NAVAL POSTGRADUATE SCHOOL

MONTEREY, CALIFORNIA

THESIS

A SURVEY OF METHODS FOR DETECTING INTENTIONAL INSIDER THREATS AGAINST DIGITAL SYSTEMS

by

Meagan A. Bridgeman

June 2021

Thesis Advisor: Neil C. Rowe
Second Reader: Victor R. Garza

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			*Form Approved OMB No. 0704-0188*

1. AGENCY USE ONLY *(Leave blank)*	**2. REPORT DATE** June 2021	**3. REPORT TYPE AND DATES COVERED** Master's thesis

4. TITLE AND SUBTITLE
A SURVEY OF METHODS FOR DETECTING INTENTIONAL INSIDER THREATS AGAINST DIGITAL SYSTEMS

5. FUNDING NUMBERS

6. AUTHOR(S) Meagan A. Bridgeman

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Naval Postgraduate School
Monterey, CA 93943-5000

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
N/A

10. SPONSORING / MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.

12a. DISTRIBUTION / AVAILABILITY STATEMENT
Approved for public release. Distribution is unlimited.

12b. DISTRIBUTION CODE
A

13. ABSTRACT (maximum 200 words)

Insider threats are a serious problem in all organizations and continue to be an issue despite technological advances. Insider threats can undermine cybersecurity by subverting controls or exploiting weak systems. These subversions can be enabled by an inadequate security policy, poor policy implementation, or new attack innovations. This thesis identifies the tactics that can be used by an intentional insider threat to subvert policies, and examines whether sufficiently reliable tools and policies are available to prevent and detect this type of behavior. We conclude with policy and technical control recommendations.

14. SUBJECT TERMS
insider threat, insider threat detection, insider threat deterrence

15. NUMBER OF PAGES
71

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	**18. SECURITY CLASSIFICATION OF THIS PAGE**	**19. SECURITY CLASSIFICATION OF ABSTRACT**	**20. LIMITATION OF ABSTRACT**
Unclassified	Unclassified	Unclassified	UU

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. 239-18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release. Distribution is unlimited.

A SURVEY OF METHODS FOR DETECTING INTENTIONAL INSIDER THREATS AGAINST DIGITAL SYSTEMS

Meagan A. Bridgeman
Lieutenant, United States Navy
BA, University of Texas at Dallas, 2014

Submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE IN COMPUTER SCIENCE

from the

NAVAL POSTGRADUATE SCHOOL
June 2021

Approved by: Neil C. Rowe
Advisor

Victor R. Garza
Second Reader

Gurminder Singh
Chair, Department of Computer Science

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Insider threats are a serious problem in all organizations and continue to be an issue despite technological advances. Insider threats can undermine cybersecurity by subverting controls or exploiting weak systems. These subversions can be enabled by an inadequate security policy, poor policy implementation, or new attack innovations. This thesis identifies the tactics that can be used by an intentional insider threat to subvert policies, and examines whether sufficiently reliable tools and policies are available to prevent and detect this type of behavior. We conclude with policy and technical control recommendations.

THIS PAGE INTENTIONALLY LEFT BLANK

TABLE OF CONTENTS

LIST OF FIGURES

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF TABLES

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF ACRONYMS AND ABBREVIATIONS

APWG	Anti-Phishing Working Group
AUC	Area under the receiver operating-characteristic curve
CERT	Computer Emergency Response Team
CMU-CERT	Carnegie Mellon University's Computer Emergency Response Team
CNN	Convolutional neural network
DDoS	Distributed denial of service
DNN	Deep neural network
DOD	Department of Defense
FIRST	Forum of Incident Response and Security Teams
FISMA	Federal Information Security Modernization Act
GDPR	General Data Protection Regulation
GFIRST	Government Forum of Incident Response and Security Teams
GINA	Genetic Information Nondiscrimination Act
HIMSS	Healthcare Information Management Systems Society
HIPAA	Health Insurance Portability and Accountability Act
HTTP	Hypertext Transfer Protocol
IBM	International Business Machines Corporation
INSA	Intelligence and National Security Alliance
LDAP	Lightweight Directory Access Protocol
LSTM	Long short-term memory neural network
LZW	Lempel-Ziv-Welch compression algorithm
NIST	National Institute of Standards and Technology
PAM	Privileged-access management
PCI-DSS	Payment Card Industry Data Security Standard
RDP	Remote Desktop Protocol
RNN	Recurrent neural network
SNA	Social-network analysis
SSH	Secure Shell Protocol
SVM	Support-vector machine

SIEM	Security-incident and event management
TCP/IP	Transmission Control Protocol / Internet Protocol
UBA	User-behavior analytics
UEBA	User-and-entity behavior analytics
USB	Universal Serial Bus

ACKNOWLEDGMENTS

To my husband, you were instrumental to my success, providing the love and support needed to get through this thesis. This thesis would not have been possible without you.

To my advisors, the guidance you provided while allowing me the freedom to explore this topic was a perfect balance. Thank you for keeping me on track and for helping to keep the scope manageable.

THIS PAGE INTENTIONALLY LEFT BLANK

I. INTRODUCTION

A. THE PROBLEM

Much study and investment in protecting networks from outside threats have been made; however, many attacks are launched from within an organization's network ("insider threats"). Such cases have steadily risen over the last several years. Insider threats can be caused by user negligence or intentional malicious actions. An intentional insider threat is defined by the Intelligence and National Security Alliance (INSA) as

> the threat presented by a person who has, or once had, authorized access to information, facilities, networks, people, or resources; and who wittingly, or unwittingly, commits: acts in contravention of law or policy that resulted in, or might result in, harm through the loss or degradation of government or company information, resources, or capabilities; or destructive acts. (Alba, 2015, para. 3)

This thesis focuses on intentional malicious actions involving people knowingly trying to curtail security policies and tools to harm an organization or gain personally. While most insider threats are due to user negligence, the cost to an organization is highest for intentional insider threats (IBM Security, 2020). In 2020, a report showed that of 204 organizations that had insider-threat activity—63% were due to negligence and 23% were due to criminal insiders. However, the estimated annualized cost for incidents due to negligence was $4.58 million, while the cost for criminal insiders was $4.08 million. While the amount of insider-threat activity is debated, the damage of even one incident can be high in sensitive domains when it involves proprietary information, classified information, or private personal information.

B. RESEARCH GOALS

The primary goal of this thesis is to determine if the continued issue with insider threats is due to problems with security implementation, security-policy adoption, or insufficient security policies. From reviewing current literature and studies, recommendations are formulated to deter or detect intentional insider threats. These recommendations are based on some new studies. Consolidating both past and current

efforts to address insider-threat detection methods will allow future work on insider-threat detection methods to be efficiently selected. Also, these recommendations aim to focus attention on the most vulnerable areas for improvement.

This thesis proposes a method that is focused on denial then detection. In particular, focusing on deterrence first will raise the barrier for intentional insider threats that lack the technical capability to circumvent a system and reduces the number of insider-threat attacks of opportunity. Policies and automated tools must be in place to detect insider threats exploiting administrator or privileged access. We hope to show the importance of investing in strategies and tools that work on fixing these threats.

C. OTHER CHAPTERS

Chapter II will explore previous methods for handling intentional insider threats and their effectiveness in deterring or detecting them. The chapter will examine tools such as network-traffic analysis, user-behavior analytics, and privilege-access management. Chapter III will explore the categories of insider threats and their motivations, including the most common methods used. Chapter IV will study the observed effectiveness of deterring or detecting intentional insider-threat activity and discussions of ethical or legal implications. Chapter V will recommend related methods and policies to improve an organization's security. Chapter VI will give conclusions and recommendations for future work.

II. PREVIOUS WORK

A. WHAT ENABLES AN INSIDER THREAT?

Many studies have been done on how outside attackers infiltrate protected networks. However in recent years, more focus has been given to protecting a network from inside an organization's protected walls. One important issue is what enables an insider threat in the first place. This involves three features: motive, opportunity, and capability.

1. Motive

Since insider threats often engage in criminal activities, it helps to consider what can motivate someone to commit a crime. Financial issues, foreign influences, and disgruntlement could motivate someone to become an insider threat. One study identified four categories of motivations: predisposition to malicious behavior, mental disorders, personality factors, and current emotional state (Gheyas & Abdallah, 2016). A predisposition to malicious behavior can be challenging to detect as it requires an insider to try to execute a malicious behavior. Some studies have used honeypots on an internal network to see if insiders attack them as an indicator of insider activity (Spitzner, 2003; Levine et al., 2003). Mental disorders such as depression, paranoia, and schizophrenia can create a higher risk of an individual becoming an insider threat. Personality factors and current emotional states are the two areas often brought up in the discussion of insider-threat detection (Schultz, 2002). Potential insider threats observed had a "history of managing crises ineffectively, pattern of frustration, sense of inadequacy, aggrandized view of their abilities and achievements, a strong sense of entitlement, views self above the rules, and actions that seek immediate gratification." (Gelles, 2016, p. 9). Relevant emotional states can be an increased hostility, intolerance of criticism, inability to take responsibility for one's actions, and strong dissatisfaction with the employer (Gheyas & Abdallah, 2016).

2. Opportunity

A crime cannot be committed without an opportunity. This opportunity can occur due to an insider's role within an organization, such as being a system administrator, or due to the improper creation, implementation, or enforcement of security controls.

A malicious insider can cause more harm with more unfettered access to a system. This potential becomes greater with insufficient controls such as auditing system logs, insufficient application of the principle of least privilege, and resiliency measures to ensure no single point of failure. Opportunity is another way to say how easy it is to commit a crime. Having inadequate security policies or unenforced policies can allow an insider to exfiltrate data or install malicious software easily. Combining technical security controls and enforcement of security policies reduces an insider threat's potential to attack successfully.

3. Capability

When security controls are adequate, an insider threat's technical abilities must increase to subvert them. Sophisticated insider threats have the technical know-how to perform advanced techniques and cover their tracks without detection by auditing systems. Also, if the sophisticated attacker knows how the detection schemes work, they can tailor their actions to circumvent them. The most dangerous insider threats have both the capability and the opportunity to commit a crime, such as a system administrator with little supervisory oversight, no separation of duties, and who is the single point of failure within the organization.

B. TYPES OF INSIDER THREAT ACTIONS

Insider threats can harm an organization in three main ways: sabotage, theft, and fraud (CERT Software Engineering Institute, 2018). The deliberate destruction, damage, or obstruction of information-technology infrastructure can be costly for an organization. Information-technology sabotage examples range from system administrators who refused to reveal administrator passwords, installing backdoors to connect remotely into servers and shutting them down, and deletion of critical data (CERT Insider Threat Center, 2010).

While employees in any position can commit information-technology sabotage, system administrators can do the most damage and are the most common perpetrators of this type of crime.

When an employee develops code for an application for their previous employer, they may mistakenly think they own the code. Policies and laws differ between organizations about whether the employee has the right to their work to show future employers. Some organizations consider this as theft of the organization's intellectual property. The Federal Bureau of Investigation defines intellectual property theft as stealing ideas, inventions, or creative expressions from people or companies. Some examples are trade secrets, proprietary information, or parts of a movie or music. An example of this type of behavior is an insider committing industrial espionage for a competing company by stealing proprietary information.

Fraud can be defined as the wrongful or criminal deception intended to result in financial or personal gain. Insider threats commit fraud to "steal because of a sense of entitlement, and some who want to exact revenge against an organization simply because they are angry" (CERT Insider Threat Center, 2010, para. 1). Fraud can be creating a fictitious employee or vendor who gets paid by an organization, altering payroll data, using personal information for identity theft, and misuse of corporate credit cards. Fraud can be committed in many ways, and in 2012, Carnegie Mellon University's Computer Emergency Response Team (CMU-CERT) found that 71 percent of insider frauds were through non-technical means (CERT Software Engineering Institute, 2012).

C. CURRENT METHODS TO DETECT INSIDER THREATS

As technology advances, so have the methods and tools to protect organizations and their information from misuse. In the last five years, advancements in artificial intelligence and machine learning have been integrated into insider-threat programs, tools, and policies.

A survey of 204 organizations and 964 people not associated with military organizations that have experienced insider-threat events reported their security measures included user training, data-loss prevention, user-behavior analytics, employee monitoring,

security-incident and event management, incident-response management, third-party vetting procedures, threat-intelligence sharing, privileged-access management, and network-traffic intelligence (IBM Security, 2020). Table 1 shows the observed frequencies of these methods to be explored in the following sections. While more technical means can subvert systems that protect organizational networks, such as using covert channels to exfiltrate data, we will only consider the more obvious ways.

Table 1. Tools and Activities That Reduce Insider Threats. Source: IBM Security (2020).

Security tools & activities	Frequency of companies	Percentage of companies
User training & awareness	112	55%
Data loss prevention	110	54%
User behavior analytics	102	50%
Employee monitoring & surveillance	96	47%
Security incident & event management	91	45%
Incident response management	89	44%
Strict third-party vetting procedures	87	43%
Threat intelligence sharing	85	42%
Privileged access management	80	39%
Network traffic intelligence	77	38%

1. User Training and Awareness

User training and awareness help any security policy. It can reduce both intentional and unintentional insider-threat incidents. While user training is well suited for preventing users from becoming unintentional insiders, it also helps by recognizing and reporting suspicious behavior. DOD insider-threat training reviews spotting suspicious behavior and uses case studies to provide real examples of this suspicious behavior. However, this is a passive form of security-awareness training. More active methods for training include phishing campaigns done internally to test an organization's staff (Intelligence and National Security Alliance, 2019).

Insider-threat training programs should depend on position within the organization. Leadership should know about the impacts that insiders can cause and the importance of

reinforcing security policies. General staff should also know how to report indicators of suspicious activity and measures to protect against becoming an unintentional insider threat. Insider-threat training should not associate malicious insiders with only well-known attributes such as social isolation. It should include employees bragging about how much damage they could do to the organization, downloading data before termination, using company resources for their own business, or trying to access other employees' accounts by social engineering (Collins et al., 2016).

While it is difficult to quantify user awareness in preventing or detecting insider threats, it is essential to any security policy. However, these programs cannot guarantee to discourage insiders from committing crimes. Intentional insider behavior often involves disgruntlement, and user training does not help stop this (Gelles, 2016). Moreover, relying upon other employees to report suspicious behaviors is difficult without a private or anonymous way to do so (Collins et al., 2016).

2. Data-Loss Prevention

The second most used mitigation is data-loss prevention (IBM Security, 2020). These are tools and techniques for protecting sensitive data from being deleted, accessed by unapproved users, or exfiltrated. These methods are commonly used to comply with regulations such as Health Insurance Portability and Accountability Act (HIPAA), Payment Card Industry Data Security Standard (PCI-DSS), or General Data Protection Regulation (GDPR). They have components of critical-data identification, leak detection, protection of data at rest, protection of data in use, and protection of data in motion.

Critical-data identification is usually done by rules such as those for identifying credit-card or Social Security numbers. Most commercial data-loss prevention solutions offer regular expression to detect such data and address compliance standards automatically. Other tools can determine a document's privacy classification from a training set of correctly classified documents and materials. This can be done using regular expressions, keywords, and hashing (Hart et al., 2011).

Data-loss protection of data at rest is done by encryption and access control. Private information should be encrypted when stored on servers and endpoints. Additional

protection can be provided by scanning to determine if sensitive data is in the wrong location. For data in use, some methods try to establish typical use patterns for users; others prohibit certain activities with the data, such as copying and pasting, printing, and screen capturing (Shabtai et al., 2012). For data in transit, other products detect and inspect data flows for suspicious activity on protocols such as email and Hypertext Transfer Protocol (HTTP). Some solutions do more detailed analysis by packet inspection.

Data-loss prevention has several weaknesses. One weakness is the reliance on automated methods on training data, where they may incorrectly classify newer proprietary information as non-critical when it should be classified as critical. Another issue is how easily the classification system can be tricked by replacing keywords or paraphrasing information to bypass filters (Shabtai et al., 2012). Machine learning can improve classification accuracy but works best in controlled environments such as a virtual machine (Kongsgård et al., 2016). Hash-based data-loss prevention systems can be tricked by replacing a single character to change the hash. Another weakness is the inability to read encrypted data. An organization can block all encrypted files from leaving the organization, but this could violate privacy policies and could hurt worker productivity, especially in health care and finance. Lastly, one survey showed the inability of automated systems to correlate a user completing a prohibited action (such as emailing a marked document) and a subsequent allowed action such as printing the document.

In summary, data-loss prevention can combat confidential information loss but is prone to error and requires careful configuration. The current tools are insufficient to combat a malicious insider threat and should be combined with other techniques such as user-behavior analytics and access controls.

3. User-Activity Profiling

User-activity profiling is a subset of user-and-entity behavior analytics (UEBA). This includes monitoring other entities such as unmanaged endpoints and cloud applications. User-activity profiling can be user-based or role-based.

User-based profiling identifies abnormal behavior of a user by comparing their actions to a baseline, looking for changes beyond some threshold that would require

investigation (Dean & Rowe, 2018). Machine-learning methods set a threshold during a training period that learns to identify normal behavior for each user. Once numeric values are more than a specified number of standard deviations outside this threshold, or if predefined suspicious actions are taken, such as deleting log files, an alert would be triggered for that user (Legg et al., 2017). This can work if malicious behavior occurs only rarely in an organization. However, this method can generate false positives through errors during training or for anomalous non-malicious changes in behavior (Dean, 2017).

User-activity profiling can be made more accurate by a role-based grouping of user behaviors rather than job titles. Users who perform the same tasks should act similarly, and behavior grouping will address people that perform multiple job roles within an organization. However, grouping by job title helps for new users with little historical data, and it may be better to group users by behavioral patterns instead of roles (Dean, 2017).

These detection schemes are useful if a malicious insider tries to do unauthorized actions such as escalating privileges or tries to access a sensitive folder. However, they are less successful at detecting legitimate actions that are misused, such as a user printing a much larger volume of documents than previously or compared to their peers. Handling such cases requires setting acceptable levels of false positives. Many algorithms have been suggested with false-positive rates between .01% to 54% depending on the data's training and which type of anomalous behavior was trying to be detected (Azaria et al., 2014). However, it is unclear if any available tool uses these algorithms since vendors withhold that information.

User-activity profiling was used by 50% of the companies surveyed in a recent report, although detection methods were unclear (IBM Security, 2020). Many tools perform "user baselining" but do not indicate what factors contribute to a threat score. Nonetheless, commercial products provide many features such as security-incident and event-management integration, malware detection and prevention, additional monitoring for privileged users, and forensic tools.

4. Alerting Tools

Security-incident and event management, incident-response management, and vetting procedures are used with the other tools to prevent, detect, or respond to possible insider-threat incidents. All three were used by over 40% of the companies surveyed.

Security-incident and event management tools aggregate security data from multiple sources to act as an alert system for security incidents. It can integrate with firewalls, network-traffic analysis, workstation-log information, intrusion-detection systems, antivirus software, and some user-activity profiling software. Security-incident and event-management tools issue alerts based on configured rules. They can detect security violations such as excessive file copying, distributed denial-of-service (DDoS) attacks, and brute-force attacks. They can enforce HIPAA, GDPR, and other regulatory laws and can be used to investigate incidents. However, they have some of the same issues that the previously mentioned tools have, as they require careful calibration to avoid excessive false positives. Depending on the rules, security-incident and event management can trigger automated responses such as locking accounts or turning off a workstation, but doing this too often hurts productivity. Another issue is that many solutions do not prevent actions from happening but only alert once the action has already occurred. Also, many insider threats act within the organization's security parameters, and rules may not detect such actions.

Incident-response tools can also help with insider threats. These tools can analyze log files of the host where the incident occurred, create after-action reports, and patch vulnerable systems. Some organizations maintain dedicated incident response teams, while others form task forces once an incident has been detected (Ahmad et al., 2012). Proper expertise in investigating the cause of an incident is important to prevent missed issues. Good incident response should include root-cause analysis, and lessons learned should be reported outside of the incident-response team. This study also showed that much more attention was given to high-impact incidents, but the lower-impact incidents should also be analyzed to find the root causes that could be more serious.

Third-party vetting was another tool used by many companies for reducing insider threats by examining something or someone carefully to determine suitability. In a 2019 survey of 107 organizations, the three most common methods for pre-employment vetting were employment verification, federal criminal court-record screening, and state and local court-record screening (George et al., 2019). Other sources of vetting found are shown in Table 2.

Table 2. Processes for Pre-Employment Vetting. Adapted from George et al. (2019).

	Process Used
1	Employment and work history verification
2	Federal criminal-court record screening
3	State and local criminal-court screening
4	E-Verify screening
5	Credit checks
6	Education verification
7	Professional references
8	Personal-address history verification
9	Fingerprint background checks
10	Personal references
11	Military-history verification
12	Civil-court record screening
13	Sexual-offender screening
14	Drug screening
15	Licensure, certification, registration verification
16	Driving-history screening
17	Regulatory-sanctions screening
18	Social-media screening and open-source research
19	International background screening
20	Other

Vetting can also be done when awarding contracts. Many regulatory organizations maintain lists to include blacklisted companies and people based on previous vetting or reports of misconduct. Sources of such lists are the Excluded Parties List, the Office of Foreign Assets Control, the U.S. Customs, the Securities and Exchange Commission, the

State Department list of foreign terrorist organizations, the Food and Drug Administration, and the World Bank (Appel, 2017).

While it is common to vet incoming personnel, few organizations do a continuous evaluation (George et al., 2019). This is an issue with malicious insider threats as many, such as Edward Snowden and Chelsea Manning, had previously cleared background investigations. The RAND Corporation recommends a robust continuous evaluation program (Luckey et al., 2019). Both CERT and RAND recommended mandatory behavioral monitoring of people who were either terminated, voluntarily or involuntarily, for a minimum of 30 days both before and after an employee has left the organization as many incidents occur then (Collins et al., 2016; Luckey et al., 2019). A weakness of the traditional vetting process is the lack of inspection of online activities (Appel, 2017). This can cause an incomplete assessment of potential employees to identify behavioral indicators that could indicate a reason for distrust.

5. Privileged-Access Management

Privileged-access management, also called privileged-access security or privileged-identity management, is grounded in the principle of least privilege. This is the concept of giving an individual only the minimum rights and privileges necessary to complete their work. Ways to implement privileged-access management include role-based rights, segregation of duties, and reliable employee termination procedures. Role-based rights restrict users further to only what is necessary to complete a role in their jobs. This can limit the amount of damage a malicious insider could do (Collins et al., 2016). For example, a Human Resources employee should not need access to files in the Sales department.

One potential problem is "privileged-access" creep, where users gain additional access rights over time that are not revoked when no longer needed, such as after a departmental transfer or promotion; audits must be done regularly. Also, using role-based access rights can be challenging to configure when an employee has multiple roles, which is typical in smaller organizations (Collins et al., 2016). Role-based access also does not address the misuse of legitimate access as by the many insiders in technical positions who have committed information-technology sabotage. Lastly, role-based access privileges can

impede organizational productivity. If users are given only the minimum amount of access usually required, it may be too cumbersome to gain the new access needed to do their jobs well. Furthermore, users may be encouraged to find ways around security controls (Stolfo et al., 2008).

A related principle is the segregation of duties, which disallows any one person from making critical changes to a system or being responsible for monitoring it. Two or more people should be required to sign off on significant changes to critical systems, do backups, and make major system changes (Collins et al., 2016). This reduces the chance of sabotage or theft; many of the most damaging incidents involved an insider who was an organization's sole system administrator. Another method is to give system administrators two segregated accounts, a "superuser" account to make essential changes and a standard user account for other day-to-day activities. This reduces the likelihood of mistaken changes and the amount of log data required to review if an incident occurred.

Limiting access rights on digital systems is especially important for terminated-employee accounts. A frequent recommendation is to monitor an account for 30 days, both before and after termination. This prevents an ex-employee from using remote-access tools to infiltrate and attack an organization after termination (Collins et al., 2016). While most issues with privileged-access management are tied to the configuration of a system, commercial services can address maintenance. They provide solutions for onboarding, auto-discovery for tracking privileged accounts, automated temporary access, and customized reporting.

6. Network Traffic Analysis to Detect Anomalies

Network-traffic analysis was only used by 38% of the companies surveyed. It can help identify anomalous behavior and identify critical assets related to insider threats. Anomalous behavior in network traffic can be abnormal traffic spikes, unusual login times, or unverified remote-access sessions. Insider-threat activity that can be identified includes accessing sensitive files, email patterns, Web browsing, excessive downloads, and suspicious software installations (Collins et al., 2016). Email is important as most insiders perpetrating theft use email to exfiltrate data (Moore et al., 2012). However, it may be

impossible with encrypted traffic, which conceals some actions from detection. It also can be incorrect to attribute activity to a user if an insider used their account. Lastly, as with all anomaly-based systems, an accurate baseline and thresholds must be established. Logs of all network activity will not protect an organization very well from malicious insider activity. However, network traffic analysis can identify an organization's most often used services, databases, and servers for reducing the damage a malicious insider can achieve (Collins et al., 2016). Identifying critical assets is essential to creating effective risk-management policies.

7. Threat Intelligence Sharing

Sharing of threat intelligence, and specifically, insider-threat intelligence, can aid other organizations in increasing their security. Many organizations and groups issue such information, such as the Forum of Incident Response and Security Teams (FIRST), the Government Forum of Incident Response and Security Teams (GFIRST), and the Anti-Phishing Working Group (APWG) (Cichonski et al., 2012). These groups provide information to response teams and security professionals about real-world cyber threats.

However, this information is focused more on outside threats to an organization than insider threats. Insider threat intelligence sharing was used by 42% of companies surveyed, but it is unclear how much of that was about negligent insiders instead of malicious ones. Another obstacle to threat-intelligence sharing is the reluctance of companies to share information that may harm their reputation, and an insider breach or intellectual-property theft could do that (Wagner et al., 2019). While the Federal Information Security Modernization Act (FISMA) requires reporting to a centralized authority all incidents in government organizations, private organizations have no such requirements (Cichonski et al., 2012).

III. TYPES OF INSIDER THREATS

Intentional insider threats come in many forms with different names. Each name categorizes insider threats based on their motivation. In this chapter, we will examine the categories of insider threats and the most common methods for theft, fraud, and sabotage of information technology systems. While workplace violence has been called an insider threat, we will not discuss here threats that are primarily physical.

A. CATEGORIES OF INSIDER THREATS

Insider threats can be categorized into four types based on motive: malicious, vengeful, virtuous, and wicked (Thompson, 2019). These categories apply to both intentional and unintentional insider threats. Intentional insider threats tend to be vengeful or malicious.

A malicious insider threat engages in deliberate destructive behavior (Thompson, 2019). Malicious insiders typically do sabotage and theft. The most common motivation suggested in the literature is a disgruntled employee. No distinct psychological profile automatically makes someone disgruntled. However, some personality characteristics increase risks, such as a strong sense of entitlement, a pattern of frustration, disappointment, and a sense of inadequacy (Gelles, 2016). These can occur for several reasons, such as being passed up for promotion or a lack of recognition. A way to reduce disgruntlement is to reduce the feelings of unmet expectations through clear communication from leadership and consistent enforcement of clear policies. Consistent enforcement will reduce the chance that employees feel that they are treated differently, or that management is held to a different standard.

The vengeful insider "willfully acts out against supervisor or co-workers and believes in the greater good of the organization/system" (Thompson, 2019) and desires to harm the organization knowingly. An example of vengeful insider behavior would be ignoring prescribed security practices to cause trouble for their supervisor. A vengeful insider can become malicious after becoming increasingly disgruntled or after seeing their actions significantly affect the organization. Both vengeful and malicious insiders are

prime targets for foreign or competitor influence, as when a foreign agent sees an employee negatively speaking about their employer on social media.

The virtuous insider is a well-intentioned employee that places an organization at risk through their behavior. An example is an employee who tries to meet deadlines by taking documents home and leaving them unprotected. This type of threat can be reduced through user training and awareness.

The wicked insider is like the virtuous insider, except this employee knows they are deliberately breaking policies to get the job done while still appearing to be well-intentioned. A wicked insider can also disregard security rules for self-interest by copying inventory reports from previous weeks to go home earlier. A wicked insider still may not know the full extent of the risk they are putting on the organization. This type of threat would be considered unintentional since the intent to harm the organization is absent.

B. OTHER MOTIVATIONS FOR INTENTIONAL INSIDER THREATS

Other factors that can cause an employee to become a malicious insider threat are financial gain, foreign influence, or ideological differences. An insider can exploit access to a system for their financial gain through fraud and intellectual property theft. Fraud is when an insider incorrectly changes an organization's data or steals information for identity theft for personal gain. Theft of intellectual property is when an insider takes confidential information from a business for personal gains, such as selling it to a competitor or as a condition for new employment. According to Carnegie Mellon University's Computer Emergency Response Team (CMU-CERT), the sector most affected by fraud is the financial, followed by healthcare and governments (CERT Software Engineering Institute, 2018). Indicators that an employee could be an insider threat due to financial motivation are defaulting on a loan, a high debt-to-income ratio, a current or previous gambling problem, or greed (CDSE, 2019). Other indicators that an employee has committed fraud or intellectual property theft would be sudden changes in their financial situation, such as expensive purchases or unexplained affluence.

Financial issues can also make people open to influence from foreign adversaries. Some insider threats became spies for money, coercion, ideology, ego, disgruntlement,

ingratiation, or thrills (Charney & Irvin, 2016). This influence can be coerced through concern over family members' safety in another place or foreign country. Private businesses can be susceptible to malicious insiders who commit theft by performing economic espionage for business competitors for money or because they are disgruntled. Ideological threats occur when beliefs or philosophies from a group or person contradict the organization they serve. These threats can occur from disagreements with an aspect of an organization's function, such as subscribing to Communism while serving in a democratic organization (Charney & Irvin, 2016). However, it can be uncertain if an insider commits the crime due to their belief in that ideology or uses it to justify their actions.

C. METHODS USED BY INSIDER THREATS

Insider threats usually achieve goals of theft or fraud by exfiltrating data. Information technology sabotage is most commonly done through destructive changes.

1. Exfiltration Methods for Theft or Fraud

Most insider crimes reported were by non-technical means (CERT Software Engineering Institute, 2012). Advanced techniques, such as using covert channels to avoid detection, are rare. The main methods reported were sending data by email, printing, faxing, smartphones, universal serial bus (USB) storage devices, and cloud storage.

Email is the most common way to forward stolen information (Gelles, 2016). Exfiltrating data by sending emails to personal accounts or accessing Web-based email services like Gmail from within an organization's network has occurred in many case studies (CERT Software Engineering Institute, 2012, 2018). Data has also been exfiltrated directly by sending it to people outside an organization. Confidential information can also be printed and carried out of the organization. Regular audits of the print server help detect this.

Faxing or scanning documents also enables exfiltrating them to a remote site (Collins et al., 2016). Smartphones can take pictures of confidential documents, record sensitive conversations, and provide a mass storage device. Smartphones can then send data to a cloud server, personal email account, or a text messaging site using cellular data

services outside the monitoring of organizational networks. Smartphones can introduce other threats, such as by connecting to otherwise isolated (air-gapped) devices to access remotely. Removable media such as CDs, DVDs, or USB storage devices can allow an insider to record much more data than paper. These devices are also easier to conceal than paper and allow data to be transferred to personal computers or people. Lastly, exfiltrated information stored in the cloud can be accessed from anywhere and with less likelihood of their actions being observed by other employees.

2. Information Technology Sabotage

Information-technology sabotage requires more technical knowledge and sophistication than exfiltration. In many cases, an insider threat that committed the sabotage was a system administrator that had privileged access to systems and used it to degrade or destroy them (CERT Insider Threat Center, 2010; CERT Software Engineering Institute, 2018; Collins et al., 2016). They used technical means such as backdoor installations or logic bombs that made a system inoperable once triggered or at some future date. They also did mass file deletion or refused to provide administrator passwords.

IV. EFFECTIVENESS OF DETERRENCE AND DETECTION OF INSIDER THREATS

Can some insider threats be deterred from acting? Studies show that deterrence can help security-policy compliance and encourage the proper use of information systems. In this chapter, we will examine the overall effectiveness of deterrence, threat profiling, and added controls on digital data in controlling insider threats.

Not much real-world data on insider threats and their methods is available, though there is some synthetic data we will discuss in Chapter 5. This limits the evaluation of the effectiveness of tools discussed in Chapter 2. However, tool effectiveness can be discussed based on increasing the cost to the insider and reducing the risk to the organization.

A. EFFECTIVENESS OF DETERRENCE

Deterrence is the idea that people can be discouraged from an illegal activity if the likelihood of being caught and its cost outweighs the benefits of committing the crime. Deterring an insider threat is complex as it must consider different motivations such as greed, revenge, and other gains. The overall effectiveness of deterrence has been debated in various studies; however, for security policy compliance, some work has empirically measured deterrence effects (Albrechtsen, 2007; Herath & Rao, 2009; Pratt et al., 2017). These studies found that a big factor in deterrence is the perception by the insider that they would likely be caught. However, these studies do not generally consider individual factors such as self-control or personality. Also, an intentional insider threat is harder to deter as their motivation for committing crimes increases. The likelihood of detection may be judged low by some insider threats like system administrators who have expert knowledge of the system investigated. Despite these challenges, some deterrence effects can be accomplished when insiders know good tools are present that will detect them and deterring even a few insider threat incidents reduces the workload on investigators and reduces organizational risk.

B. DETERRENCE AND TECHNICAL CONTROLS FOR INSIDER THREATS BASED ON MOTIVATION

Consistent enforcement of security policies is important for deterrence or detection to work. Ideally, other employees should see enforcement, such as being discussed at departmental meetings and being consistently enforced despite the employee's position. Some studies suggest that making the computer- security team more visible through interactive training makes people more likely to comply with security policies (Albrechtsen, 2007). Consistent enforcement is a prerequisite for both deterrence and detection to increase the insider's perception of the likelihood of being caught. A second issue with deterrence is that people are less likely to follow security policies when they think they will decrease their efficiency. If managers and other employees do not report security violations, this weakens deterrence by reducing the likelihood of being caught further. The reduction in efficiency will also affect detection if the tools and methods recommended require significant work from investigators. Lastly, increasing the severity of punishment is largely ineffective at deterring unwanted behaviors (Herath & Rao, 2009). Therefore, the effectiveness of deterrence depends mostly on the likelihood of an insider threat being caught.

1. Deterring Financial Gain Motivations

Insider threats motivated by financial gain can be deterred and detected more easily than other types discussed later. To deter an insider threat that is motivated by money, the insider must believe the risks of being caught outweigh the potential monetary gain. This can be aided by telling employees of monitoring and detection efforts made by the organization, which can positively affect an employee's likelihood of complying with security policies (Herath & Rao, 2009). However, this tends to deter more those insiders who commit acts based on opportunity rather than those in significant debt. Also, little work has considered the tradeoff between the cost and the potential gain for insider continuing the illicit activity. For these insiders, the potential gain can also be reduced by better detection or increasing the act's difficulty. More regular or randomized audits could increase the detection rate, and most instances of fraud have been found through audits (CERT Software Engineering Institute, 2018). Often clever insiders make small changes

over many months to avoid immediate detection, which is called a "low and slow" approach (CERT Software Engineering Institute, 2012). Increasing the rate of audits increases the likelihood of detecting suspicious activity and lowers the potential damage, such as the amount that could be stolen or fraudulently charged.

Other monitoring methods have been proposed, such as the Behavioral Analysis of Insider Threat (BAIT) framework to identify malicious insiders trying to exfiltrate data for crimes such as identity theft (Azaria et al., 2014). However, this method uses semi-supervised learning and requires labeled malicious insider data, and it cannot identify new attack methods. Other methods that have been proposed have either weak results, such as a precision of only 42%, or have a high probability of overfitting for the most popular synthetic dataset provided by CMU-CERT. Machine learning techniques with deep neural networks show better promise.

2. Deterring Foreign Influence and Espionage

Foreign influence or espionage is difficult to deter since the motivator could be money, coercion, ideology, disgruntlement, ingratiation, or thrills (Charney & Irvin, 2016). Usually, insider threats that are spies provide information for a foreign handler. Their only deterrence would be that they could be caught. However, many believe they are too smart to be caught or fall into a "gambler's fallacy" where they would be very unlucky to be caught more than once (Pogarsky & Piquero, 2003). This resetting effect could be strongest in crimes like low-impact data exfiltration, where an insider threat could be caught and counseled against breaking a security policy. Another issue with foreign influence or espionage is that the insider can become trapped after providing information once and is coerced or blackmailed into continuing to provide information (Charney, 2010). Detecting these insiders requires user-behavior monitoring from logs that record how many times a file has been read, copied, printed, or downloaded. Other useful tools could be Web-site monitoring to see if users are visiting extremist sites, booking foreign travel, trying to escalate their privileges, or doing social engineering of co-workers. For Web site monitoring, natural-language processing could discern a page's topic instead of only looking at the site name.

3. Deterring the Ideologically Motivated

Ideological threats are harder to deter; they involve conflicts between personally held beliefs and those of the organization. The best countermeasures are safeguards to minimize the damage that this threat could cause. An intentional insider threat motivated by their ideology would more likely commit vengeful acts such as sabotage over fraud. Extensive access controls, mandatory multi-person checks on critical system changes, and proper separation of duties are needed for this kind of threat since most information-technology sabotage involves a sole system administrator of an organization. This means requiring at least two people to approve major system updates; ideally, the second person should be someone with enough technical background to know if a system change is necessary. Administrators should not be authorized to make or modify backups, reconfigure the network, install software on critical systems, create new users, or do other major changes on their own. While this does increase the work to make such changes, it can considerably reduce the chances of a security incident. It also helps to separate duties properly, so no one audits a system they otherwise control (CERT Software Engineering Institute, 2018). Other measures that could detect some ideological threats are similar to foreign influence or espionage, including host-based employee monitoring.

4. Deterring Ego-based Motivation

Disgruntlement, ingratiation, and self-importance are motivators that can be difficult to deter but more easily detectable. Disgruntlement has been cited as a major motivator of information-technology sabotage. Deterring a disgruntled employee is unlikely, but such employees are made through a series of events that may be controllable. De-escalation techniques could reduce feelings of disgruntlement. Otherwise, disgruntlement could be detected through host-based user-behavior monitoring and community-behavior monitoring. User-behavior profiling can identify preparatory behavior to more serious acts; community-behavior profiling can help identify outliers to what is considered normal for a group of people, such as accountants. Obvious outliers from community behavior are usage outside of normal business hours and printing or downloading large numbers of documents. Other outlying behaviors could be sending more

emails outside the organization or downloading extensively from the Internet. Identifying outliers in individual behavior is more difficult, but some work associated Facebook "likes" with identifying depression, impulsivity, and life satisfaction with insider threats (Youyou et al., 2015). Some studies further support looking into personality traits because malicious insider threats show these traits in a higher proportion than in the general population (Liang et al., 2016). However, this study extracted feature keywords that may not have been observable until after the incident occurred. Some suggest that it is not these traits but how the individual handles personal failure that creates an insider threat (Charney, 2010). Some possibly observable traits are predatory behavior, personal or work-related conflicts, and problems with financial status. Interestingly, insiders also had more positive traits such as being agreeable, professional, and dedicated to family or work (Liang et al., 2016).

Community-based models could reduce false positives with insider threats by allowing wider tolerance of day-to-day changes that may affect many employees, such as an increased file accesses due to an upcoming audit. Clustering similar behavior instead of roles is better for finding anomalous activity. Even if insider threats only spend a small fraction of their time committing crimes, this activity should stand out when compared against normal activity (Azaria et al., 2014; Dean, 2017). Reducing the damage that a disgruntled employee can cause also requires enforcement of the principle of least privilege. Privileged access management, segregation of duties, and two-person integrity should help to reduce the overall damage that a single insider threat could cause. Also, it is important to ensure that accounts from terminated employees are no longer active and the employees lack remote-access capability.

Motivators of ingratiation or self-importance cannot be easily influenced and would be difficult to detect. Detection of ingratiation may be possible with social network analysis (SNA) to determine relationships between users. In particular could detect people that were not expected to have much two-way correspondence, such as a sales employee talking to someone in finance. These are weak clues; however, over time, a person who does much ingratiation can become disgruntled by feeling like they are not being recognized or appreciated.

Misuse of authorized access is harder to recognize. Special procedures will be needed for auditing it. These policy implementation measures do create an added workload for administrators to properly configure and audit regularly. Table 3 summarizes the motivations by which an insider threat could be deterred or detected with the recommended method.

Table 3. Deterrence and Detectability of Insider Threats by Motivation

Motivator	Deterrable	Detectable	Detection Method
Financial Gain	✓	✓	Auditing/User Behavior Analytics
Foreign Influence/ Espionage	?	✓	Reporting/ User Behavior Analytics /Access Controls
Ideological	×	✓	User Behavior Analytics /Access Controls
Disgruntlement	✓	✓	User Behavior Analytics /Access Controls
Ingratiation	×	?	Social Network Analysis/ User Behavior Analytics
Self-Importance	×	✓	User Behavior Analytics

C. LEGAL AND ETHICAL IMPLICATIONS OF INCREASED EMPLOYEE MONITORING

Many tools and processes can monitor employee performance. More broadly called "people analytics", these tools and processes raise legal and ethical issues. Employers may legally monitor employees through software within the scope of employment. "Within scope" is generally defined as 1) anything that can be seen as job-related, 2) disclosed to

the individual, 3) not genetic information, 4) does not include surveillance of employees at personal locations away from work, and 5) is not judged as a "highly offensive" intrusion into their personal lives (Bodie et al., 2016).

A challenge with data collection on employees is the possibility of introducing bias prohibited by law against employees based on race, sex, religion, especially as defined in the Americans with Disabilities Act (ADA), which includes mental health conditions (Tursunbayeva et al., 2021). Disclosing how employees are monitored and what information is collected is important for complying with many state-mandated regulations that prohibit intercepting communications without specific consent from an employee's phone if the employee reasonably expected privacy (Bodie et al., 2016). This consent must be more than broad consent to monitoring forms, which may be insufficient when significant privacy breaches occur. Genetic information cannot be collected according to the Genetic Information Nondiscrimination Act (GINA) of 2008, and other laws prevent adverse employment actions against employees who do not participate in employee wellness programs.

Employees can now be monitored in their personal spaces by keystroke logging, camera recording, microphone monitoring, and other methods to verify where and how they are working (Tursunbayeva et al., 2021). This now includes working from home, raising legal issues for determining what can be considered a "personal location" and what acceptable monitoring can be done. This includes the monitoring necessary for detecting insider threats. Other legal issues can occur if the organization aggregates data, but courts have generally found an aggregation of non-private information to be unproblematic (Bodie et al., 2016). Organizations should also prevent accidental disclosure of personal data due to improper protection, such as failing to limit access to employee resumes or background investigations.

Increased employee monitoring can affect employees by reducing their performance, raising questions about employee protection from manipulation, and reducing job satisfaction. Ethical issues can occur in gaining employee trust for monitoring, avoiding coercive power in "at-will" employment states, and the lack of legal or ethical precedents for new monitoring technology (Tursunbayeva et al., 2021). In "at-will"

employment states, employers can terminate an employee at any time for any reason except for a few illegal reasons, such as racial discrimination. However, employees can show decreased productivity if the perception of surveillance reduces the trust in the organization (Catrantzos, 2012). Mistrust is further exacerbated when employees do not know what information is collected and why, and do not see the benefits of such monitoring. Consent for monitoring should be requested from employees, and organizations must be aware that it may be illegal if consent is a condition of employment. Overall, employee tolerance for electronic monitoring has increased over time and is more accepted the more transparent an organization is about what information they collect.

Another issue with increased monitoring is that it affects job satisfaction even after increased pay (Holt et al., 2017). Low job satisfaction is a major cause of disgruntlement and possibly becoming an insider threat (Liang et al., 2016). It is important that an organization be transparent and fair about what is collected. This develops trust between the organization and the employee and gives the organization some feedback on possible legal issues in privacy law. Other recommendations are to ensure that all data collected is job-related, avoid keeping data no longer useful and avoid aggregating too much information about an employee. Organizations should also publish a policy about how long information is kept on terminated employees (Tursunbayeva et al., 2021).

V. RECOMMENDED METHODS AND POLICIES

We recommend closer integration of policy and technical controls to deter or detect intentional insider threats. This includes security policies that are easy for users to understand and follow, consistent vetting and termination procedures, more automated configuration options for administrators, centralized aggregation of data, and insider-threat reporting standards. Figure 1 shows the recommended hierarchy for increasing the likelihood of deterring or detecting insider threat incidents.

Figure 1. Hybrid Policy and Technical Control Hierarchy

A. POLICY IMPLEMENTATION AND ENFORCEMENT

Proper policy implementation and enforcement can address an insider threat's motive and limit the insider's opportunity to cause harm. Fair and easy-to-follow security policies positively correlate with higher employee satisfaction and increase the likelihood of following security policies (Herath & Rao, 2009). However, in one survey of health industry security personnel, 27% said their insider-threat management programs were informal or rarely enforced; 24.2% indicated they had no insider-threat management program at all (HIMSS, 2018). The lack of enforcement could also be due to a lack of

formalized policies for handling insider threat incidents. Such policies could deter insider threat activity by demonstrating that even minor incidents, such as a policy infraction, are detected and enforced fairly. Consistent enforcement of these policies must prevent risky behavior from becoming the norm, as the social influence of peers and supervisors has a significant impact on employee security behaviors (Herath & Rao, 2009).

For developing an insider-threat program, suggestions range from smaller quick-fix guides to complete programs with checklists (CERT Software Engineering Institute, 2018; Cybersecurity & Infrastructure Security Agency, 2020; NIST, 2020). One benefit of guides is that important security fundamentals are documented in writing, such as the principle of least privilege and proper access management. For smaller organizations with a limited capability to use most insider-threat tools mentioned later, a risk analysis of the most important incidents they wish to defend against and a review of tools in place to see if they can be configured may suffice. For example, existing network firewalls can be configured for content filtering to prevent proprietary data from being sent outside the organization.

Enforcing policies is easier if employees want to comply. However, a common issue is a perception that complying with security policies unnecessarily hinders worker efficiency, especially in more flexible environments where employees hold multiple roles (Herath & Rao, 2009). Ideally, a good policy would communicate how simple actions can have a significant impact. For instance, an organization should look into alternatives to complex password policies that encourage users to write down their passwords. Two-factor authentication or personalized certificates would reduce an insider's ability to steal credentials from co-workers. It could also reduce the ease of an administrator in creating fake accounts by requiring more resources to be allocated to the account.

A policy that reduces an insider's ability to commit crimes is proper termination procedures. Many insider-threat incidents occurred within 30 days of the insider leaving the organization. Some insiders showed warning signs before being terminated, and others had issues that should have been found during a pre-employment investigation. While thorough background investigations before employment are costly and time-consuming, monitoring behavior while a person is employed should be a routine security policy when

the cost of intellectual-property theft, fraud, or sabotage is significant. Monitoring around the time of termination, even if it were voluntary, is especially important to minimize risk to the organization and allow for possible prosecution. If the employer decides to provide, or is required by law to provide, a notice instead of immediate termination, remote access for that employee should be immediately disabled. In one review of 550 malicious information technology sabotage cases, 54% of the insiders did their attack using remote access tools, while 27% made their attack on-site (CERT Insider Threat Center, 2011). The most common protocols used by insiders were Secure Shell (SSH), Telnet, and Remote Desktop Protocol (RDP).

Another consideration is that for proper implementation of security policies, the Human Resources department, the legal department, and employee supervisors must coordinate with the security team. They should also be continuously evaluating employees to monitor for sudden changes that might affect an employee's likelihood of becoming an insider threat, such as major life events, repeated security violations, or unexplained affluence. Supervisors and Human Resources could reduce intentional insider threats by having formal and fair channels to resolve such work-related issues. One survey of 44 insider incidents found 34 behavioral incidents before the insider's attack (Claycomb et al., 2013). In a larger study, 97% of insiders that committed information technology sabotage had indicative behavioral incidents that supervisors or co-workers knew about before the attack (Moore et al., 2012).

B. TECHNICAL CONTROLS FOR INSIDER THREAT DETECTION OR DENIAL

Technical controls such as user-behavior analytics, privileged-access management, auditing, and centralized aggregation allow organizations to deny or detect malicious insider threats. They primarily address an insider's opportunity to commit a crime and increase the required technical capability.

1. User Behavior Analytics for Employee Monitoring

Behavioral monitoring can reduce the insider's opportunity to attack by alerting anomalous behavior, including noticing preparatory behavior common before an attack. If

monitoring is clearly communicated to employees, it would also have a deterrent effect. Also, subverting behavioral monitoring tools would require an insider to have significant technical capability to disable or otherwise subvert the tool to go unnoticed. In a review of sabotage incidents, 90.9% of the cases had observable indicators before the attack, with an average of 5.4 separate events that involved both behavioral and technical events (Claycomb et al., 2013). Employee monitoring should be conducted by trained incident response employees who and the legal department should be involved to ensure no privacy laws are being broken.

Employee monitoring includes data from system and network use. One major source is email, often used for data exfiltration when insiders steal information by sending attachments outside the organization (Gelles, 2016; Moore et al., 2012). For organizations in certain fields such as finance, housing, and sales, it is impractical to block all emails with attachments going outside of the organization; each organization must balance data protection and worker efficiency. However, indicators such as emailing outside of normal working hours or emailing many attachments should be flagged as abnormal and reviewed.

Web sites should also be monitored. Malicious sites and sites useful for data exfiltration that are not required for jobs, such as DropBox or Gmail, can be blacklisted and checked for. Some user-behavior tools can analyze a Web page's topic and recognize suspicious topics such as violent extremism.

Tracking of suspicious insider activities on workstations must go beyond anti-malware tools. A variety of host-based sensors should be used. For a confirmed suspicious insider, host-based user-monitoring capabilities could be keystroke monitoring, applications monitoring (e.g., email, chat, data export, Web browser), screen captures, records of USB port activity, and records of file editing, and host logs such as security logs (Spooner et al., 2018). This information can be collected by an application installed on each user's computer that sends the data back to a centralized location for analysis. An organization's servers can add information in logs of email, chat, printing, file transfers, and faxes. For instance, print servers log information on which user printed, which printer the job was sent, and the time the print request was sent. Indicators of anomalous behavior could be printing a large volume of sensitive material or repeated tries to access an

unauthorized file on a file server. Network-based sensors can log remote connections, email content, and Web browsing, although they may be unable to inspect encrypted content.

Table 4 summarizes the types of logs that could be collected. These logs should be sent to the centralized logging system daily for analysis. Targeted monitoring should be used for users that have been flagged by the monitoring system.

Table 4. Possible Employee Monitoring Methods. Adapted and Extended from: Spooner et al. (2018).

Type of Users	Capability	Client	Server	Network
All users	Applications executed by user	X	X[2]	X[2]
	Chat content	X	X	X[3]
	Document and file content for sensitive information	X	X	X[3]
	Email content	X	X	
	File access history	X	X	
	File editing	X	X	
	Kernel-process logging	X		
	Print logging	X	X	X
	Remote-access logging	X[1]	X	X
	Removable media activity	X		
	USB port activity	X		
	User behavioral analysis	X	X	X
	Web browser activity	X		X
	Workstation event and security logging	X		
	Logon and logoff information	X	X	X

Type of Users	Capability	Client	Server	Network
All users	Lightweight Directory Access Protocol (LDAP)		X	
Targeted monitoring	Clipboard (copy, cut, and paste) activity	X		
	Keystrokes	X		
	Screen capture of workstation display	X		

X[1]: If using company-owned equipment.

X[2]: If using Web-based application.

X[3]: Less effective if encrypted.

While collecting and analyzing all the data shown may be too costly for many organizations, a user-behavior analysis tool should try to report all obviously suspicious clues like copying many files. Not all malicious insider threats give much notice before their attack, as many suspicious clues are given less than one day before their attack (Claycomb et al., 2013). This suggests a need to integrate security staff with human resources and supervisors so that behavioral indicators can be correlated with online actions. Another recommendation is more automated options for configuring user behavior analytics tools to avoid errors. Many data-loss prevention tools offer autoconfiguration for checking compliance with policies such as the Health Insurance Portability and Accountability Act (HIPAA), Payment Card Industry Data Security Standard (PCI-DSS), or General Data Protection Regulation (GDPR). Other recommendations from the National Institute of Standards and Technology specify best security practices, but they are numerous and require careful implementation. Efforts should be made to automate this process further.

Choosing thresholds for alerts and metrics on user behavior requires balancing the organization's risk tolerance and acceptance of false positives. In establishing a threshold for a new employee, averages for the employee's role may suffice until enough data about the employee is gathered to set baseline behaviors. However, employees who perform the same role may not interact with the system in the same way. Community-behavioral

clustering allows for more tailored monitoring of an employee's behavior and also allows flexibility for smaller organizations where employees have more than one role (Dean & Rowe, 2018).

Several anomaly-based detection methods have been tested for insider-threat incident cases. Data required biased sampling because it contained many more normal than malicious users (Kim et al., 2019). Several methods showed promise using supervised, unsupervised, and semi-supervised machine-learning algorithms. Table 5 shows reported scores for various insider threat detection methods. SVM means "support-vector machine," UBA means "user-behavior analytics," CNN means "convolutional neural network," RNN means "recurrent neural network," and "LSTM" means long short-term memory recurrent neural network.

Table 5. Reported Scores for Insider-Threat Detection Methods

	Method	Recall / True Positive Rate	False Positive Rate	Precision	AUC	Study
Supervised	LTSM-CNN	-	-	-	0.94	(Yuan & Wu, 2021)
	LTSM-RNN	0.924	0.068	0.951	-	(Meng et al., 2018)
	Artificial Neural Network	0.913*	-	0.493*	-	(Le et al., 2019)
	Random Forest	0.796*	-	0.902*	-	(Le et al., 2019)
	UBA	1.0	-	0.42	-	(Legg et al., 2017)
Semi-Supervised	SVM + multinomial Naïve Bayes	0.6	-	0.3	-	(Azaria et al., 2014)
	LSTM Auto-encoder	0.91	0.098	-	-	(Sharma et al., 2020)

	Method	Recall / True Positive Rate	False Positive Rate	Precision	AUC	Study
Un-Supervised	Isolation Forest	0.72	-	-	0.76	(Gavai et al., 2015)
	LTSM-Diag.	0.95	-	-		(Tuor et al., 2017)
	Isolation Forest	0.96*	-	-	0.87*	(Aldairi et al., 2019)
	One-class SVM	0.99*	-	-	0.89*	(Aldairi et al., 2019)
	Lempel-Ziv-Welch (LZW)	0.983	0.036	-	-	(Parveen & Thuraisingham, 2012)

* Chose "daily" option for consistency, some better statistics were available for longer periods of observation.

Supervised learning had generally higher precision and recall rates with enough data; however, it cannot identify new attack vectors and requires careful feature selection. Unsupervised learning offers more flexibility because it can be used to identify new anomalies. It is well-suited for real-time analysis and can have a high true-positive rate; however, it takes longer to run and can have a high initial false-positive rate.

The most promising methods found for identifying insider threats used long short-term memory (LSTM) and deep neural network (DNN) algorithms for supervised learning. Deep neural networks can automatically extract important combinations of features. Long short-term memory and convolutional neural networks (CNN) are particularly promising. On the CMU-CERT insider threat synthetic dataset version 4.2, one study claimed they could accurately detect insider threats with an area under the receiver operating characteristic curve (AUC) to 0.9449 (F. Yuan et al., 2018). On version 6.2 of the same dataset, several studies used a recurrent neural network (RNN), a generalization of an LSTM, to capture the features to predict a user's next action (Meng et al., 2018; Tuor et al., 2017). The model performs better over time as more data is analyzed to know what is considered normal and performs similarly to the LSTM model, but they claim that the LSTM can better generalize despite sensitivity to weighted adjustments at the beginning of

training. Deep neural networks permit dynamic feature choices on highly imbalanced data. Supervised machine learning like neural networks needs labeled malicious behavior to train the model. Insider threat datasets are available, but few are publicly available, and most are synthetic, which can allow for important context-based events (i.e., the content of email messages) to be missed. Unsupervised deep learning could detect new threat methods, and the topic should be explored further. Unsupervised neural networks could recognize multivariate features automatically and learn historical trends to predict abnormal user behavior better.

2. Privileged Access Management Monitoring

System administrators that become insider threats pose more significant risks to organizations by their increased technical skill, expert knowledge of the system and monitoring capabilities, and the privileged access they already possess. Administrators and other privileged accounts should be monitored more closely than a regular employee, and there should always be more than one administrator for every critical system. Specific actions such as creating a backup, deleting a backup, and making network changes should require approval from several people. Proper separation of duties is also critical to reduce potential harm from sabotage (CERT Software Engineering Institute, 2018). Also, system administrators should not audit the systems they are responsible for.

User behavioral monitoring of privileged users should send an immediate alert from the centralized system to a manager or other system administrator on suspicious actions. All privileged actions should be logged and sent to a centralized analysis system to look for misuse. Privileged accounts should not be shared, and for temporary access, such as when a system developer needs to make an approved change, clear policies must be in place for how these instances are tracked.

3. Auditing Requirements

User access management, policy violations, and access control all require careful monitoring to identify vulnerabilities that could allow an insider-threat incident. User access management should be done at least quarterly to verify the absence of unauthorized (ghost) accounts and those of previously terminated employees. Logging should be done

automatically for account creations, changes, enabling, disabling, and removal. Especially vulnerable accounts can be recently created fake accounts, shared accounts, training accounts, and contractor accounts (CERT Software Engineering Institute, 2018; NIST, 2020).

Also, privileged access should be reviewed periodically to determine if an employee still needs that access. This applies to access controls on files and folders within an organization. Auditing should also record changes to access controls, including when file-access permissions have been delegated to non-administrators, such as a manager granting access to the accounting files to other employees. Auditing should verify that non-technical incidents are recorded in the user-behavior analysis systems.

4. Centralized Log Analysis and Correlation

Centralized logging and timely analysis are essential for managing insider-threat activity. A dedicated server must be used since the computational requirements needed to perform this analysis will be time-consuming and must be running continuously in the background. Access to the server must be strictly managed. Security-incident and event management tools can visualize different anomaly detection methods such as user and role-based activity, threshold or volume-based alerts, or new unidentified patterns and trends (Spooner et al., 2018). Key sources of information to be aggregated and analyzed are:

- User-behavior anomalies or threshold violations.
- Attempted and actual policy violations by users.
- Network anomalies such as visiting suspicious websites, making abnormal remote connections, and attempting sensitive file exfiltration.
- File access actions involving high-value assets as classified by data-loss prevention tools.
- Print-server anomalies such as high-volume printing and printing material that violates access rights.

- Abnormal event sequences where a prohibited action is followed by an allowed action (such as a prohibited attempt to copy, followed by an allowed action of printing).

Security-incident and event management tools require careful configuration and detailed rule tuning since the policy will vary between organizations. It will help to develop a "baseline" rule set for insider-threat detection analogous to a baseline rule set for intrusion detection. The system should have a user-friendly interface so that both technical and non-technical managers can use it to see high-priority alerts and take needed actions. Most security and event management systems allow customized priorities, but this capability should be audited regularly and require at least two administrators to change. Priorities can be assigned based on the severity of the action or the number of events involved. Examples of high-priority events are disabling of logging capabilities, abnormal remote access, unauthorized program download or attempted execution, attempted log modification, privilege escalation, repeated login attempts, or any significant user behavior alerts. High-priority events should require more than one person to approve the resolution. Some Automated responses can, in some cases, reduce the potential damage, such as automatic account locking or enablement of full log capture.

A centralized correlation system should correlate events over multiple timeframes. Critical information in the form of alerts and correlations should be displayed quickly to allow time to respond, as after a meeting where the insider was denied a promotion (Claycomb et al., 2013). Most instances of insider-threat activity had indicators longer than 24 hours, so a daily analysis schedule may be fine for most organizations. This allows for a granular view of changes in behavior daily and helps avoid missing key events since insiders generally will perform very few if any, malicious actions within a given day. Having only high-priority events alert immediately and other events correlate daily allows for an appropriate balance between computing requirements and timely security actions. Figure 2 summarizes our suggested system information flow.

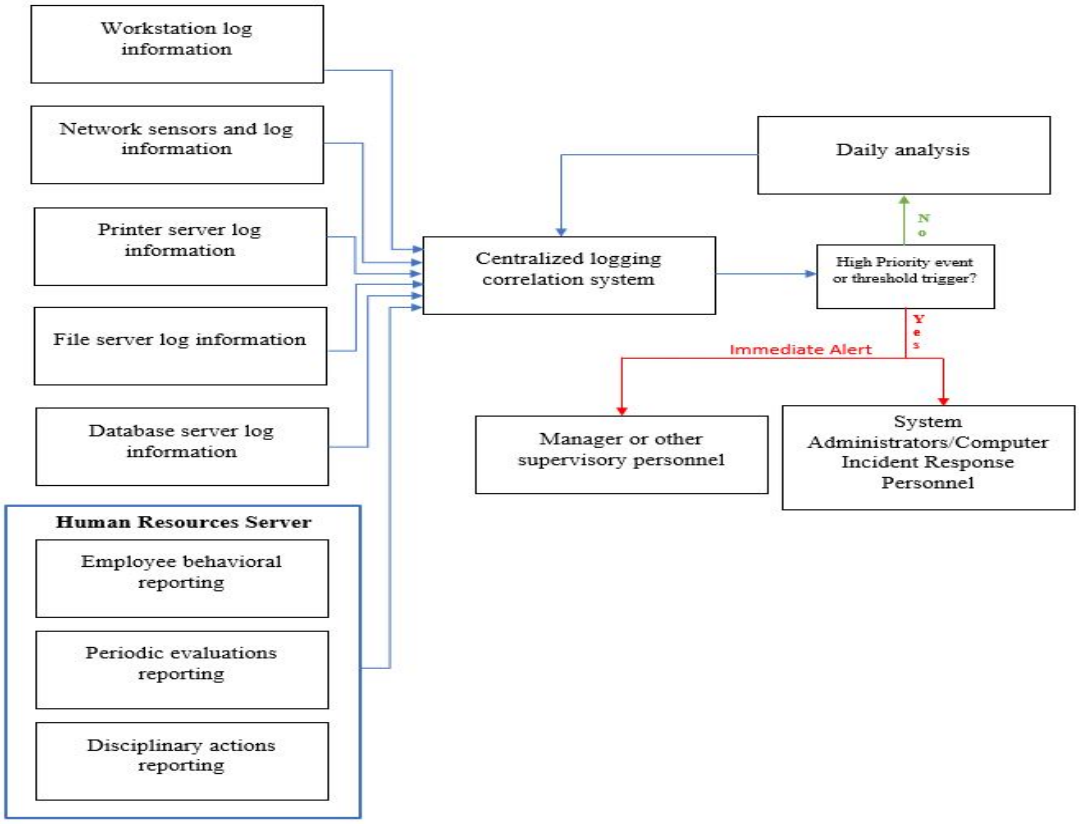

Figure 2. Proposed Centralized System Information Flow

5. Optional Controls

Additional technical controls can be added for each organization on a case-by-case basis. In most cases, small changes can address insider-threat incidents with tools already in place. For instance, for data-loss protection, methods can detect discrepancies between user assigned sensitivity levels, and data loss prevention inferred sensitivity levels (Kongsgård et al., 2017). Manual classification can be time-consuming and is prone to user-error; this method was meant to act as an indicator that employees were incorrectly classifying documents. This method used a dataset from the Digital National Security Archive to train a classifier using regression to distinguish between critical and non-critical documents with an overall accuracy of 0.84. This method also accounts for attempts at manipulating the system to classify new material improperly. An insider may still go undetected if their actions mimic normal employee actions. However, misclassification can still be a risk to an organization and should be audited for consistency by a manual review of a subset of documents. This prevents an automated

classification system from introducing errors into the pool of documents used to refine training. Also, new types of data that the data-loss prevention tool has not seen before may require manual classification with confirmation to avoid misclassification of new material. Misclassification can occur for new intellectual-property documents or by user-error. Some data-loss prevention tools can also move data stored in an improper location to the correct location and securely deleting the improperly stored version (Spooner et al., 2018). Data-loss prevention tools should be used when access control is insufficiently detailed or the cost of data leakage is high.

Honeypots or honeytokens are a novel way to identify potential insider threats that open, copy, download, email sensitive information, or try to do so. A honeypot is a computer or system designed to act as a trap for attackers to target instead of an organization's actual system. A honeytoken is similar, except instead of a computer, it can be a file or a fake login and password. Honeypots and honeytokens do require careful design so that the deceptive method is believable, easily visible, protected against sophisticated disabling, enticing to any potential insider, and not interfering with normal user activity. Common types of enticing data are login account information, financial information, social security numbers, and seemingly proprietary information. Honeytokens can act as a beacon or redirect an insider to a honeypot or honeynet with more tempting information to see if they exploit it (Spitzner, 2003). Creating and deploying honeytokens can be done automatically (Bowen et al., 2009). Software-based decoys can automatically generate bait software (Park & Stolfo, 2012). Honeytokens on user systems can help detect a masquerade attack, but this is prone to false positives with large numbers of bait files (Salem & Stolfo, 2011). Honeypot implementation does add more layers to be managed by the security team, and they may interfere with normal user activity. If such deceptions are already in place to protect against potential outside attacks, they can also help with insider-threat detection.

C. THE NEED FOR REAL-WORLD DATA

Many tools and models mentioned here were evaluated on a few synthetically created datasets and small-group discussions. More real-world data is needed on how intentional insiders are subverting controls and which controls are the best in detecting

them. A centralized agency should collect this information for both the federal and private sectors. It should be anonymized to encourage businesses and government organizations to provide data and to reduce possible privacy concerns. Reporting and sharing insider-threat data should also be made easy for organizations to do, so they will not just fire the insider and fail to offer lessons learned.

The most popular dataset used in evaluating insider-threat detection is the CMU-CERT insider-threat dataset. It was synthetically created to simulate five scenarios in a fictional company. The five scenarios were:

1. A user who did not previously use removable drives or work after hours began logging in after-hours using a removable drive and uploading data to wikileaks.org. The user leaves the organization shortly afterwards.
2. A user surfs a job website and solicits employment from a competitor. Before leaving the company, they use a thumb drive extensively to steal data.
3. A system administrator becomes disgruntled. They download a keylogger and use a thumb drive to install it on their supervisor's machine. They use the collected key logs to steal the supervisor's password and send out an alarming mass email under the supervisor’s name, causing panic in the organization. The system administrator leaves the organization immediately.
4. A user steals a password, logs into another user's machine, searches for interesting files, and emails what they find to their home email address. This behavior occurs more frequently over three months.
5. A member of a group with many layoffs uploads documents to Dropbox, planning to use them for personal gain.

Such synthetic data can be valuable in the testing of methods and tools. However, these scenarios given are well-known and smart insiders may expect the clues involved will be monitored in most organizations. Another issue with the CMU-CERT dataset is that

all email, Web pages, and other text content were randomized. This makes it impossible to discern topics of interest or do sentiment analysis which can show early indicators of disgruntlement. The data also appeared quite "clean" even when inconsistencies were artificially introduced (Glasser & Lindauer, 2013). It would be desirable to get real-world data as well for study. While privacy issues can arise with using de-identified data, more information about how a malicious insider threat behaves over time is valuable in developing models. Anonymized data of real-world incidents, though rare, could provide a different kind of data to significantly increase our understanding of insider threats.

D. CASE STUDY EXAMPLE

In April 2020, at the start of the coronavirus (COVID-19) pandemic, Christopher Dobbins disrupted shipments of critical medical personal-protective equipment. He was sentenced a few months later in July to federal prison for causing more than $200,000 in damages from destroying 2,371 electronic shipping records and editing 115,581 shipping records (U.S. Attorney Northern District of Georgia, 2020). Mr. Dobbins had administrator access to the organizations' systems and sabotaged the system after he was terminated. He used a fictitious user account he had made while still employed. After he had successfully altered or deleted shipping records, he deactivated his fake user accounts and logged out of the system. Privileged-access management tools would have caught the creation of new fake accounts and, if two-person authorization were in place, he could not have made so many fake accounts unnoticed. Also, if the organization had implemented the principle of least privilege, a regular user account should not have accessed that many shipping records. Also, a new administrator account should have been flagged. Probably the most helpful would have been increased employee monitoring before and after termination to look for anomalous activity and things like the creation of fake user accounts.

E. INSIDER THREAT REPORTING

Insider threat reporting comes from internal, external, or national sources. A recommendation for internal reporting is to use fellow employees more to notice negligent employees or intentional insiders. Many insider threats show suspicious behavior before the attack that is not detectable by technical means. However, relying upon other

employees to report suspicious behavior is difficult without a private or anonymous way to do so (Collins et al., 2016). One solution is to provide a third-party hotline to report incidents, so employees can report higher authority personnel, such as managers, without fear of reprisal. Of course, internal reporting benefits from close integration between security personnel, human resources, management, and the legal department.

Organizations should try to build trusted relationships with other organizations to share insider-threat information, building on their other information sharing (Wagner et al., 2019). Even though manual sharing is widely used, 44% of organizations surveyed reported that slow or manual processes impeded threat intelligence sharing, and 37% reported that it kept them from sharing at all (see Figure 3). It helps to automate some information sharing for cross-organizational coordination, though full automation can risk accidental spillage of protected information (Cichonski et al., 2012). Figure 3 shows some reasons why organizations chose not to share threat intelligence. While this is for all types of threat intelligence, these same issues occur for insider-threat sharing. While this is for all types of threat intelligence, similar concerns can be made when discussing insider threat sharing.

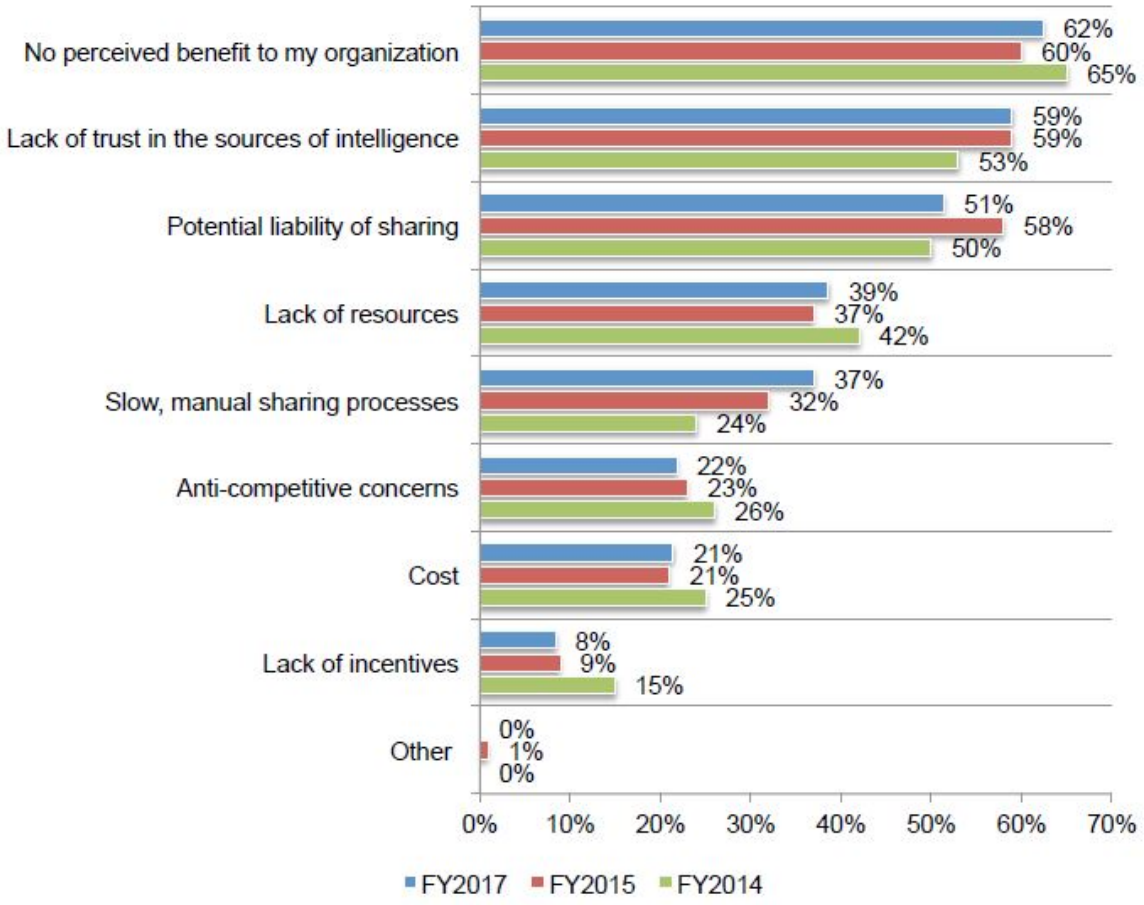

Figure 3. Reasons for Not Sharing Threat Intelligence. Source: Ponemon Institute LLC (2018).

Insider threat data sharing should also be handled at a national level. On May 12, 2021, the White House issued an executive order requiring information-technology and operational-technology service providers contracted by the federal government to share threat information (Exec. Order No. 14028, 2021). The order is meant to reduce the difficulty of sharing information between the federal government and the private sector. Similar coordination should be made for organizations that do not fall under this requirement. This would encourage the development and streamlined anonymization of real-world data and would greatly help the development of intent-based behavioral monitoring. Companies can be incentivized to share information by being able to receive intelligence threat information from other organizations, thus reducing their need for privatized solutions.

Our proposed countermeasures for insider threats are summarized in Table 6. An organization should first focus on low-cost measures that enable proper enforcement and provide a basis for more complex activities. Policy implementation is the most critical, followed by technical controls, then reporting.

Table 6. Summary of Insider Threat Enabler and Recommended Method for Mitigation or Detection

	Motive	Opportunity	Capability
Policy	✓	✓	
User behavior analytics	✓	✓	✓
Privileged-access management		✓	✓
Auditing	✓	✓	✓
Security-incident and event management		✓	✓
Threat reporting	✓	✓	✓

VI. CONCLUSION AND FUTURE WORK

The policy of deterring first then detecting is not new for countering intentional insider threats but must be planned carefully. Using ideas from deterrence models for security policy compliance with modern technological controls, the chances of being caught are high for insiders. However, deterrence effectiveness cannot be adequately evaluated without better real-world data on insider threats.

Overall, deter-then-detect should be straightforward for organizations to implement because most deterrence can be accomplished through a few monitoring capabilities and policy enforcement. The challenges are primarily in security implementation. Insider threats are a rising problem, with 68% of organizations surveyed by Ponemon Institute saying they observed more frequent insider attacks over the last 12 months; the number of incidents reported has tripled since 2016 (IBM Security, 2020). The average annual cost of handling malicious insiders or credential theft was $6.87 million per organization in 204 organizations surveyed across 13 industry sectors. Robust policies and modern technical controls are definitely recommended.

We also need more studies on intentional insider threats. Hundreds of studies are available on handling insider threats, but they focus either on the psychological side or the technical side alone. We noted a lack of discussion about whether insider threats are increasing enough to need increased investment. However, we believe that many tools discussed here would help protect an organization against insider threats.

Future work on insider threats should focus on improving detection tools using real-world data of insider-threat incidents, particularly long short-term memory models. Additional work could focus on comparing the different LSTM models against each other to find the most efficient model when comparing overall recall and precision against the number of extracted features required. Another recommendation is to focus more on combining psychological and technical techniques. A combined study with representatives from both disciplines developing tools and studying how they influence user behavior would help.

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

Ahmad, A., Hadgkiss, J., & Ruighaver, A. B. (2012). Incident response teams—Challenges in supporting the organizational security function. *Computers & Security*, *31*(5), 643–652. https://doi.org/10.1016/j.cose.2012.04.001

Alba, E. (2015, December 3). Explanation of INSA-developed insider threat definition. *INSA*. https://www.insaonline.org/explanation-of-insa-insider-threat-definition/

Albrechtsen, E. (2007). A qualitative study of users' view on information security. *Computers & Security*, *26*(4), 276–289. https://doi.org/10.1016/j.cose.2006.11.004

Aldairi, M., Karimi, L., & Joshi, J. (2019). A Trust aware unsupervised learning approach for insider threat detection. *2019 IEEE 20th International Conference on Information Reuse and Integration for Data Science (IRI)*, 89–98. https://doi.org/10.1109/IRI.2019.00027

Appel, E. J. (2017). *Internet searches for vetting, investigations, and open-source intelligence* (2nd ed.). CRC Press. https://doi.org/10.1201/b10523

Azaria, A., Richardson, A., Kraus, S., & Subrahmanian, V. S. (2014). Behavioral analysis of insider threat: A survey and bootstrapped prediction in Imbalanced data. *IEEE Transactions on Computational Social Systems*, *1*(2), 135–155. https://doi.org/10.1109/TCSS.2014.2377811

Bodie, M., Cherry, M., McCormick, M., & Tang, J. (2016). The law and policy of people analytics. All Faculty Scholarship. https://scholarship.law.slu.edu/faculty/3

Bowen, B. M., Hershkop, S., Keromytis, A. D., & Stolfo, S. J. (2009). Baiting inside attackers using decoy documents. In Y. Chen, T. D. Dimitriou, & J. Zhou (Eds.), *Security and Privacy in Communication Networks* (pp. 51–70). Springer. https://doi.org/10.1007/978-3-642-05284-2_4

Catrantzos, N. (2012). *Managing the insider threat: No dark corners*. CRC Press.

CDSE (2019). *Potential risk indicators: Insider threat*. Center for Development of Security Excellence. https://www.cdse.edu/documents/toolkits-insider/INTJ0181-insider-threat-indicators-job-aid.pdf

CERT Insider Threat Center (2010). *Insider threat deep dive: IT sabotage*. https://insights.sei.cmu.edu/insider-threat/2010/09/insider-threat-deep-dive-it-sabotage.html

CERT Insider Threat Center (2011). *Insider threat control: Using a SIEM signature to detect potential precursors to IT sabotage*. https://apps.dtic.mil/sti/pdfs/ADA636508.pdf

CERT Software Engineering Institute (2018). *Common sense guide to mitigating insider threats* (Sixth). Software Engineering Institute. https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=540644

CERT Software Engineering Institute (2012). *Insider fraud in financial services* [Brochure]. https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=28204

Charney, D. L. (2010). True psychology of the insider spy. *The Intelligencer: Journal of U.S. Intelligence Studies*, *22*(Fall/Winter 2010), 8.

Charney, D. L., & Irvin, J. A. (2016). The Psychology of espionage. *The Intelligencer: Journal of U.S. Intelligence Studies*, *22*(1), 7.

Cichonski, P., Millar, T., Grance, T., & Scarfone, K. (2012). *Computer security incident handling guide: Recommendations of the National Institute of Standards and Technology* (NIST SP 800-61r2). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-61r2

Claycomb, W. R., Huth, C. L., Phillips, B., Flynn, L., & McIntire, D. (2013). Identifying indicators of insider threats: Insider IT sabotage. *2013 47th International Carnahan Conference on Security Technology (ICCST)*, 1–5. https://doi.org/10.1109/CCST.2013.6922038

Collins, M., Theis, M., Trzeciak, R., Strozer, J., Clark, J., Costa, D., Cassidy, T., Albrethsen, M., & Moore, A. (2016). *Common sense guide to mitigating insider threats* (Fifth). Software Engineering Institute. https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=484738

Cybersecurity & Infrastructure Security Agency (2020). *Insider threat mitigation guide*. https://www.cisa.gov/sites/default/files/publications/Insider%20Threat%20Mitigation%20Guide_Final_508.pdf

Dean, J. (2017). *Systematic assessment of the impact of user roles on network flow patterns* [Doctoral dissertation, Naval Postgraduate School]. http://hdl.handle.net/10945/56119

Dean, J., & Rowe, N. (2018, December). *Utility of user roles in comparing network flow behaviors*. 2018 *International Conference on Computational Science and Computational Intelligence (CSCI)* (pp.42-47), https://doi.org/10.1109/CSCI46756.2018.00016

Exec. Order No. 14,028, 3 C.F.R. 26647 (2021).

Gavai, G., Sricharan, K., Gunning, D., Hanley, J., Singhal, M., & Rolleston, R. (2015). Supervised and unsupervised methods to detect insider threat from enterprise social and online activity data. *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications*, *6*(4), 47–63.

Gelles, M. (2016). Insider threat: Prevention, detection, mitigation, and deterrence. Butterworth-Heinemann is an imprint of Elsevier.

George, A., Arey, B., Ertzinger, B., Michaelson, B., Heck, D., Johnson, D., Demmon, J., Speciale, J., & Smith, K. (2019). *Best practices in vetting prospective and current employees.* Public-Private Analytic Exchange Program. https://www.dhs.gov/sites/default/files/publications/ia/ia_best-practices-vetting-prospective-current-employees-v2.pdf

Gheyas, I. A., & Abdallah, A. E. (2016). Detection and prediction of insider threats to cyber security: A systematic literature review and meta-analysis. *Big Data Analytics*, *1*(1), 6. https://doi.org/10.1186/s41044-016-0006-0

Glasser, J., & Lindauer, B. (2013). Bridging the gap: A pragmatic approach to generating insider threat data. *2013 IEEE Security and Privacy Workshops*, (p. 98–104). https://doi.org/10.1109/SPW.2013.37

Hart, M., Manadhata, P., & Johnson, R. (2011). Text classification for data loss prevention. In *Privacy Enhancing Technologies* (Vol. 6794, pp. 18–37). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-22263-4

Herath, T., & Rao, H. R. (2009). Protection motivation and deterrence: A framework for security policy compliance in organizations. *European Journal of Information Systems*, *18*(2), (p. 106–125). https://doi.org/10.1057/ejis.2009.6

HIMSS (2018). *2018 HIMSS cybersecurity survey*. HIMSS. https://www.himss.org/sites/hde/files/d7/u132196/2018_HIMSS_Cybersecurity_Survey_Final_Report.pdf

Holt, M., Lang, B., & Sutton, S.G. (2017), Potential employees' ethical perceptions of active monitoring: the dark side of data analytics, *Journal of Information Systems*, *31*(2), pp. 107–124.

Intelligence and National Security Alliance (2019). *Components of effective insider threat training*. https://www.insaonline.org/wp-content/uploads/2019/10/INSA_WP_Training-Programs.pdf

IBM Security (2020). *Cost of insider threats global report 2020*. Ponemon Institute. https://www.observeit.com/cost-of-insider-threats/

Kim, A., Oh, J., Ryu, J., Lee, J., Kwon, K., & Lee, K. (2019). SoK: A Systematic review of insider threat detection. *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications*, *10*(4), 46–67. https://doi.org/10.22667/JOWUA.2019.12.31.046

Kongsgård, K. W., Nordbotten, N. A., Mancini, F., & Engelstad, P. E. (2016). Data loss prevention based on text classification in controlled environments. In *Information Systems Security:* (Vol. 10063, pp. 131–150). Springer International Publishing.

Kongsgård, K. W., Nordbotten, N. A., Mancini, F., & Engelstad, P. E. (2017). An Internal/Insider threat score for data loss prevention and detection. Proceedings of the *3rd ACM on International Workshop on Security and Privacy Analytics*, (p. 11–16). https://doi.org/10.1145/3041008.3041011

Le, D. C., & Zincir-Heywood, N. A. (2019). Machine learning based insider threat modelling and detection. *2019 IFIP/IEEE Symposium on Integrated Network and Service Management (IM)*, 1–6.

Legg, P. A., Buckley, O., Goldsmith, M., & Creese, S. (2017). Automated insider threat detection system using user and role-based profile assessment. *IEEE Systems Journal*, *11*(2), (p. 503–512). https://doi.org/10.1109/JSYST.2015.2438442

Levine, J., LaBella, R., Owen, H., Contis, D., & Culver, B. (2003). *The use of Honeynets to detect exploited systems across large enterprise networks*. IEEE Systems, *Man and Cybernetics Society Information Assurance Workshop*. West Point, NY. https://doi.org/10.1109/SMCSIA.2003.1232406

Liang, N. (Peter), Biros, D. P., & Luse, A. (2016). An empirical validation of malicious insider characteristics. *Journal of Management Information Systems*, *33*(2), (p. 361–392). https://doi.org/10.1080/07421222.2016.1205925

Luckey, D., Stebbins, D., Orrie, R., Rebhan, E., Bhatt, S., & Beaghley, S. (2019). *Assessing continuous evaluation approaches for insider threats: How can the security posture of the U.S. Departments and agencies be improved?* RAND Corporation. https://doi.org/10.7249/RR2684

Meng, F., Lou, F., Fu, Y., & Tian, Z. (2018). Deep learning based attribute classification insider threat detection for data security. *2018 IEEE Third International Conference on Data Science in Cyberspace (DSC)*, (p. 576–581). https://doi.org/10.1109/DSC.2018.00092

Moore, A. P., Hanley, M., & Mundie, D. (2012). A pattern for increased monitoring for intellectual property theft by departing insiders. *Software Engineering Institute, Carnegie Mellon University*, *Report Number: CMU/SEI-2012-TR-008*, 10. https://doi.org/10.1184/R1/6571703.v1

NIST (2020). *Security and privacy controls for information systems and organizations* (NIST Special Publication (SP) 800-53 Rev. 5). National Institute of Standards and Technology. https://doi.org/10.6028/NIST.SP.800-53r5

Park, Y., & Stolfo, S. J. (2012). Software decoys for insider threat. *Proceedings of the 7th ACM Symposium on Information, Computer and Communications Security - ASIACCS '12*, 93. https://doi.org/10.1145/2414456.2414511

Parveen, P., & Thuraisingham, B. (2012). Unsupervised incremental sequence learning for insider threat detection. *2012 IEEE International Conference on Intelligence and Security Informatics*, (p. 141–143). https://doi.org/10.1109/ISI.2012.6284271

Pogarsky, G., & Piquero, A. R. (2003). Can punishment encourage offending? Investigating the "resetting" effect. *Journal of Research in Crime and Delinquency*, *40*(1), 95–120. https://doi.org/10.1177/0022427802239255

Ponemon Institute LLC (2018). *Third annual study on exchanging cyber threat intelligence: There has to be a better way* (p. 47). Ponemon Institute. https://www.infoblox.com/wp-content/uploads/infoblox-white-paper-ponemon-infoblox-2018-final-report.pdf

Pratt, T. C., Cullen, F. T., Blevins, K. R., Daigle, L. E., & Madensen, T. D. (2006). The Empirical status of deterrence theory: A Meta-analysis. In *Taking stock: The status of criminological theory* (pp. 367–395). Transaction Publishers.

Salem, M., & Stolfo, S. J. (2011). Decoy document deployment for effective masquerade attack detection. In T. Holz & H. Bos (Eds.), *Detection of Intrusions and Malware, and Vulnerability Assessment* (pp. 35–54). Springer. https://doi.org/10.1007/978-3-642-22424-9_3

Schultz, E. E. (2002). A framework for understanding and predicting insider attacks. *Computers & Security*, *21*(6), 526–531. https://doi.org/10.1016/S0167-4048(02)01009-X

Shabtai, A., Elovici, Y., & Rokach, L. (2012). *A survey of data leakage detection and prevention solutions*. Springer U.S. https://doi.org/10.1007/978-1-4614-2053-8

Sharma, B., Pokharel, P., & Joshi, B. (2020). User behavior analytics for anomaly detection using LSTM autoencoder—Insider threat detection. *Proceedings of the 11th International Conference on Advances in Information Technology*, 1–9. https://doi.org/10.1145/3406601.3406610

Spitzner, L. (2003). *Honeypots: Catching the insider threat*. 19th Annual Computer Security Applications Conference. Las Vegas, NV. https://doi.org/10.1109/CSAC.2003.1254322

Spooner, D., Silowash, G., Costa, D., & Albrethsen, M. (2018). Navigating the insider threat tool landscape: Low cost technical solutions to jump start an insider threat program. *2018 IEEE Security and Privacy Workshops (SPW)*, 247–257. https://doi.org/10.1109/SPW.2018.00040

Stolfo, S., Bellovin, S., Hershkop, S., Keromytis, A., Sinclair, S., & Smith, S. (Eds.). (2008). *Insider attack and cyber security* (Vol. 39). Springer Science+Business Media, LLC.

Thompson, E. (2019). *The insider threat: Assessment and mitigation of risks*. CRC Press Taylor & Francis Group.

Tuor, A., Kaplan, S., Hutchinson, B., Nichols, N., & Robinson, S. (2017). Deep learning for unsupervised insider threat detection in structured cybersecurity data streams. *ArXiv:1710.00811 [Cs, Stat]*. http://arxiv.org/abs/1710.00811

Tursunbayeva, A., Pagliari, C., Di Lauro, S., & Antonelli, G. (2021). The ethics of people analytics: Risks, opportunities and recommendations. *Personnel Review*. https://doi.org/10.1108/PR-12-2019-0680

U.S. Attorney Northern District of Georgia (2020, October 20). *Former employee of medical packaging company sentenced to federal prison for disrupting PPE shipments*. https://www.justice.gov/usao-ndga/pr/former-employee-medical-packaging-company-sentenced-federal-prison-disrupting-ppe

Wagner, T. D., Mahbub, K., Palomar, E., & Abdallah, A. E. (2019). Cyber threat intelligence sharing: Survey and research directions. *Computers & Security*, *87*, 13. https://doi.org/10.1016/j.cose.2019.101589

Youyou, W., Kosinski, M., & Stillwell, D. (2015). Computer-based personality judgments are more accurate than those made by humans. *Proceedings of the National Academy of Sciences*, *112*(4), 1036–1040. https://doi.org/10.1073/pnas.1418680112

Yuan, F., Cao, Y., Shang, Y., Liu, Y., Tan, J., & Fang, B. (2018). Insider threat detection with deep neural network. In Y. Shi, H. Fu, Y. Tian, V. V. Krzhizhanovskaya, M. H. Lees, J. Dongarra, & P. M. A. Sloot (Eds.), *Computational Science – ICCS 2018* (pp. 43–54). Springer International Publishing. https://doi.org/10.1007/978-3-319-93698-7_4

Yuan, S., & Wu, X. (2021). Deep learning for insider threat detection: Review, challenges and opportunities. *Computers & Security*, *104*, 102221. https://doi.org/10.1016/j.cose.2021.102221

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center
   Ft. Belvoir, Virginia

2. Dudley Knox Library
   Naval Postgraduate School
   Monterey, California