. Figures 1a and 1b illustrate this quandary. We have not seen an answer to the question yet, but researchers continue to pursue this and other pertinent topics related to detecting and preventing malicious insiders, as evidenced by the works collected here.

Current Works

This special issue collects five papers that discuss different aspects of insider threat detection. Four were selected for invitation to this journal from papers submitted to and presented at the 2nd IEEE Symposium on Security and Privacy Workshop on Research for Insider Threat (WRIT 2013)[2].

- The first article 'Insider Threat Defined: Discovering the Prototypical Case" [4] presents the results of a survey conducted among 30 cybersecurity experts, exploring their opinions on the attribute of a prototypical insider and insider threat case. The survey is based on attributes in an Entity-Relationship Model developed in the authors' previous study of 42 different definitions of insider and insider threat. [5].

- In the second article, "Differentiating User Authentication Graphs" [6], the authors build upon their previous work [7] to represent user authentication activity as a set of user-specific graphs over an enterprise network, finding that certain types of user behavior have distinguishable graph attributes. More specifically, they demonstrate significant distinction between system administrators and non-privileged users. The paper also explores the differentiation of other functional organization-based user categories. In addition, due to the operational value user authentication graphs have in reflecting user behavior, the authors discuss the development of a system for visually presenting the graphs. This system enables exploration and validation of both appropriate and anomalous user behavior relevant to both intrusion and insider threat detection.

- The next article "Multi-Source Fusion for Anomaly Detection: Using Across-Domain and Across-Time Peer-Group Consistency Checks," [8] presents two methods for robust anomaly detection in multi-dimensional data. The authors extend previous work [9] to describe information fusion

[2]WRIT 2013 was held in San Francisco, California, USA, on May 24, 2013. `http://www.sei.cmu.edu/community/writ2013/`

across multiple levels in a layered architecture. The goal is to ensure accurate and reliable detection of anomalies from heterogeneous data. The paper considers the problem of detecting anomalous entities (e.g., people) from observation data (e.g., activities) gathered from multiple contexts or information sources over time. The authors illustrate the performance of their proposed approaches on an insider threat detection problem using a real-world work-practice data set.

- In the fourth paper "Invalidating Policies using Structural Information" [10], the authors present a step towards detecting the risk of insider attacks by invalidating policies using structural information of an organizational model. Based on this structural information and a description of the organization's policies, their approach invalidates the policies and identifies exemplary sequences of actions that lead to a violation of the policy in question. Based on these examples, the organization can identify real attack vectors that might result in an insider attack. This information can be used to refine access control systems or policies. The authors provide case studies showing how mechanical verification tools, i.e. model-checking with MCMAS and interactive theorem proving in Isabelle/HOL, can be applied to support the invalidation and thereby the identification of the attack vectors. This paper builds on previous work by the authors [11].

- The final paper, "Generating Test Data for Insider Threat Detectors" [12], addresses the difficult challenge of obtaining suitable data for research, development, and testing of insider threat prediction and detection. Building on previous work [13], the authors outline the use of a synthetic data generator to enable research progress, while discussing the benefits and limitations of synthetic insider threat data, the meaning of realism in this context, comparisons to a hybrid real/synthetic data approach, as well as future research directions. The primary task described in the paper is the generation of data to simulate the aggregated collection of logs from host-based sensors distributed across all the computer workstations within a large business or government organization over a 500 day period. The authors note that producing synthetic data with a high level of human realism is considerably more difficult than producing synthetic data simply to test performance of classification systems.

In conclusion, we would like to extend our special thanks to Dr. Ilsun You who is the EiC of JoWUA. Without his help, this special issue could not have been published on time.

Bill Claycomb, Phil Legg, and Dieter Gollmann
Guest Editors
June 2014

References

[1] D. M. Capelli, A. P. Moore, and R. F. Trzeciak, *The CERT Guide to Insider Threats*. Addison Wesley, 2012.

[2] K. Ilgun, R. Kemmerer, and P. Porras, "State transition analysis: a rule-based intrusion detection approach," *IEEE Transactions on Software Engineering*, vol. 21, no. 3, pp. 181–199, Mar 1995.

[3] R. Maxion, "Keynote: Detecting cognitive state for operators of cyber-physical systems: Design of experiments," in *Proc. of the 43rd Annual IEEE/IFIP Conference on Dependable Systems and Networks Workshop (DSN-W'13)*, June 2013, pp. 1–1.

[4] D. A. Mundie, S. Perl, and C. H. J.D., "Insider threat defined: Discovering the prototypical case," *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications (JoWUA)*, vol. 5, no. 2, pp. 7–23, June 2014.

[5] D. A. Mundie, S. Perl, and C. L. Huth, "Toward an ontology for insider threat research: Varieties of insider threat definitions," in *Proc.o f the 3rd Workshop on Socio-Technical Aspects in Security and Trust (STAST'12), New Orleans, Los Angeles, USA*. IEEE, June 2013, pp. 26–36.

[6] A. D. Kent, L. M. Liebrock, and J. Wernicke, "Differentiating user authentication graphs," *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications (JoWUA)*, vol. 5, no. 2, pp. 24–38, June 2014.

[7] A. D. Kent and L. M. Liebrock, "Differentiating user authentication graphs," in *Proc. of the 2013 IEEE Security and Privacy Workshops (SPW'13), San Francisco, California, USA*. IEEE, May 2013, pp. 72–75.

[8] H. Eldardiry, K. Sricharan, J. Liu, J. Hanley, B. Price, O. Brdiczka, and E. Bart, "Multi-source fusion for anomaly detection: using across-domain and across-time peer-group consistency checks," *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications (JoWUA)*, vol. 5, no. 2, pp. 39–58, June 2014.

[9] H. Eldardiry, E. Bart, J. Liu, J. Hanley, B. Price, and O. Brdiczka, "Multi-domain information fusion for insider threat detection," in *Proc. of the 2013 IEEE Security and Privacy Workshops (SPW'13), San Francisco, California, USA*. IEEE, May 2013, pp. 45–51.

[10] F. Kammüeller and C. W. Probst, "Invalidating policies using structural information," *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications (JoWUA)*, vol. 5, no. 2, pp. 59–79, June 2014.

[11] F. Kammüller and C. W. Probst, "Invalidating policies using structural information," pp. 76–81, May 2013.

[12] B. Lindauer, J. Glasser, M. Rosen, and K. Wallnau1, "Generating test data for insider threat detectors," *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications (JoWUA)*, vol. 5, no. 2, pp. 80–94, June 2014.

[13] J. Glasser and B. Lindauer, "Bridging the gap: A pragmatic approach to generating insider threat data," pp. 98–104, May 2013.

Author Biography

Bill Claycomb Bill Claycomb is the Lead Research Scientist for the Insider Threat Program at Carnegie Mellon University's Software Engineering Institute. His primary research topics on insider threats include indicator development and validation, sociotechnical models, data integration, and linguistic patterns. Dr. Claycomb is also involved in efforts related to cloud computing security, incident response, systems modeling, and vulnerability analysis. He is an adjunct faculty member at CMU's Heinz College, teaching in the School of Information Systems and Management.

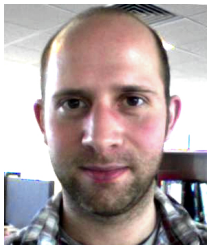

Phil Legg Phil Legg is a Post-Doctoral Research Associate of the University of Oxford, UK. He is a member of the Cyber-Security research group at the Department of Computer Science. His current research interests include machine learning techniques, data visualization, image processing and computer vision. Prior to his current role, he has also worked on projects that include sports video visualization, medical imaging and mobile application development.

Dieter Gollmann Prof. Dieter Gollmann received his Dipl.-Ing. in Engineering Mathematics (1979) and Dr.tech. (1984) from the University of Linz, Austria, where he was a research assistant in the Department for System Science. He was a Lecturer in Computer Science at Royal Holloway, University of London, and later a scientific assistant at the University of Karlsruhe, Germany, where he was awarded the 'venia legendi' for Computer Science in 1991. He rejoined Royal Holloway in 1990, where he was the first Course Director of the MSc in Information Security. He moved to Microsoft Research in Cambridge in 1998. In 2003, he took the chair for Security in Distributed Applications at Hamburg University of Technology, Germany. Prof. Gollmann is an editor-in-chief of the International Journal of Information Security and an associate editor of the IEEE Security & Privacy Magazine. His textbook on 'Computer Security' has appeared in its third edition.