Detecting Malicious Insiders in Military Networks

Mark Maybury
The MITRE Corporation
202 Burlington Road
maybury@mitre.org

Keywords: insider threat, malicious insider, information assurance, cyber indications and warning, observables taxonomy, assets, data fusion, attack graphs, honeypots

ABSTRACT

Given that a network is only as strong as its weakest link, a key vulnerability to network centric warfare is the threat from within. This paper summarizes several recent MITRE efforts focused on characterizing and automatically detecting malicious insiders within modern information systems. *Malicious insiders (MI)* adversely impact an organization's mission through a range of actions that compromise information confidentiality, integrity, and/or availability. Their strong organizational knowledge, varying range of abusive behaviors, and ability to exploit legitimate access makes their detection particularly challenging. Crucial balances must be struck while performing MI detection. Detection accuracy must be weighed against minimizing time-to-detect and aggregating diverse audit data must be balanced against the need to protect the data from abuse. Key lessons learned from our MI research include the need to understand the context of the user's actions, the need to establish models of normal behavior, the need to reduce the time to detect malicious behavior, the value of non cyber-observables, and the importance of real-world data collections to evaluate potential solutions.

1. The Threat: Malicious Insiders

An *insider* as anyone in an organization with approved access, privilege, or 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. Analysis of the behavior of dozens of malicious insiders (DSS 1999, Herbi and Wiskoff 2002) and our detailed analysis of six representative cases (Maybury et al. 2004) such as CIA's Aldrich "Rick" Ames, FBI's Robert Philip Hanssen (2003), and DIA's Ana Belen Montes (2001) illustrates the diversity of the insider threat challenge. Each of these cases is unique in terms of their position, motive, foreign handlers, impact, sentence, computer skill, polygraph experience, cyber security violations, counter intelligence activities, physical and cyber access, cyber extraction and exfiltration, cyber communication, and the transfer of materials to foreign handlers. The devastating impact of these three individuals alone included the violation of confidentiality, undermining of intelligence integrity, adverse influence of US policy, the revelation of sources and methods, and the death and compromise of field agents. MI motives are diverse, ranging from financial to thrill to ideological. In each of these cases, handlers were professional foreign service agents. Two of the three passed polygraphs. While the computer skills of each of these insiders ranged significantly, all left trails of suspicious cyberactivity while performing cyber access, exfiltration, and/or communication. All engaged in counter intelligence to evade detection and/or destroy incriminating evidence. In each case we found opportunities to observe individual incidents and/or to detect anomalous behavior from correlated observables.

2. Approach and Research Agenda

Our study of MIs supports the finding that there is no single silver bullet solution to the problem. Accordingly, we take a wholelistic approach which incorporates prevention, detection, and reaction. In this paper we focus on our efforts aimed at automated detection. Our analysis has led us to explore several fundamental hypotheses including:

1. While some MIs can be detected using a single cyber observable, other MIs could be detected only by using multiple and heterogeneous observables.

2. Fusing information from heterogeneous information sources (e.g., logs from printers, authentication, card readers, telephone calls) and various levels of the IP stack (e.g., application vs. network traffic) allows more accurate and timely indications and warning of malicious insiders. Even with a single sensor, if you monitor a broad range of activities it will increase your detection rates.

3. Observables together with domain knowledge (e.g., user role, asset value to mission) can help detect inappropriate behavior (e.g., need to know violations).

Our basic approach is consistent with an overall strategy that aims to prevent, detect, and react to insider threats while balancing privacy and security. Our research methodology includes conducting studies under the auspices of an independ-

ent review board (IRB) together with measures of anonymization and aggregation to ensure the protection of privacy.

The remainder of this paper first summarizes our experience in an insider threat challenge workshop to assess the ability of several distinct sensors approaches to detect three simulated insiders on a live network. We then describe an initiative to develop a broad set of context-sensitive rules and fuse individual indicators into overall threat scores to highlight potentially abusive behavior. Input is based on passive, network-based sensors that monitor how users interact with information taking advantage of models of context of users and information. We conclude by identifying lessons learned from our investigations as well as future research directions.

4. Insider Threat Challenge Workshop

In order to enhance understanding of and accelerate solutions for the insider threat, a collaborative, six month challenge workshop was held to characterize and create analysis methods to counter sophisticated malicious insiders (Maybury et al. 2004). Following a careful study of past and projected cases, several prototype techniques were developed to provide early warning of insider activity, including novel algorithms for structured analysis and data fusion. The algorithms were assessed in an operational network against three distinct classes of human insiders (an analyst, application administrator, and system administrator), measuring timeliness and accuracy of detection, which we subsequently describe.

4.1 Simulated MIs: Pal, Jill, and Jack

Grounding our efforts in realistic insider behavior, we explored detecting three types of insiders in detail in this activity. The first was a historical insider modeled as a prototype of past need-to-know violators. We call this insider Pal. A second insider, named Jack, was a projected insider who would aim to disrupt, damage, or destroy the network or elements thereof. In the course of defining and simulating these insiders, the scenario team implemented a third category of insider, an application administrator, called News Admin or Jill. Only Pal's behavior model was disclosed to sensor builders prior to the experiment. For detail about these insiders including a log of specific actions taken by the insiders see Maybury et al. (2004). The three malicious insider cases were simulated on MITRE's Demilitarized Zone (DMZ) network. The DMZ consists of over 300 hosts with a range of missions utilizing services such as web (HTTP), news (NNTP), file transfer (FTP), messaging (SMTP), mail (POP, IMAP), database (SQL), and question answering. We instrumented 18 of 31 nodes on the NRRC (Northeast Regional Research Center) subnetwork which had 75 on-line, active users during the evaluation.

A semi-automated process captured, filtered, and anonymized the malicious insider collection to address security and privacy concerns. Figure 1 illustrates the heterogeneous nature of the collection consisting of over 11 million records which spans physical sensors (e.g., employee badge readers), network level sensors (e.g., Snort rules modified to detect inappropriate connections or behavior), host sensors (to detect user access and command sequences), and applications (e.g., mail server logs, web server logs, network news logs). A Common Data Repository (CDR) was established as a central database storing the over 11 million anonymized, time stamped audit-log records collected over three months.

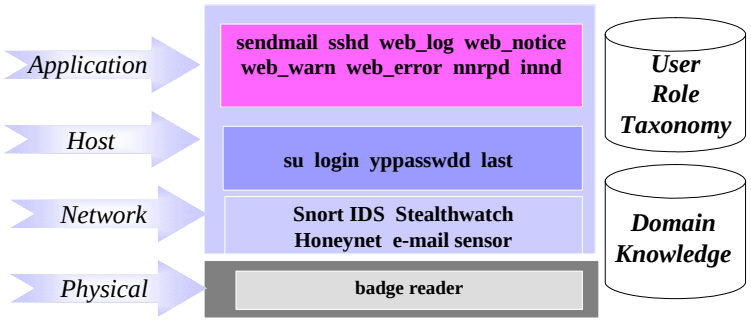

Figure 1. Heterogeneous and Multilevel Data Sources

4.2 Event and Observable Taxonomy

In order to access, exploit, or damage assets, a MI will necessarily need to perform (or have another person or process perform) a series of actions to gain privileges, access or manipulate assets. Derived from our analysis of MI cases, Figure 2 shows a taxonomy of cyber events which have associated observables that hold promise for the foundation of a detection system. The taxonomy distinguishes observables in the cyber domain from those in the physical domain. The taxonomy includes observables such as results of the polygraph, records of security violations, missing or misleading reports on finances, foreign travel or foreign contacts, physical facility access, personal finances, materials transfer, counter intelligence, social behavior, and communications. In this research we focused exclusively on cyber observables, including other observables that could be readily converted to a cyber signal (e.g., digitized facility access logs).

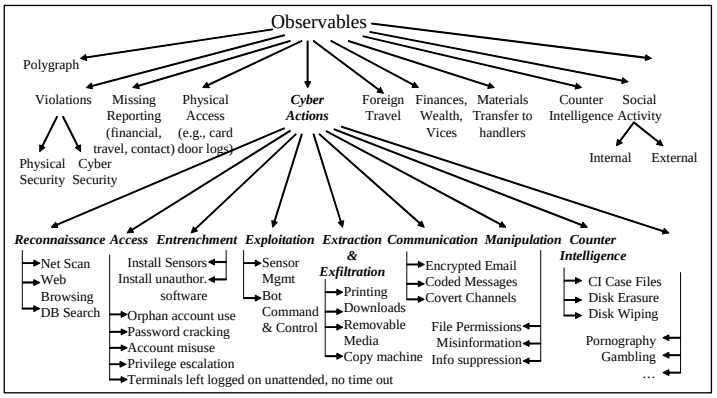

Figure 2. Cyber Event/Observable Taxonomy

The core of the taxonomy incorporates a range of cyber observables encompassing a range of classes of cyber actions indicated in bold italics in Figure 2. These include activities of network, system, and information reconnaissance, access to assets (e.g., media, hosts, accounts), entrenchment (e.g., installing sensors or unauthorized software), exploitation (e.g., commanding and controlling entrenched assets such as software bots or zombie ma-

chines), extraction and exfiltration (e.g., of hardcopy, media, information), communication (e.g., encrypted messaging, encoded messages, covert channels), manipulation of cyber assets (e.g., changing file permissions, suppressing or altering information content), counter intelligence (e.g., wiping disks), and other cyber activities associated with unethical or addictive behavior (e.g., on line gambling). Some observables have been used in some historical cases as a tip-off of malicious activity; others serve as direct indicators of inappropriate behavior.

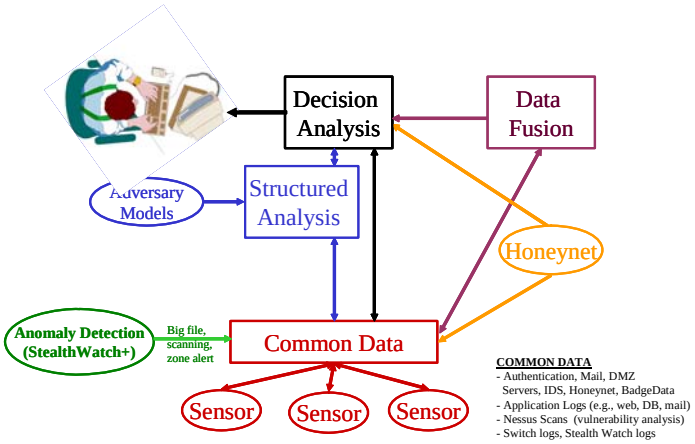

Figure 3. Integrated Architecture for Insider Detection System

4.3. Insider Detection

While the live network instrumentation describe in Section 3 provided an unprecedented and essential set of MI experimental data, the thrust of our activity was developing novel algorithms to detect MIs. Figure 3 illustrates the high level architecture of a proof of concept system that was designed, implemented, and tested to detect MIs. Distributed, heterogeneous sensors provide input to a Common Data Repository (CDR) from which a range of analyses are performed including data fusion and structural analysis to identify potential suspects on a watch list or issue an alert of an insider threat. As illustrated in Figure 3, our technical approach is novel in the following respects:

- A Common Data Repository (CDR) captures and anonymizes heterogeneous sensor input.
- Multilevel monitoring occurs at the packet level, system level, and application level.
- StealthWatch sensors detect abnormal insider behavior on the network such as scanning, file transfer, or internal network connections.
- Distributed honeynets acquire attacker properties, pre-attack intensions, and potential attack strategies.
- A real-time, top-down structural analysis drawing upon functional models of MIs maps pre-attack indicators to models of potential MIs.
- Traditional and non-traditional indicators (e.g., logs of network activity, physical access, PBX, help desks), including non-digital sources, are fused bottom-up.

Sensor inputs are then exploited by a decision analysis component to determine watch list membership and insider detection. We next consider each of the primary detection strategies.

4.4 HoneyTokens

Honeypots are realistic but dummy systems that reflect true production systems and are designed to attract malicious users to inappropriately access resources. Combined with subtlely advertised enticements to potential insider threats, honeypots provide a mechanism to determine what motivates the inside attacker and what capabilities the attacker possesses.

A novel idea developed during the workshop and applied in the insider detection process is the notion of a *honeytoken*. A honeytoken is a semi-valuable piece of information whose use can be readily tracked. This could be a credit card number, an Excel spread sheet, a database entry, or a login and password. A honeytoken is an entity that has no authorized use. Honeytokens can be used for the initial detection of insider threats, then those threats can be redirected to honeynets to confirm if a violation has occurred, potentially learning more about the threat.

In the Pal scenario described in Section 3, the honeytoken takes the form of a web page which lists (fictitious) operatives in the geographic region of interest to the MI. The data fusion group, detailed in a subsequent section, exploits the detection of honeytoken access as one of a range of indicators of malicious behavior. In other examples of honeytokens (e.g., login password) it is possible to use the false information to track activities (e.g., in a controlled account) to more readily discover MI actions, capabilities, and intentions.

4.5 Stealthwatch

One of the sensor types deployed in support of the insider threat workshop was StealthWatch. StealthWatch (http://www.lancope.com) provides traffic profiling (e.g., data flow analysis) and host profiling such as zone profiling (e.g., monitoring which hosts can talk to which). StealthWatch creates a concern index (CI) based on flow analysis, reconnaissance activity, and anomaly detection. Any host reaching an accumulated CI level above 20,000 generates an alarm. Stealthwatch was deployed to monitor scanning, downloads, and inside connections. StealthWatch enabled an analyst to identify Jack within a day of his malicious activity, which initiated on February 11th. However, StealthWatch did not detect Pal and Jill who exhibited no scanning, download or insider connection anomalies.

4.6 Structured Analysis

Using the cyber observable framework introduced at the beginning of this article in Figure 2, the structured analysis approach considered a range of cyber observables associated with adversaries to perform top-down, real-time, model-based detection of MIs. Knowledge of network mission and network configuration together with likely adversarial actions were used to filter and highlight adversary behavior. MI actions modeled included reconnaissance (e.g., via web browsing or net scans), access (e.g., privilege escalation), en-

trenchment (e.g., via deployment of sensors), extraction and exfiltration (e.g., unauthorized downloads), and communications (e.g., coded messages or covert channels). A range of sensors and data logs are relevant to detecting many of the cyber actions. For example email communication patterns would provide insight into the social network of a malicious insider. Conversely, large downloads might signal data exfiltration.

The structural analysis group (SAG) modeled two insiders, Pal and Jack, considering temporal characteristics of protocols such as event proximity (e.g., immediate vs. days vs. years) and observable ordering. The Pal detector exhibited 3% false positives and no false negatives and the Jack detector had 1% false positives and 50% false negatives.

The structured analysis approach detected all three simulated MIs, Pal, Jill, and Jack, at various times. Pal was put on the watch on December 11th, 2003, two days after Pal initiated behavior. Six days later an alert was issued for Pal. The News Administrator Jill was put on the watch list on January 24th, 2004, even before her suspicious behavior started on February 11th. Four days later, on February 15th, SAG issued an alert for Jill. Finally, SAG issued an alert on 17 February for Jack, six days after he initiated malicious administrator behavior.

4.7 Data Fusion

In contrast to the preceding model-based, top-down analysis of observables, a data fusion approach explored bottom-up combination of observables from multiple sensors including the combination, correlation and analysis of both traditional and novel indicators stored in the Common Data Repository. This low-level data came from available resources on the MITRE DMZ network along with inputs from other teams (StealthWatch and honeynets). Thus data fusion occurred over input from the physical, network, host and application levels. Fusion can occur along a number of dimensions such as the type of sensor (e.g., card reader, authentication, printer, telephone calls) and the level of IP stack (e.g., from network to application).

A proof-of-concept data fusion engine was developed that:

- Fused multiple indicators that were related by IP, user name, etc. into a single indicator
- Allowed customer selectable weights for each element in the observable taxonomy
- Allowed site specific criteria for generating an alert and for submission of names to a watch list based on:
  - Accumulated indicator weight
  - Accumulated indicator breadth, i.e., the number of unique sources of evidence

In addition, a novel sensor was developed that performed email consistency checking. The email sensor analyzed PGP signatures and PGP-encrypted attachments and generated error messages to report if the PGP data did not decode properly. A simple, flexible data fusion engine was implemented in the JESS rule system.

Figure 4 illustrates a series of logged messages from the operation of the data fusion engine that report various inferences being made on the basis of sensor input for the scenarios on the DMZ network introduced in Section 3.

The data fusion engine was able to find indications of malicious activity in two out of the three scenarios, with little tuning. One in-the-wild incident was discovered, in particular a user who was excessively downloading news. Perhaps equally significant is the substantial data reduction that resulted. The data fusion engine needed to examine 7.4 million records to analyze the activity of the 75 users on the system. 259 indicators were generated for 24 of those users. The data fusion approach exhibited very low false positive rates (0 or 3%) with moderate to high false negative rates (33 to 66%) using a threshold of the maximum change between any two detected users.

Cyber-Access, ***user324***, weight 1, at 2003-12-10 11:14:38, from news.mitre.org
su to user9676 failed for non-admin user user324 on /dev/pts/0

Physical-Access, ***user295***, weight 5, at 2003-12-15 19:19:37,
After hours badge access for user295

Cyber-Extraction-Exfiltration, ***user2649***, weight 5, at 2004-01-06 15:37:28, from nrrc-springfield.mitre.org,
Data was uploaded to an external server via FTP protocol

Cyber-Reconnaissance, user295, weight 10, at 2004-01-09 20:57:18, from nrrc-springfield.mitre.org,
User user295 searching in non-need-to-know country korea

Cyber-Communication, ***user9***, weight 15, at 2004-02-10 22:14:48, from cvw.mitre.org,
User user9 received email with masqueraded content from user11649@yahoo.com
---
Cyber-Reconnaissance, ***user1***, weight 5, at 2004-02-10 13:54:15, from nrrc-plymouth.mitre.org,
Ongoing CI violation -- 066.170.227.074 currently has 49613 alerts of this type…

Cyber-Extraction-Exfiltration, ***user295***, weight 8, at 2004-02-12 23:54:58, from dmzsrv1.mitre.org,
User user295 sent encrypted email to user9983@comcast.net

Cyber-Extraction-Exfiltration, ***user1***, weight 15, at 2004-02-20 12:25:03, from nrrc-erie.mitre.org,
user1 sent email with masqueraded content to user1@mitre.org

Figure 4. Data Fusion Log

4.8 Evaluation: Accuracy and Speed

We designed evaluation methods to assess both the accuracy and timeliness of insider threat detection. Accuracy was measured in terms of false positives and false negatives ad defined:

- False positives = # users incorrectly detected as malicious insiders
- False negatives = total # actual malicious insiders - # correctly detected

For each sensor, the maximum change between any two detected users is computed, and then false positive and false negative

measures are computed at several points along the graph. Detection results were measured at three weighted thresholds: maximum change, 50% of the highest factor, and 10% of the highest factor.

In contrast to accuracy, timeliness were measured as the difference between the time when an insider's malicious activity begins, the time at which they are put on a watch list, and the time at which an insider threat alarm occurs. We consider these measures stricter than would be found in an operational environment which would include a human analyst in the loop to make alert decisions (Stealthwatch results did include a human analyst, whereas the other methods were autonomous).

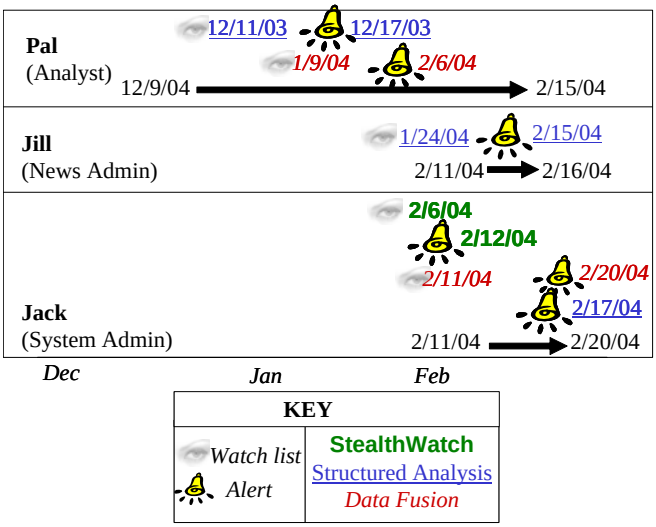

Figure 5. Timeliness of Detection Results

Figure 5 summarizes the timeliness performance of the StealthWatch, structured analysis, and data fusion detectors for the three insiders: Pal, Jill and Jack. One objective was to reduce the time from defection to the time of detection from years to months to weeks to days if not minutes. In Figure 5, the eye icon indicates the day when the sensor put the insider on the watch list. The bell indicates the first day when an alert is issued. The black vectors indicate the start (left of the vector) and stop (right of the arrow head) times of the malicious behavior by the three human MIs. The bold font dates are associated with the StealthWatch sensor (which only detected Jack), the underlined dates are for the structured analysis method, and the italicized ones are for the data fusion performance. In summary, when reviewed across all sensors and methods, of our three MIs, two were detected within one week of their initiation of suspicious activity and the third was detected within two weeks. Even removing some more obvious indicators such as the scanning behavior of Jack, because a multiplicity of sensors provide evidence for inferences Jack would still be detected.

5. Context Sensitive Malicious Insider Detection

One goal is to help experts identify insider threats as quickly as possible, ideally before they can do damage. Developing profiles of a typical employee's online behaviors—how they search, how they access networks, how they print documents—and "adding context" to those actions, can help pinpoint behaviors that fall outside normal boundaries. Adding context involves human effort along with the electronic detective work. This includes understanding users' role, normal relationships with others, and normally information usage patterns.

One of the key challenges illustrated in the insider threat challenge workshop is the fact that inappropriate behavior for one user could be considered appropriate for another. In a related initiative, we have developed a broad set of context-sensitive rules and have fused individual indicators into overall threat scores to highlight potentially abusive behavior. Input is based on passive, network-based sensors that monitor how users interact with information taking advantage of models of context of users and information.

Analysts must dig deeper to find out if the insider has a legitimate reason to access the materials in question by considering what is this staff's role, with whom do they normally interact, and what information are they trying to get. For example, if a computer technician starts searching on terms that an analyst would use, this might send up a red flag. Accordingly we have Been developing and testing information-use sensors and user attribution techniques, along with the development and testing of context-sensitive rules that help users determine when insiders transfer files in unusual ways or suddenly change their information-seeking behavior.

Accordingly some of the key elements of the technical approach include:

- Monitor how users interact with data by sensing and translating the network protocols tied to information use
- Establish methods to attribute events to users (vs. IP addresses).
- Deploy software agents to collect user and information context
- Develop a broad set of context-sensitive rules to highlight potentially-abusive behavior, and
- Combine indicators into a scoring system to prioritize threats.

For a given user, sensors issue an alert, which is linked to an algorithm that helps us find the probability that the user is malicious. The analyst is given a *threat score*, a ranking of all the employees in the organization based on this probability. Our accomplishments to date include:

- The development and successful testing of our information-use sensors, user attribution techniques, and data anonymization routines.

- The collection of a large, realistic, real-world background data set representing over 300 days of activity and over 3000 users.
- The crafting and execution of 8 malicious insider scenarios and the tools to integrate them into the background data set for testing purposes.
- The development and testing of context-sensitive rules that detect printing anomalies, reconnaisance, search and acquisition of distant or unusual information, unusual file movement, changes in information seeking behavior, and evasive information seeking

Future plans include the refinement of a user interface and analysis tool to help analysts further refine their searches for malicious actors.

6. Summary

Malicious insiders pose perhaps the most serious threat to organizational cyber assets. Malicious insider behavior is distinct from that of classical external intruders and cannot be detected using traditional intrusion detection methods. In this article, we report results from a challenge workshop that demonstrated how an integration of multiple approaches promises early and effective warning and detection for a range of insider threats. We also report our efforts to create context sensitive malicious insider detection.

However, while this research makes initial contributions to the malicious insider, it equally raises many new research directions. These include the need for more refined malicious insider models, more elaborate cyber actions/observables taxonomies, more comprehensive test corpora, and more sophisticated detection algorithms. Effective counter MI programs should encompass protection, detection, and reaction elements and must address challenges of data fusion, sensor accuracy, real time detection, and privacy.

Acknowledgments

Special appreciation to Jeff Sebring, Greg Stephens, and Rich Pietravalle who provided information security expertise and to Penny Chase, Laurie Damianos, and Tony Ricciatti who acted as malicious insiders.

The challenge workshop effort was performed at The MITRE Corporation at the Northeast Regional Research Center (NRRC) which is sponsored by the Advanced Research and Development Activity in Information Technology (ARDA), a U.S. Government entity which sponsors and promotes research of import to the Intelligence Community which includes but is not limited to the CIA, DIA, NSA, NGA, and NRO.

References

1. Anderson, Robert H.; Bozek, Thomas; Longstaff, Tom; Meitzler, Wayne; Skroch, Michael; and Van Wyk, Ken. August, 2000. Research on Mitigating the Insider Threat to Information Systems - #2. Workshop Proceedings. http://www.rand.org/publications/CF/CF163.

2. [DSS 1999] Recent Espionage Cases 1975-1999 (Defense Security Service). Security Research Center. Defense Security Service. Monterey, California September 1999. http://www.dss.mil/training/espionage

3. [Hanssen 2003] A Review of the FBI's Performance in Deterring, Detecting, and Investigating the Espionage Activities of Robert Philip Hanssen August 14, 2003 Office of the Inspector General. http://www.usdoj.gov/oig/special/03-08/index.htm

4. Hayden, Lt Gen Michael V. July 1999. The Insider Threat to U.S. government Information Systems. National Security Telecommunications and Information Systems Security Committee (NSTISSAM) INFOSEC 1-99. http://www.nstissc.gov/Assets/pdf/NSTISSAM_INFOSEC1-99.pdf

5. Herbig, Katherine L. and Wiskoff, Martin F. July 2002. Espionage Against the United States by American Citizens 1947-2001. Defense Personnel Security Research Center PERSEREC-TR 02-5. http://www.ncix.gov/news/2002/oct/Espionage.pdf

6. Jones, Anita K. (chair). November 1-2, 2001. White Paper: Cyber-Security and the Insider Threat to Classified Information. Computer Science and Telecommunications Board, National Research Council. http://www7.nationalacademies.org/CSTB/whitepaper_insiderthreat.html

7. Matzner, Sara Nov. 2004. Approaches to Insider Threat Mitigation. *ISSA Journal*, Feature article pp.6-8.

8. Matzner, Sara and Tom Hetherington. Summer 2004. Detecting Early Indications of a Malicious Insider", *IA Newsletter*, 7(2): 42-45.

9. Maybury, Mark, Sebring, Jeff, Chase, Penny, Chiekes, Brant, Pietravalle, Richard, Costa, Mick, Zarrella, Guido; Gaimari, Bob; Brackney, Dick, Lehtola, Penny, Matzner, Sara; Hetherington, Tom; Marin, Jack; Wood, Brad; Sibley, Conor; Longstaff, Tom; Spitzner, Lance; Haile, Jed; Copeland, John; and Lewandowski, Scott. 2004. Insider Threat Challenge Workshop: Final Report. MITRE Technical Report 04B-14.

10. Montes, Anna. September, 2001. Affidavit. http://news.findlaw.com/hdocs/docs/montes/usmontesaff901.pdf

11. Shaw, Eric D.; Post, Jerrold M.; and Ruby, Kevin G. Inside the Mind of the Insider. http://www.securitymanagement.com/library/000762.html

12. Spitzner, Lance. “Honeypots: Catching the Insider Threat” ACSAC, Las Vegas, Dec 2003

13. Webster, William H. (chair). March 2002. A Review of FBI Security Programs Commission for Review of FBI Security Programs. U.S. Department of Justice. http://www.fas.org/irp/agency/doj/fbi/websterreport.html