ANALYTIC APPROACHES TO DETECT INSIDER THREATS

DECEMBER 9, 2015

TABLE OF CONTENTS

EXECUTIVE SUMMARY

All organizations face security risks. With the growth of information technology-enabled infrastructure, these risks are manifested in the cyber domain. To detect and mitigate the risks, organizations rely on continuous security assessment and monitoring programs. These programs must be conducted in compliance with applicable laws and the organization's ethical, and privacy policies.

Of these security risks, some estimates show that over 50% are posed by insiders—individuals with access to organizational resources. This whitepaper identifies steps that organizations may use to enhance their security posture to detect potential insider threats. In many cases, this detection can be done using existing organizational security infrastructure that leverages modern network architectures. Similar to the rest of the security infrastructure, the whitepaper reminds organizations that insider threat capabilities must operate within an appropriate legal, ethical, and privacy framework and the techniques proposed within this whitepaper should be tailored accordingly.

The whitepaper expands upon published insider threat agent attack research[1] by providing analytic indicators[2] for early detection. It is important to note that an individual analytic[3] by itself is neither a definitive indicator of an attack nor sufficient to distinguish between attack types. The white paper also identifies the data required for those analytics to operate. The whitepaper presents a sample system architecture that illustrates the infrastructure components and data they provide. Then, the whitepaper discusses modern "big data" architectures that are capable of capturing and managing the data volumes from these components, and making that data accessible to streaming and batch analytic tools which power the insider threat analytics. To reduce implementation costs, the whitepaper focuses on leveraging tools that typically exist within an organization's security infrastructure and identifies additional classes of automated tools that can facilitate the integration of analytics.

The presentation of this material is structured in a manner that facilitates organizational tailoring of the guidance based upon information technology limitations, legal authorities, corporate policies, business concerns, and workplace culture. In addition, all of this material is aligned with the following five core recommendations of the whitepaper:

1. Implement an insider threat program to provide an integrated approach to addressing insider-based risks within an appropriate legal, ethical, and policy framework to ensure privacy-protections.

---

[1] Research sources including those in the bibliography refer to "attacks" as behaviors or activity that can cause damage regardless of the intent of the threat agent, a person who accidentally or maliciously takes steps to cause harm, or the type of potential damage. This whitepaper uses the term "attack" in this sense.

[2] **Analytic indicator -** analytics' output that suggests the presence of an insider threat; may prompt decision making e.g., further analysis, analytic refinement, legal response.

[3] **Analytic -** automated process run against data to identify meaningful patterns or relationships in the data.

2. Deploy a continuous assessment capability as part of a well-governed and securely-operated insider threat program.
3. Deploy analytics to discover potential insider threats; focus detection on the organization's most valued assets**.**
4. Provide investigative tools to help analysts and management correlate the indicators, understand the observed activity, and determine if it is a false positive.
5. Facilitate attribution of individuals through a comprehensive identity management system for individuals.

A. INTRODUCTION

1. BACKGROUND

In a recent survey by Forrester Research (Shey, Mak, Balaouras, & Luu, 2013), 2,134 Information Technology (IT) executives and technology decision makers from around the globe were surveyed about the current state of security and privacy. When asked what the most common cause of a breach was in the last 12 months, most respondents (36%) identified inadvertent misuse by an insider, and another 25% indicated that breaches were caused by a malicious insider. One 2015 survey estimates the overall cost to an organization to remediate one successful insider attack is $445,000. Given an average of 3.8 successful insider attacks per year, the annual cost to an organization can reach $1.7 million (Schulze, 2015). These insiders have easier access to information, systems, and physical facilities when compared with outside threats, and, often, insiders can have strong motives for abusing this access to benefit themselves or cause harm to an organization.

For the purposes of this whitepaper, insider threat is defined as:

> ***Insider threat is the potential for a current or former employee, contractor, or business partner to accidentally or maliciously misuse their trusted access to harm the organization's employees, customers, assets, reputation, or interests.***

Within the whitepaper, this definition is used to include a number of insider threat types, consider the behaviors or activity that can cause damage associated with each threat type, and identify the analytics and data requirements to detect these behaviors. This decomposition allows an organization to focus on those threat types of concern to its operations, within the legal and policy framework under which it operates. Note that within this whitepaper, a person who accidentally or maliciously takes steps to cause harm is referred to as an agent, a behavior or activity that can cause damage is referred to as an attack, and an automated process run against data to identify meaningful patterns or relationships in the data is referred to as an analytic.

Furthermore, this whitepaper defines an insider threat *program* as a concerted effort by an organization to detect insider threats and respond to insider attacks. Insider threat analysts use information from multiple sources to put user behaviors and activities into context and determine if damage to an organization is likely. Based on this analysis, and consideration of policy, legal, ethical, privacy, and other factors, the organization might pursue a variety of responses. An insider threat program can be implemented via external, internal, or manual processes, or some combination thereof.

Many organizations do not have an insider threat program, but the need for one has never been more apparent. When building an insider threat program, it is critical for organizations to engage stakeholders, such as senior management, legal, and human

resources, from the program's inception to implementation and refinement. Also, numerous online resources are available to assist. For example, the CERT® Insider Threat Center at the Carnegie Mellon University Software Engineering Institute (SEI) (CERT Division) and the CERT® Program's *Common Sense Guide to Mitigating Insider Threats* (Silowash, et al., 2012) are good starting points.

2. OUTLINE

This whitepaper provides suggestions for security programs regarding continuous assessment and monitoring to detect potential insider threats based on assumptions about the capability of an organization's Information Technology (IT) system (Appendix B). For reasonable efficiency, this monitoring requires automated analytics based upon data gathered from systems and the security infrastructure. Specifically, this whitepaper will:

- Present the policy, privacy, ethical, legal, and cost considerations in the context of a high-level model for insider threat programs (Section **B**);
- Expand upon current literature defining insider threat agents and their associated attack types (Section **C**);
- Present the state of the art and propose advances in current strategies and technologies to provide analysts with an improved threat detection capability (Section **D**);
- Describe the analytic process and investigation of potential insiders (Section **0**);
- Identify how modern architectures can enable the collection of data and invocation of big-data analytics to detect insider threats (Section **F**); and
- Provide recommendations on how to use these technologies in the context of a comprehensive insider threat program (Section G).

This whitepaper presents the findings in a manner that can be adapted to the needs of both small and large organizations by taking into account applicable national laws, the laws of countries and localities in which they do business, as well as corporate policies, business concerns, and workplace culture.

The effective detection of insider threats and events, especially in cyber domains, is an emerging discipline. The intent of this whitepaper is to bring together many sources to comprehensively describe the current state of the art. It draws on research and case studies where available, as well as the judgment and hands-on experience of many experts from industry, academia, and government. It is acknowledged that much research remains to be done, and that this whitepaper is neither exhaustive nor the final reference. However, in addition to supporting insider threat programs today, this whitepaper can also provide a solid starting point for future discussion and research needed to mature the art and science of insider risk management.

B. INSIDER THREAT PROGRAM OVERVIEW

1. INTRODUCTION

The success of an insider threat capability depends on people and processes as much as technology. Employee education and awareness of the damage that can be done due to insider threats, as well as security analyst and investigator knowledge of the organization's mission, culture, and relationships with its customers, employees, and society are necessary to ensure that an insider threat program is effective. Additionally, transparency regarding an insider threat program may help establish both a legal foundation for the program and allow an insider threat program to be conducted without adversely affecting employer/employee relations, privacy, and civil liberties.

Upon this basis, the insider threat program can gather the necessary data; detect patterns and behaviors, and instantiate responses for items of concern. Depending upon the threat, responses can include using automated IT reconfiguration, referring the case to law enforcement, providing human resources intervention, or providing training to employees. Recommendations in the following materials focus on how to build the analytics and data architecture to make sense of the data and inform an appropriate response, both integral to achieving a core insider threat capability: The CERT Guide to Insider Threat (Cappelli, Moore, & Trzeciak, 2012), The Common Sense Guide to Mitigating Insider Threat (Silowash, et al., 2012), Insider Threat Program Best Practices. 2013 46th Hawaii International Conference on System Sciences (Guido & Brooks, 2013), etc.

Any insider threat program should establish and adhere to guiding policies and principles related to data collection, handling, and access control. First, when establishing data handling, processing, and storage capabilities, protection mechanisms must be put in place to address privacy and civil liberties concerns. Second, authorized individuals with access to sensitive data should be well trained, and controls should be put in place to ensure they do not abuse their privileges. Third, to limit the deployment of expensive investigative resources, organizations should use as many relevant, independent analytics for each attack type as is feasible. Fourth, all analytics and associated data must tie back to an individual's identity and/or patterns to ensure that a narrow focus is maintained by the insider threat analysts. Lastly, the insider threat program must provide the governance and training necessary to execute courses of action for positive indicators, including addressing potential false positives, ensuring access control and integrity protections for insider threat information, and identifying investigative processes.

2. POLICY, PRIVACY, AND ETHICAL CONSIDERATIONS

Unlike other types of computer security programs, which typically target malware threats or internal computer systems, insider threat programs focus on people, which can raise a host of policy, privacy and ethical concerns. An effective insider threat program must balance a variety of interests, including the protection of an organization's proprietary, sensitive, and classified assets, as well as the preservation of customer and employee

privacy and civil liberties. While protecting an organization's data and assets is a central goal of an insider threat program, an organization that does not adequately protect employee data or uses it in ways that employees have not authorized, risks losing employee trust or facing litigation. The relationship between an organization and its employees should include transparency about business practices that result in the use and disclosure of employee information, particularly where it involves others outside of the organization, such as the government. Any enhancements to an insider threat program must be made in a way that maintains trust between employees and the employer, between the organization and its customers, and between the organization and the public.

Transparency may also be required to lawfully operate an insider threat program. For example, employee consent and notice, e.g., user agreement and login banners, can be a critical component to conducting an insider threat program consistent with international norms and domestic laws of any country in which the organization operates. In the case of organizations handling data outside the United States or data belonging to employees in other countries, consent alone may not be sufficient.

Transparency can also be critical for other reasons, such as maintaining employee and public relations and adherence to an organization's code of conduct. Monitoring systems that are unduly intrusive or indiscriminate can also have negative business, political, or legal consequences. The results can include disgruntlement that actually increases the insider threat, or reduction of employee engagement and loyalty. Additionally, different cultures and regions will have different expectations regarding the relationship between an organization and employees, and these expectations can impact how an organization will need to implement its insider threat program. This is particularly germane to multinational organizations.

3. LEGAL CONSIDERATIONS

An effective insider threat program relies on data drawn from multiple sources, data that in some cases is protected by U.S. law and the laws of other countries in which the organization does business. For instance, in the U.S., federal and state statutes protect personnel files and other human resources information, healthcare data, and intercepted electronic communications. An organization that handles the data of foreign employees or has offices abroad, it may also be subject to foreign data privacy protections, e.g., the European Union Privacy Directive. Furthermore, the use of some information may be constrained by an organization's own internal policies or contracts, e.g., employee agreements concerning the confidentiality of their information. In addition, legal issues, such as monitoring of personal devices that employers allow to be used in the workplace, are emblematic of the new complex legal issues faced by organizations implementing insider threat programs.

In many instances, legal and policy restrictions allow the use of protected data in an insider threat program within certain parameters. Safely navigating the various legal and policy restrictions requires close engagement with the organization's legal counsel. Furthermore, the assistance of privacy counsel can be critical to a multinational organization that must

comply both with U.S. laws concerning the collection, use, and disclosure of information and with international standards and rules on privacy and civil liberties in the various international jurisdictions in which it operates. Although not the main focus of this whitepaper, considerations from a legal and policy perspective are provided in detail in *Appendix A: Legal Considerations for Insider Threat Monitoring Programs.*

4. COST CONSIDERATIONS

Each organization must do its own cost-benefit analysis based upon its own risk profile and the organization's risk tolerance. In small organizations, limited resources might make it necessary to integrate insider threat analytic capabilities within the existing Security Information and Event Management (SIEM) system and security operations center (SOC) architecture. For larger organizations, other constraints, such as policy and legal considerations, might require segregation of insider threat analytics and data from network defense operations.

A goal for any organization should be to balance the various costs and expected benefits of early detection. Costs include infrastructure expenses associated with building enhanced insider threat detection, costs imposed by regulatory or other external authorities, and costs associated with protecting privacy-sensitive data. Benefits include preventing loss of intellectual property, avoiding brand damage and litigation, or even improving the defensive posture against external attacks. An organization should focus on those benefits that are significant to the organization i.e., the return on investment from detecting individuals stealing office supplies may not warrant the cost.

C. INSIDER THREAT AGENT AND ATTACK TYPES

This whitepaper's broad definition of insider threat includes agents with a variety of motives, intent, desired effects, and levels of negligence; however, the important consideration is the damage, or outcome, these agents can cause. This broad definition can be refined to identify classes of analytics that can best detect an agent's presence within an enterprise. The refinement or breakdown of the definitions summarized in **Figure 1** is based on (Casey, Insider Threat Field Guide, 2015). Both hostile and non-hostile agent types are identified and mapped to the attack vectors they are most likely to employ (indicated in the table by an "X"). [For definitions of insider agent types, see Threat Agent Library Helps Identify Information Security Risk (Casey, 2007).] Not all insider threats will fall cleanly into one of these insider agent types; but, such decomposition helps organizations to determine the scope and capabilities of their insider threat program. Doing so allows organizations to focus resources on specific agent types based on business risk and legal and policy considerations. The analytic recommendations for insider threat detection focus on the attack vectors. Note that these threat-based attack types are listed alphabetically rather than by priority because each organization can and should prioritize them according to their unique environment.

		Insider Agent Types												
	Intent→	Non-Hostile			Multiple		Hostile							
		Reckless Individual	Untrained / Distracted Individual	Outward Sympathizer	Vendor	Partner	Irrational Individual	Thief	Disgruntled Individual	Activist	Terrorist	Organized Crime	Competitor	Nation State
Attack Types	Accidental Leak	X	X	X	X	X	X		X					
	Espionage				X	X		X	X	X		X	X	X
	Financial Fraud				X	X		X	X			X		
	Misuse	X	X	X	X	X	X		X	X				
	Opportunistic Data Theft				X	X		X	X	X		X	X	X
	Physical Theft						X	X	X		X	X		
	Product Alteration	X	X		X	X			X	X		X	X	X
	Sabotage						X		X	X	X		X	X
	Violence						X		X		X			

Figure 1: *Insider threat agents and their associated attack types*

To further elaborate on each of these attack types, notional examples of each are in the following paragraphs.

Accidental Leak: Phillip's department has been hit hard with layoffs and the remaining employees now have to rapidly take on new duties, sometimes without adequate preparation. Overwhelmed by the new tasks, Phillip posts internal specification documents to the company's vendor information site hoping to preemptively answer vendor questions, never realizing the site is publicly accessible and company intellectual property is now widely exposed and quickly indexed by internet search engines.

Espionage: Esteban is a star Ph.D. candidate in material sciences at a top university and conducts his research there in conjunction with several major research companies. He is best friends with Pat, who often helps Esteban with his heavy school expenses. In return, Esteban often shares the files on his "way-cool discoveries" with Pat, despite the school's policies and non-disclosure agreements (NDA). Pat actually works for a nation-state and uses the files to steal leading-edge technology from the research. They continue their contact and sharing after Esteban graduates and gains employment at one of the research companies.

Financial Fraud: Marie is a long-time and trusted employee in the Procurement office. With the help of a cohort, she uses company accounts to pay faked electronic invoices in the names of former vendors, and the money is actually directed into their personal accounts. She knows the old vendor accounts are never purged from the system, so the invoices appear legitimate and will not raise suspicion. Over several years, Marie and her accomplice siphon off a substantial portion of the company's profits.

Misuse: Jon is an immigrant from a region experiencing frequent violent conflicts, and his family still lives in the area. When a large conflict flares, he feels he must help protect his family, but the distance prevents direct involvement. Instead, he installs a hacker toolkit onto some of the company servers and uses it to cyber-attack the country opposing his own. That country detects the attack and not only retaliates by cyber means, but also seizes the company's local offices.

Opportunistic Data Theft: Koharu is a biologist who has just landed a position with a prestigious bio-engineering firm. While at her current company, she developed several valuable technologies and as the developer she thinks of those technologies as her own property. Before she leaves her current company, Koharu downloads key files on those technologies still accessible to her and takes them with her to facilitate a quick start at her new company, but also giving the new company an unfair competitive advantage.

Physical Theft: Tom has run up large gambling debts with a local crime syndicate and cannot pay them back. To erase the debt, he agrees to help the syndicate steal shipments of his company's secret, high-value hardware prototypes. He uses his manager-level network access to find the manifests and shipping schedules and relays the information to the syndicate, which then easily hijacks the shipments in route.

Product Alteration: Adrien works for a software engineering firm as a software configuration manager. He uses his position to clandestinely add botnet malware to the company's financial products, and he manages the profitable botnet from his home. The botnet is eventually discovered, and the company is revealed as the source of the malware. Company operations are severely disrupted for many months dealing with both the public relations issues and the criminal investigations.

Sabotage: Chuck is passed over for a promotion for the third time at the electrical power company he works for. He feels cheated and resentful and decides to get revenge for what he believes is poor treatment by his company. He purchases a hacker toolkit that enables him, over several months, to build and quietly install software “time bombs” that will erase the hard drives and memory of every company computer it can reach, including the industrial control systems he works on. When the time bombs activate during a morning rush hour, the company goes offline for the several hours it takes to recover from the massive outage, severely impacting everyone in the surrounding area.

Violence: Mathew recently got a new supervisor at work and does not like him. He stops completing assignments on time and openly voices his dislike of the supervisor. After several months of increasing hostility, the supervisor gives Mathew a written warning and

places him on probation. This greatly angers Mathew, and he threatens loudly to "get even with" the supervisor and his peers after work.

D. Analytic Indicators

1. Context

To mitigate the threat agents described above, this whitepaper will discuss the analytics used to indicate specific attack types in the next section. First, it is important to understand the broader context associated with insider threat detection. This context takes into account: data and information needed for analysis; availability of said data and information; decision to deploy analytics; analysis findings which may indicate an attack or the need for further collection and analysis; and appropriate responses. These factors drive and inform decisions that should be governed by an organization's leadership and must be underpinned by the privacy, civil liberties, ethics, legal and policy considerations of a well-managed insider threat program. This insider threat context is depicted in **Figure 2.**

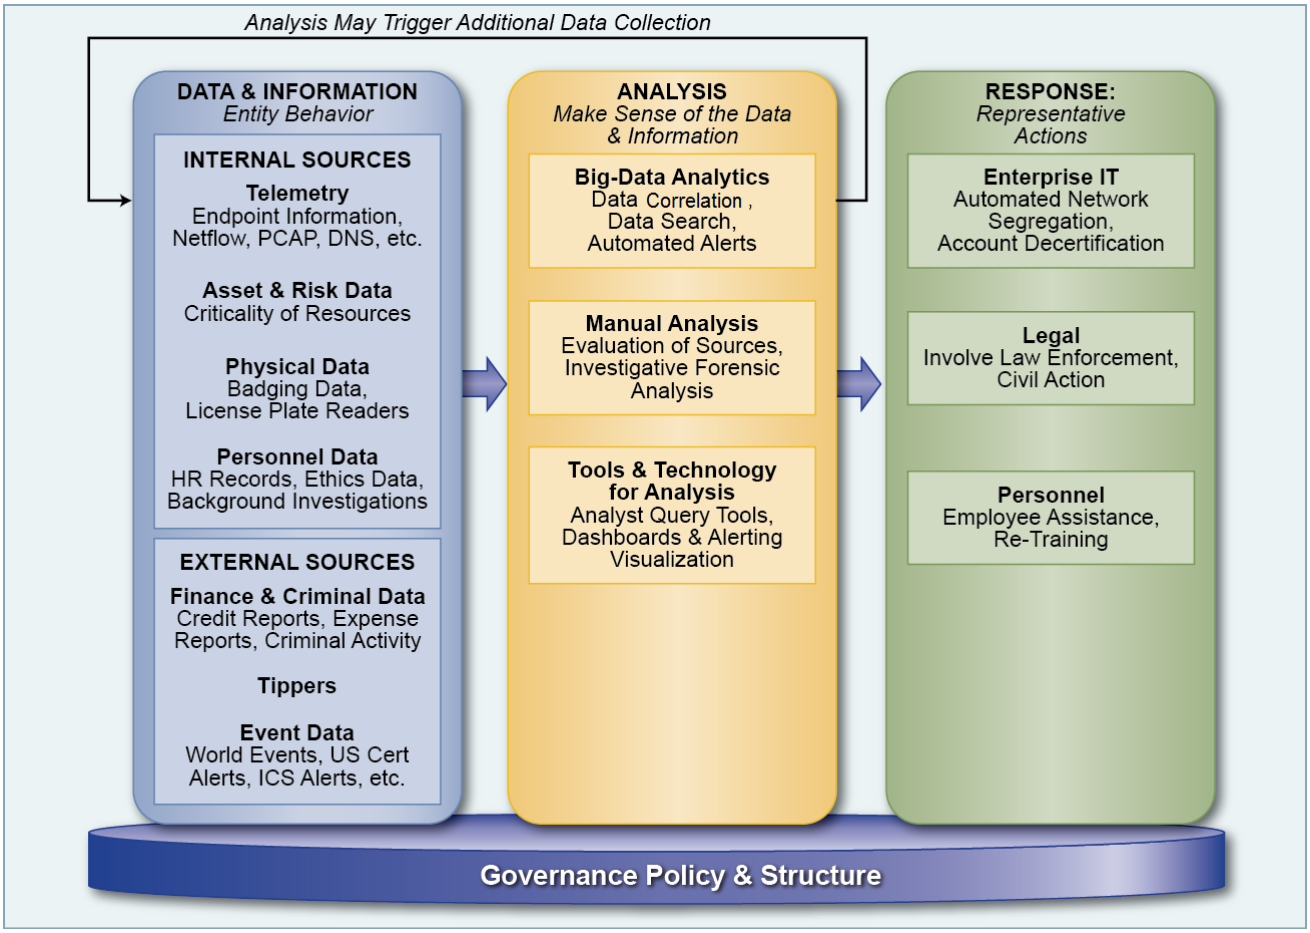

Figure 2: *Insider Threat Context*

A continuous assessment capability will take advantage of a variety of analytics that use many forms of telemetry[4] and other data to detect these indicators. Data and information includes both internal and external sources. Internal sources include telemetry derived from sensor platforms, enrichment data derived from telemetry and other on-line organizational resources, and access policy and configuration data associated with system resources. These are typically part of an existing SIEM/SOC architecture, and insider threat programs can leverage the existing capabilities to implement additional analytic indicators. Internal sources, such as content-based data used in an insider threat program, require protection that is not part of SIEM/SOC architecture. External sources include data acquired from outside the organization, as well as tip lines that would not typically be part of the SIEM/SOC architecture. Within the proper legal and policy constraints and with protections supporting privacy and civil liberty concerns, these categories of data can be utilized for insider threat analysis.

An organization's response posture will also influence the choice of analytic indicators deployed for detecting insider threat activities. Responses will vary widely based upon the

[4] **Telemetry –** data obtained from network components, security components, and other devices and applications (e.g. log files, network traffic) for the purpose of monitoring user activity.

intent of the threat agent. An untrained employee may require a response aligned with re-training. An administrator who is purposefully and improperly changing permissions and accesses may require an HR response. An employee selling trade secrets may require a law enforcement response. While not a complete list, the varied nature of these insider threat responses is driven by the behaviors identified by the analytics. Less compelling indicators of an increased but not imminent insider threat might prompt a general tightening of defenses for organizational resources. The procedures for determining and implementing a response must be governed by the organization's leadership and policies and guided by ethics, legal, policy and business considerations.

Organizations should be cautious in determining when to respond to the analytic indicators described in this whitepaper. In many instances, initial indicators will require additional data related to uncertain or non-specific indicators. Analysts should have the ability to configure the data sources to collect additional data needed to understand the context of such initial indicators to determine an appropriate response.

2. ANALYTIC OVERVIEW

An understanding of the characteristics of analytics which may indicate an attack is crucial to ensuring a continuous assessment capability. It is important to note that *an individual analytic by itself is neither a definitive indicator of an attack nor sufficient to distinguish between attack types*. The analytics that detect anomalous behavior often identify unusual, but not necessarily malicious or even damaging, activity. These analytics can have high false positive rates that can inundate insider threat analysts with many more false positives than true indicators of malicious insiders [see for example, The Base Rate Fallacy (Axelsson, 2000)]. Therefore, it is critical to consider combining the results of several independent indicators before making a decision to conduct an expensive investigation into a potential insider threat. Indeed, some of the more intrusive or expensive analytics might be reserved for deployment when other, simpler analytics indicate anomalous activity. How the analytics are deployed and how the indicators are combined requires careful analysis and will depend on the specific environment and circumstances of each organization. Additionally, the policy, privacy, civil liberties, ethics, legal authority, and other business considerations under which the insider threat program operates shape which analytics can be implemented for a given organization.

Analytics useful to an insider threat program are outlined below and described in detail in sections 3, 4, and 5. For clarity, the analytics are organized into three classes: Activity-Based Analytics, Content-Based Analytics, and Inferential Analytics**.** In turn, each class is divided into categories of analytic indicators. It is important to note that this white paper identifies many important analytic indicators but does not include analysis of the effectiveness of specific implementations.

Activity-Based Analytics (section 3): use content and event-based information derived from telemetry to understand user activity.

- **System –** Analyze changes or trends in IT asset behavior, data, or access patterns.

- **Facility –** Analyze changes in the time or locality of physical access patterns.
- **Business Capabilities –** Analyze business or mission capabilities either internally for changes and failures or externally for leaks or capability duplication.

Content-Based Analytics (section 4): use content extracted from telemetry to examine user characteristics.

- **Social –** Analyze social interactions and communications.
- **Health –** Analyze network activity and content to derive potential indicators of mental health issues.
- **Human Resources –** Analyze network activity and content for indicators of external life events or internal complaints against the agent.

Inferential Analytics (section 5): use telemetry to refine the understanding of user behavior in light of other information sources.

- **Financial –** Analyze network activity and content to derive indicators of unexpected changes in wealth or affluence.
- **Security –** Analyze telemetry for indicators of security violations.
- **Criminal –** Previous disposition in court and criminal activity.

3. ACTIVITY-BASED ANALYTICS

Activity-based analytics utilize telemetry directly to provide indicators of user activity. Implementers should consider the dual-nature of these analytics as potentially valuable for both insider and external threats and should leverage existing SIEM system and SOC capabilities whenever possible. The analytics range in complexity depending on whether the analytics are looking for discrete, typically unauthorized, events, or whether they build statistical models to characterize user behavior types. Analytics that detect deviations from established or learned baselines, and comparisons of individuals against peer behavior types combine the complexity of behavior analytics with the ability to produce real-time alerts. Activity-based analytics can be used for real-time alerting, understanding trends, and for forensic activity. Implementers should be cautious of putting too much weight on any single analytic, as the false base-rate fallacy[5] can result in a significant number of false positives.

[5] The base-rate fallacy may affect the operational effectiveness of an intrusion detection system (IDS). Due to the base-rate fallacy problem, the limiting factor for IDS performance is not the ability to correctly identify behavior as intrusive, but rather to correctly suppress false alarms or false positives (Axelsson, 2000).

a. SYSTEM INDICATORS

The following analytic indicators utilize telemetry to identify potential system anomalies that can be attributed to individuals. Some system anomalies will indicate activity that will have to be analyzed further to distinguish whether the anomaly is the result of insider or external activity. Additional weight should be given to anomalies that indicate a risk to sensitive resources—whose exposure, corruption, degradation would cause significant damage to the organization.

Authentication and Authorization: Required to access sensitive organizational resources, especially those deemed critical to the organization's mission. These resources include data, services, and capabilities and are available in operational as well as backup systems.	• **Analytic Indicator – Authentication and Authorization Failure:** Failed authentication or attempts to access unauthorized data indicate an individual's desire for data outside their work role. Low complexity analytics that indicate failed or unauthorized attempts to access resources are primary indicators for the majority of attack types. They can be implemented to provide real-time alerts.

Data Access Patterns: The specific time and frequency of accesses to sensitive data.	• **Analytic Indicator – Changes in Data Access Patterns:** Changes in data access patterns indicate interest in resources not previously associated with their work role, potentially for unauthorized purposes. Moderate complexity analytics that detect temporal changes in an individual's access behaviors via any accounts or methods accessible to the user are primary indicators for most attack types. Such analytics can develop models over time to characterize potentially suspicious changes in an individual's behavior.

	• **Analytic Indicator – Access Inconsistent with User Class:** Users accessing information that is not commonly accessed as part of their work role may be using this access for unauthorized purposes. Analytics that model user accesses to system resources are primary indicators for most attack types. These high complexity analytics compare an individual's access pattern against that of others, either by work role or by work habits. Such analytics can develop models over time to characterize atypical behavior that might be viewed as suspicious.

Network Patterns: The specific network protocols, sources and destinations, size of packets, and frequency of sessions associated to user applications.	• **Analytic Indicator – Changes in Network Patterns:** Deviations from established behavior indicate changes in a user's objectives, attitude, or skill. Moderate complexity analytics that detect temporal changes in network traffic associated to user activity are primary indicators for financial fraud and opportunistic data theft, and supporting indicators for other attack types. Such analytics can develop models over time to characterize potentially suspicious changes in an individual's behavior. • **Analytic Indicator – Network Patterns Inconsistent with User Class:** Deviations from established baselines in network traffic associated to user activity relative to their work role or peers indicate possible carelessness or risky or abusive behavior. High complexity analytics that compare network traffic associated to an individual's activity against that of others, either by work role or by work habits are primary indicators for accidental leaks, espionage, and opportunistic data theft. Such analytics can develop models over time to characterize atypical behavior.

Data Exfiltration: Unexpected or unauthorized removal of sensitive data from organizational systems.	• **Analytic Indicator – Data Exfiltration:** Analytics of varying complexity that detect large or unusual quantities or types of data leaving an enterprise through print services, in email, or via removable media are indicators for all attack types related to data loss. They are primary indicators for spontaneous acts and supportive indicators when extensive planning can detect the insider earlier in their planning. These analytics detect exfiltration as it occurs and can provide real-time alerts to egregious activity. Detection of more subtle exfiltration requires developing models to detect atypical patterns.

Unauthorized Data Access Methods: Unusual or unauthorized connections to facilitate access to, or removal of data from an organization's official systems.	• **Analytic Indicator – Unauthorized Data Access Methods:** Simple analytics that detect unauthorized connections of devices or between systems, or unauthorized activity related to data movement are primary indicators for a number of attack types. They can be implemented to provide real-time alerts.

Privilege Change: Attempts to gain privileges within the system.	• **Analytic Indicator – Privilege Change:** Unusual or unauthorized activity to gain additional privileges indicates unauthorized activity or disregard for established processes. Simple analytics that detect unauthorized privilege escalation attempts are primary indicators for most attack types. These analytics can be implemented to provide real-time alerts.

Erroneous Defensive Posture Changes: Rapid modifications in system configuration to an unsecure state, followed by attempts to restore the system to its proper state.	• **Analytic Indicator – Erroneous Defensive Posture Changes:** Rapid changes in system state indicate inadequate training or system probing. Medium to high complexity analytics that detect unusually rapid changes in commands or defensive posture associated to user errors and attempts to resolve those errors are primary indicators for misuse and opportunistic data theft, and supporting indicators for accidental leaks and product alteration attacks. Depending on the sophistication of the analytic and type of changes being detected, these analytics might be able to produce real-time alerts.

Command Usage: Unexpected or unusual command usage.	• **Analytic Indicator – Improper Command Usage:** Repeated occurrences of unexpected or unusual command usage relative to peers indicate lack of training or probing of system response. Analytics of varying complexity that detect improper command usage can support detection of accidental leaks. Depending on the sophistication and comprehensiveness of the analytic, it might be able to produce real-time alerts.

Knowledge Access: Aggregation of knowledge over time that exceeds user's need to know.	• **Analytic Indicator – Knowledge Access**: Accesses to gain excessive knowledge inconsistent with a user's role may indicate malicious user objectives. Moderately complex analytics that detect changes in search patterns, including massive searches and directory combing, or access to a wide variety of subjects, especially to subjects outside of the user's work role, are primary indicators of espionage, misuse, and opportunistic data theft, and a supporting indicator for financial fraud. These analytics can be implemented to generate real-time alerts for specific information quantities or especially critical combinations of information access.

Audit Log Modification: Modification or deletion of audit logs.	• **Analytic Indicator – Audit Log Modification:** Simple analytics that detect deletion or modification of security data and audit logs are primary indicators for many of the attack types. These analytics can be implemented to generate real-time alerts.

b. FACILITY INDICATORS

The physical location of an individual within a facility can be approximated by telemetry.

Time of Access Pattern Changes: Time of access patterns reflects the work schedule of the user.	• **Analytic Indicator – Time of Access Pattern Changes:** Unexplained or unusual changes in a user's work schedule could indicate an attempt to perform activity outside of scrutiny of coworkers or supervisors. Moderately complex analytics that learn patterns of times of access to system resources and can be configured to alert real-time to changes in these patterns are important indicators for most of the attack types. Analytics can be configured for real-time alerts of deviations that indicate highly unlikely events, such as logging into internal resources without badging in, or multiple days of continuous access.

Locality of Access Pattern Changes: The locations where users typically access system resources.	• **Analytic Indicator – Locality of Access Pattern Changes:** Unexplained or unusual changes in the location(s) from which a user typically accesses a system indicate attempts to bypass audit or security mechanisms, or to avoid scrutiny of coworkers or supervisors. Moderately complex analytics that learn patterns of the physical location where access to system resources originate and can be configured to alert real-time to significant or extremely low likelihood changes in these patterns are important indicators for most of the attack types.

c. BUSINESS CAPABILITIES INDICATORS

Business capabilities are activities that adversely impact the business functions of an organization, perhaps in subtle ways.

Failure Correlation: Business product flaws or unauthorized functionality associated to an individual.	• **Analytic Indicator – Failure Correlation:** A pattern or history of product flaws inconsistent with the user's skill level can indicate carelessness, or an effort to diminish or damage an organization's reputation, or to facilitate attacks against the organization's customer base. Highly complex analytics on performance metrics trends can be used to provide primary indicators of product alteration or sabotage attacks, and could be supporting indicators for system misuse or workplace violence.

Malware Deployment: Unauthorized deployment of malicious code.	• **Analytic Indicator – Malware Deployment:** Simple analytics to detect unauthorized installation of malicious code in real-time can detect unauthorized activity related to espionage, system misuse, or sabotage attacks. This analytic will likely be part of a SIEM system.

Deletion or Modification of Data or Infrastructure: Unauthorized deletion or changes to data or infrastructure.	• **Analytic Indicator – Deletion or Modification of Data or Infrastructure:** Unauthorized deletion or corruption indicates behavior likely to cause damage. Simple analytics that detect unauthorized deletion or modification of critical data or infrastructure are primary indicators for financial fraud, system misuse and product alteration, and supporting indicators for other attack types. These analytics can provide real-time alerts for data or infrastructure components identified as critical.

Analysis of Competitor: Advances in competitor capabilities or products that appear to take advantage of unreleased organizational information.	• **Analytic Indicator – Analysis of Competitor:** Analysis of competitor capabilities that indicates unauthorized disclosure, whether from analysis of the competitor, public media or other information sources (e.g., tips), could be the result of insider activity. High complexity analytics to determine competitor capabilities and associate unexpected advances to potential loss of organizational intellectual property provide primary, albeit after-the-fact indicators for espionage and system misuse.

Analysis of Public Media: Discovery that data that has been improperly disclosed to the public.	• **Analytic Indicator – Analysis of Public Media:** Simple to moderately complex analytics that look for organizational intellectual property releases from public media. Such analytics provide primary, albeit after-the-fact indicators of espionage, system misuse or physical theft, and supporting indicator of accidental leaks.

Attribution of Disclosure: Access to sensitive data or intellectual property illegitimately obtained by a competitor or publicly disclosed.	• **Analytic Indicator – Attribution of Disclosure:** Access to data involved in unauthorized disclosures, whether from analysis of the competitor, public media or other information sources (e.g., tips), indicates potential involvement in the disclosure. Highly complex analytics that identify and track access to data that has been disclosed publicly or to an adversary are focused to address losses whether from accidental leaks, espionage, system misuse, or physical theft attacks. They are primarily used after-the-fact on potentially archived data.

Retrieval: Attempted retrieval of corrupted or deleted data from backup or archive.	• **Analytic Indicator – Retrieval:** Recovery of data indicates an attempt to recover from inadvertent deletion, or to cover traces for temporary modifications or deletion. Moderately complex analytics to detect unusual retrieval of data provides a primary indicator for opportunistic data theft and supporting indicators for accidental leaks, financial fraud, and system misuse.

4. CONTENT-BASED ANALYTICS

Content-based analytics use content extracted from telemetry to examine user characteristics. These analytics are related to an employee's outward, observable personal behavior, not network behavior, and are routinely practiced by Human Resources and management. Here the focus is on approximations of this behavior available from telemetry and content. The typical content-based analytic is medium to high complexity and generally looks for trends rather than generating real-time alerts. They focus on user-generated content and language to infer user behavior characteristics.

Implementers must keep in mind the natural and normal variations between people and within any individual day-to-day when developing and using these indicators. Alert thresholds and interpretation may vary widely with culture, region, and even the type of work done by the individuals. Careful design is needed to create behavioral indicator systems with an acceptable accuracy, while also respecting the privacy, civil liberties, and trust of the individuals being monitored.

a. SOCIAL ANALYTICS

Social factors reflect how poorly a person interacts with those around him or her and often correlates to their willingness to act outside of socially acceptable limits. In particular, rapid changes in those interactions are often indicators of great stress, which in turn can indicate a person is more susceptible to do something harmful to the

organization. The indicators in this section are not direct indicators of an attack, but attempt to identify a person whose actions may warrant additional, direct scrutiny for harmful activity.

Disregard: Having disregard for authority and accepted practices, as well as the impact of actions on others.	• **Analytic Indicator – Disregard:** Regularly ignoring or rejecting policies, workplace culture, etc. could signify disregard for security requirements and defenses. Analytics that detect Disregard are either primary or supporting indicators for a majority of attack types.

Personal Inflexibility: An inability to properly adapt to stress or adversity.	• **Analytic Indicator – Personal Inflexibility:** Noticeable lack of resilience in dealing with challenging situations or changes in the environment, or undue stress from these causes. Analytics that model resiliency are important supporting indicators for all but one of the attack types.

Unusual Contacts: Contact with individuals or groups possibly connected to harmful activity, which may indicate the contact has undue influence over the employee.	• **Analytic Indicator – Unusual Business Travel:** A change in business travel to foreign countries not usually visited in the normal course of business, or meetings with representatives of those countries. Analytics that detect unusual business contacts provide primary indicators of espionage or product alteration attacks, as well as supporting indicators for other attack types. • **Analytic Indicator – Personal Travel:** Discovery of frequent personal travel to foreign countries not usually visited in the normal course of their employer's business, or, meetings with representatives of those countries. This is especially noteworthy if the employee has attempted to hide the travel. Analytics that detect unusual or hidden personal contacts provide primary indicators of espionage or product alteration attacks, as well as supporting indicators for other attack types. • **Analytic Indicator – Unauthorized or Inappropriate Associations:** Unauthorized or inappropriate association with hostile groups, or participation with them, can indicate a change in allegiance that can precede acts to damage an organization. Moderately complex analytics on user contact information provide supporting indicators for many attack types.

Withdrawal: Socially moving away from others and reducing contact.	• **Analytic Indicator – Withdrawal:** Noticeable and extended reduction in personal and social network interactions. Analytics that detect withdrawal are primary indicators for physical theft or workplace violence attacks, and supporting indicators for other attack types.

Workplace Events: A major workplace event that negatively changes an employee's status, circumstances, or satisfaction, such as a program cancelation, a corporate reorganization, or poor performance reviews.	• **Analytic Indicator – Workplace Events:** Many workplace events can be construed as negative by employees. Specific indicators may be unique to each workplace and current environment. Negative performance reviews or warnings are particularly significant. Analytics that detect adverse behavior changes related to workplace events, such as layoffs, are primary indicators for physical theft or workplace violence attacks, and important supporting indicators for other attack types.

Workplace Satisfaction: The degree to which an employee is content with their work situation.	• **Analytic Indicator – Workplace Satisfaction:** An employee who is noticeably disgruntled at work. A sudden change from positive to negative satisfaction is particularly significant. Analytics that model increasing disgruntlement are important indicators for most attack types, and are primary indicators for physical theft, product alteration, sabotage, and violence.

b. HEALTH ANALYTICS

Erratic behavior associated to potential unmitigated mental health issues might be detectable in network behavior or content. Access to these indicators should be restricted to authorized users, as HIPA and other privacy policies and laws might be applicable.

Mental Instability: Unpredictable or extreme behavior associated with unmitigated mental instability.	• **Analytic Indicator – Mental Instability:** Analytics that characterize user behavior and detect unstable or unpredictable behavior, or changes in behavior typically associated with mental instability are primary indicators of workplace violence or sabotage, and supporting indicators for system misuse of product alteration.

Impulse Control: Performing potentially inappropriate behavior without forethought.	• **Analytic Indicator – Impulse Control:** Analytics that characterize a user's propensity towards kleptomania or detect behavior changes associated with impulse control disorders can provide primary indicators of financial fraud, opportunistic data theft, or physical theft.

c. HUMAN RESOURCES ANALYTICS

Human Resources typically provide personnel action information about users that can impact their performance or attitudes towards coworkers. Simple to highly complex analytics can provide indicators from content or language available in telemetry to infer significant events or detect responses correlated with such events. The results of these analytics can infer protected information regarding individuals and should only be accessible by authorized personnel.

Major Life Event: Major life events outside the work environment (e.g., change in marital status, birth of a child, or death of a relative).	• **Analytic Indicator – Major Life Event:** Major life events can impact work behaviors and can create stresses that precipitate bad decisions. Analytics that detect those stressors or adverse behaviors associated to major life events are primary indicators for many attack types.

Complaints Against the User: Complaints from peers or coworkers about a user's behavior.	• **Analytic Indicator – Complaints Against the User:** Complaints from coworkers can trigger retaliation against both the complainant and the organization. Dissatisfaction towards organization, assets, or individuals could also lead someone to find another entity more appreciative of their skills, knowledge and access leading to theft of data and economic espionage. Simple analytics that attempt to detect the fact of such complaints provide supporting indicators for sabotage, violence, theft of date, and economic espionage

Negative Reviews: Poor performance, security reviews or negative feedback regarding performance.	• **Analytic Indicator – Negative Reviews:** Studies have shown high correlation between people receiving negative messages or feedback at work and then acting out in a negative manner towards the organization. Dissatisfaction towards organization, assets, or individuals could also lead someone to find another entity more appreciative of their skills, knowledge and access leading to theft of data and economic espionage. Simple analytics that attempt to detect the fact of negative reviews provide supporting indicators for sabotage or violence.

5. INFERENTIAL ANALYTICS

Inferential analytics use telemetry to refine the understanding of user behavior in light of other information sources. The analytic indicators in this class of analytics are more abstract than either of the previous classes. They indicate analytics that update a baseline established by external reporting typically available to an insider threat program as part of a continuous assessment effort. In some cases, fairly simple analytics on the result of content-based analytics can be used detect changes from assumed behavior based on a current baseline; accuracy can suffer with time between baselines. Inferential analytics can aid in investigations of suspicious individuals.

Since the results of these analytics provide unverified claims about potentially protected information, they should be used with caution, and the use of such analytics should be restricted to properly trained and authorized insider threat personnel.

a. FINANCIAL ANALYTICS

Unexpected or unexplained changes in financial status can indicate undue influence that adversely impacts an organization. Periodic Financial disclosures provide individual baselines, but indicators of financial changes from content or behavior can be integrated as part of a continuous assessment program.

Change in Means: Sudden affluence or sudden or excessive debt.	• **Analytic Indicator – Observed Temporal Change in Means:** Sudden affluence might indicate an individual is being bribed or induced to perform damaging activity; sudden and excessive debt may be associated with stress that can cloud judgment. Moderately complex analytics that detect changes in means are primary indicators for a number of attack types. • **Analytic Indicator – Observed Change in Means Relative to Peers:** Deviation in affluence relative to peers can create tension that impacts decision making. Moderately complex analytics that correlate social indicators to content or behavior associated to living beyond one's means provide primary indicators for a number of attack types. • **Analytic Indicator – Financial Reporting:** Significant changes in means derived from financial reports or other external sources can indicate hidden wealth or financial stresses. Correlating social and behavioral indicators against potential damaging activity models associated to recent disclosures provides primary indicators for a number attack types.

b. SECURITY ANALYTICS

Past security violations indicated in external reporting, or observed within an operational generally indicate an increased likelihood for additional, perhaps more severe violations and insider attacks.

Security Violations: Unauthorized activity impacting the security of an organization.	• **Analytic Indicator – Change in Violation Patterns:** Low to moderately complex analytics that detect changes in the frequency or severity of security violations are primary indicators for the types of violations that can be inferred, to include sabotage and theft, and can be supporting indicators for most other attack types. Some of these analytics can be configured to provide real-time alerts. • **Analytic Indicator – Duration and Regularity of Security Events:** Simple analytics that track security events over time can provide indicators for many of the attack types.

	• **Analytic Indicator – Unauthorized or Inappropriate Use of Tools:** Specific security violations related to the use of tools, (e.g., network sniffers and network analytic tools) that are specifically detrimental to a business activity can indicate attacks to install unauthorized functionality.

c. CRIMINAL ANALYTICS

Past criminal activity outside the workplace may warrant additional scrutiny of a user within the workplace.

Restraining Order: Legal restraining order issued against a user.	• **Analytic Indicator – Restraining Orders:** Content-based analytics that indicate the fact of legal action, and/or behavior that reflects unacceptable activity associated to an existing restraining order indicates a primary indicator for violence. • **Analytic Indicator – Wage Garnishments:** Legal actions to collect debt indicate financial stress that may warrant additional scrutiny of an individual relative to financial fraud attack. Analytics that track a user under such orders to detect changes in behavior can provide primary indicators for financial fraud. • **Analytic Indicator – Violence Outside Workplace:** Stresses and violent behavior outside the workplace can indicate an increased likelihood of violence toward the organization or coworkers. Content based analytics to detect negative changes in sentiments or behaviors provide primary indicators of violence attacks.

Recent Increase in Criminal Events: Recent criminal activity.	• **Analytic Indicator – Recent Increase in Criminal Events:** Criminal activity can damage an organization's reputation, or represent a threat to the organization or its employees. Content based analytics that detect indicators of criminal activity, or detect high risk behavior or negative sentiments associated to recent criminal activity are primary indicators for physical theft and supporting indicators for product alteration.

6. IMPORTANT ANALYTICS FOR ATTACK TYPES

For the analytic indicators described above, **Figure 3** below provides a summary of the applicability of each analytic indicator to the attack types. An organization interested in a specific attack type can read down the column to identify analytic indicators that would be useful in detecting that activity. Analytic indicators with a "1" in the cell indicate primary indicators that are more likely to be correlated with the activity; cells labeled "2" are considered supporting indicators; and cells containing no number may or may not provide useful information to an analyst. Looking at the chart by rows indicates the utility of implementing various analytics and could be useful in making investment decisions. Shading in the table is provided as a visual aid.

Analytic Indicator Category	Analytic Indicator	Attack Types								
		Accidental Leak	Espionage	Financial Fraud	Misuse	Opportunistic Data Theft	Physical Theft	Product Alteration	Sabotage	Violence
Activity-Based Analytics										
System	Authentication and Authorization Failure		1	1	1	1		1	1	
	Changes in Data Access Patterns		1	1	1	1		1	1	
	Access Inconsistent With User Class		1	1	1	1		1	1	
	Changes in Network Patterns	2	2	1		1			2	
	Network Patterns Inconsistent with User Class	1	1			1				
	Data Exfiltration	1	2	2		1			2	
	Unauthorized Data Access Methods	1	1		1	1				
	Privilege Change		1	2	1	1		2		2
	Erroneous Defensive Posture Changes	2			1	1		2		
	Improper Command Usage	2								
	Knowledge Access		1	2	1	1				
	Audit Log Modification		2	1	1			1	1	
Facility	Time of Access Pattern Changes		2	2		2	1	1	1	1
	Locality of Access Pattern Changes		2	2		2	1	2	1	1
Business Capabilities	Failure Correlation				2			1	1	2
	Malware Deployment		1		1				1	2
	Deletion or Modification of Data or Infrastructure			1	1	2		1	2	
	Analysis of Competitor		1		1					
	Analysis of Public Media	2	1		1		1			
	Attribution of Disclosure	1	1		1		1			

Analytic Indicator Category	Analytic Indicator	Attack Types								
		Accidental Leak	Espionage	Financial Fraud	Misuse	Opportunistic Data Theft	Physical Theft	Product Alteration	Sabotage	Violence
	Retrieval	2		2	2	1				
Content-Based Analytics										
Social	Disregard	2			1	1	1	2		1
	Personal Inflexibility	2	2	2	2	2		2	2	2
	Unusual Business Travel		1				2	1		2
	Unusual Personal Travel		1			2	2	1		2
	Unauthorized or Inappropriate Associations		1			2	2	1		2
	Withdrawal					2	1	2	2	1
	Workplace Events		2		2	2	1	2	2	1
	Workplace Satisfaction			2		2	1	1	1	1
Health	Mental instability				2			2	1	1
	Impulse Control			1		1	1			
Human Resources	Major Life Event			1			1	1	1	1
	Complaints Against the User								2	2
	Negative Reviews		2	1	1	2	1	2	1	1
Inferential Analytics										
Financial	Observed Temporal Change in Means		1	1			1	1	1	2
	Observed Change in Means Relative to Peers		1	1			1	1	1	2
	Financial Reporting		1	1				1	1	2
Security	Change in Violation Patterns		2	2	2		1	2	1	2
	Duration and Regularity of Security Events	1		2	2		1	2	1	2
	Unauthorized or Inappropriate Use of Tools							2		
Criminal	Restraining Orders				2					1
	Wage Garnishments, etc.			1						
	Violence Outside Workplace				2				1	2
	Recent Increase in Criminal Events						1	2		

Figure 3: *Analytic Indictors for Attack Types*

E. Analytic Process & Investigations

Automated security tools can provide hints of insider behavior in much the same way they help protect organizations from external threats. Like external tips from human sources, these automated indicators require deeper evaluation to determine the type and nature of the threat. Once tipped—either from automated or human sources—the investigator must ask penetrating questions and hunt among data sources to retrieve, correlate, and understand the risk posed by the presumed insider threat.

Investigators face many challenges in answering the "Who, What, When, Where, and How" questions for any security event. Insider threat analysts attempting to answer "Why" questions related to understanding the context of anomalous user activity have even greater challenges. For example, because compromised IT infrastructure can be used to obtain valid user credentials, investigators must carefully distinguish between activity that is deliberately caused by a malicious insider and a masquerade operation that uses legitimate credentials from an unwitting user's compromised account. The telemetry may be able to provide clues of an insider versus an external threat, but this may not be definitive. Thus, the cyber investigation should be supplemented with classic human investigative methods to determine ground truth.

Investigators also must know and adhere to many non-IT related requirements, such as knowing when information collected during an investigation is subject to legal chain-of-custody requirements.

Because of these special needs, a security analyst that does insider investigations may require additional training and should be separately identified within an organization's overall SOC team or ideally from outside the SOC. While the primary incident investigation may be done by the SOC, there should be a clear handoff to an "insider-qualified" analyst when a potential insider is detected.

F. Data Sources For Analytics

1. Data from Security and Network Components

Current cyber security practice uses a number of functional IT components that are employed in four basic activities: prevention, detection, response, and recovery. **Figure 4** provides sample architecture of functional components—each of which is capable of consuming configuration data and generating telemetry. The configuration data expresses policy in terms of permitted and denied access to resources—users, devices, networks, applications, data, etc. The telemetry is used to inspect whether the expressed policy is being enforced and whether unexpected or anomalous activities are underway. In addition to telemetry, enterprise business and IT applications capture and communicate user or customer-generated content (e.g., email, chat, documents, files, etc.) that must be protected and may also be analyzed for security-relevant information.

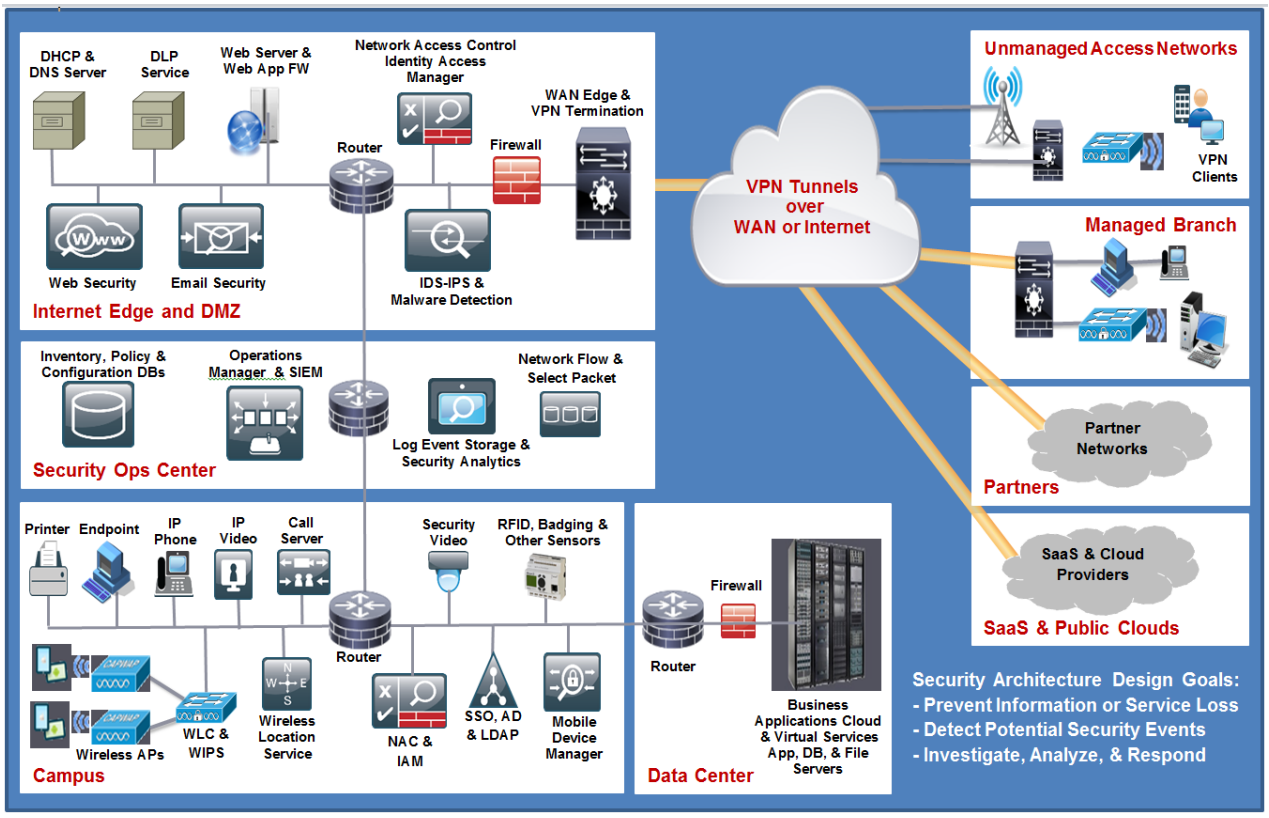

Figure 4: *Sample Architecture for Data Sources*

The sample functional architecture in **Figure 4** depicts a typical, large-scale, distributed enterprise with an emphasis on the security and network components. The components are arranged in blocks which depict a typical Internet edge, campus, and data center with secure connections to organizationally managed branch locations, partner sites, and SaaS/cloud providers. There is also a link to unmanaged access networks where endpoints configured with VPN clients can securely connect to organizational assets. All of these functions are connected to a Security Operations Center (SOC) which provides resources for managing the organization's asset inventory and configurations as well as collecting and analyzing the telemetry generated from them. In most cases, the security operational status is evaluated from the telemetry and presented in real-time through a Security Information and Event Management (SIEM) system. Additional sensitivities and specific authorities associated to monitoring users for insider threat might require that insider threat analysts be separate from SOC operators. The architecture described includes such a separation, but in smaller organizations or when the insider threat program is part of a SOC, these could be combined.

The functional components shown in **Figure 4** are meant to be illustrative regarding potential security data sources and comprehensive. An organization's IT and security infrastructure may exclude, connect, or combine these functions in various ways according to their needs and budget. Not all components may be available in an organization's system architecture. Some of these components might be organized into separate systems; if this is the case, data from these separate systems can be provided to an insider threat analyst as external enrichment data. For example, badging information and security video may be

included in a facility's physical access control system and isolated from the operational IT network. Regardless of configuration, each of the functional components shown in **Figure 4** can provide useful telemetry to investigators. The specific data sources and data types available within each block of **Figure 4** are provided in **Figure 5** below. The data types are grouped into the following categories:

- **Alerts**: Real-time notifications that a specific, pre-configured event has occurred. Examples: a login threshold was exceeded; an intrusion signature was tripped, etc.
- **Content**: Data held within containers intended for communication, transport, or processing. Examples: email contents, user files, web pages, corporate data base information, network data payloads (includes audio/video streams), etc.
- **Flows**: Network metadata that describes transport of data across a TCP/IP network, but does not contain the information being communicated. Example: Netflow, IPFIX, etc.
- **Log**: Records of events that have occurred within a device, system, or application. These records usually take the form: [timestamp] -> [event]. Examples: syslog, login logs, etc.
- **Identity**: Data that may contains characteristics that describe a specific individual. Examples: user names, certificates, etc.

Sample Architecture Block	Data Sources	Description	Data Types				
			Alerts	Content	Flows	Logs	Identity
Internet Edge and DMZ	DHCP Server (Dynamic Host Control Protocol)	Assigns IP addresses to network hosts.				X	
	DNS server (Domain Name Service)	Resolves DNS names (hosts, URLs) to IP addresses.				X	
	DLP Service (Data Loss Prevention)	Inspects outgoing content for sensitive information.	X	X		X	
	Web Server	Provides web-published content to internal and external users.		X		X	
	Web Application Firewall	Protects web server from attack.	X			X	
	Web Security	Inspects Web connections to external sites. May proxy connection to scan for malicious content or references from remote sites.	X	X		X	X

Sample Architecture Block	Data Sources	Description	Data Types				
			Alerts	Content	Flows	Logs	Identity
	Email Security	Inspects Email to filter out spam, phishing, and malicious content—either directly included/attached or referenced in the email body.	X	X		X	X
	IDS-IPS (Intrusion Detection and Prevention Service)	Inspects network content to screen and/or potentially block security intrusion events.	X			X	
	NAC (Network Access Control)	Manages initial connections of devices to a network. Can trigger endpoint compliance scanning.	X			X	
	IAM (Identity Access Manager)	Manages assignment of identity information to users.	X			X	X
	VPN Termination/WAN Edge	Validates clients connecting over secure virtual private network connections.				X	X
	Firewall	Controls connections to network resources and services.	X			X	
	Router	Provides connectivity services to network components.		X	X		
Security Ops Center	Inventory, Policy, & Configuration Data Bases	Data stores for Enterprise IT & Security Management. Includes hardware/software/network configurations, User IDs & Permissions, Network IDs & Permissions, etc.					X
	Operations Manager & SIEM	User interface to Security analytics and databases.	X			X	X
	Log Event Storage & Security Analytics	Infrastructure to store and process system logs and event data.		X	X	X	
	Network Flow & Raw Packet Capture	Infrastructure to store and process network flow data and analyst-selected content.		X	X		

Sample Architecture Block	Data Sources	Description	Data Types				
			Alerts	Content	Flows	Logs	Identity
Campus	Endpoint	Connects users to services with visual interfaces. Endpoint types include desktops, laptops, pad devices, and smartphones. Must be instrumented to be managed and generate usable telemetry in BYOD environments.	X	X	X	X	X
	Printer	Provides print services. **if recording/storage enabled*		X*		X	
	IP Phone	Specialized endpoint devices for voice. **if recording/storage enabled*		X*		X	
	IP Video	Specialized endpoint devices for video conferencing. **if recording/storage enabled*		X*		X	
	Call Server	Manages IP voice/video/teleconferencing endpoint call setup.				X	
	Wireless LAN Controller	Connects wireless endpoints to Local Area Networks and assigns IP address. Can trigger endpoint compliance scanning.	X			X	
	WIPS (Wireless Intrusion Prevention System)	Inspects wireless connections of intrusion events.	X			X	
	Wireless Location Service	Provides location and movement information for wireless equipped devices.	X			X	
	Mobile Device Manager	Manages admission of wireless endpoints to LANs.	X			X	
	SSO, AD, & LDAP (Single Sign On, Active Directory, Lightweight Directory Access Protocol)	Provides identity and permission information for users within an enterprise.	X			X	X
	Security Video	Collects and stores video from enterprise security cameras.		X			
	RFID, Badging, & Other Sensors (Radio-Frequency Identification)	Collects device-unique information regarding user or equipment identity/credentials and movement in/out of an enterprise.	X			X	

Sample Architecture Block	Data Sources	Description	Data Types				
			Alerts	Content	Flows	Logs	Identity
	Wireless Access Points		X			X	
	Router	Provides intra-campus connectivity services to network components.		X	X		
	NAC	Manages initial connections of devices to a network. Can trigger endpoint compliance scanning.	X			X	
	IAM	Manages assignment of identity information to users.	X			X	X
Data Center	Application Tier	Business applications & data bases	X	X		X	X
	Infrastructure Tier	Infrastructure software/hardware/storage	X	X	X	X	X
SaaS & Public Clouds	Saas & Cloud Provider	Outsourced infrastructure & business applications support (e.g., Amazon, Salesforce) ** as contracted / negotiated*		X*		X*	X
Partners	Partner Networks	Business collaboration connectivity (e.g., Development contractors, sales partners) * as contracted / negotiated		X*		X*	X
Managed Branch	Instrumentation similar to Campus Block	See Campus Block data sources above.	X	X		X	
Unmanaged Access Networks	VPN Clients (Virtual Private Network)	Provide endpoint access similar to Campus Endpoints using authenticated clients over 'open' networks.	X	X	X	X	X

Figure 5: *Data Sources and Data Types for Sample Architecture Blocks*

2. DATA PROCESSING FLOW AND KEY DATA ELEMENTS

Each of the components shown in **Figure 4** and described in **Figure 5** can generate telemetry (logs, flows, alerts, etc.) or user content (email, chat, web queries, files, etc.) that can help determine the security state of the enterprise, including a potential or actual compromise by an insider. These components can produce enormous amounts of security-relevant data that needs to be captured, ingested, normalized, enriched, stored, and analyzed. Fortunately, modern "big data" architectures are capable of capturing and managing this flood of data and making it accessible to streaming and batch analytic tools as well as providing interfaces for ad hoc investigative queries. An overview of this information and processing flow is provided below in **Figure 6**.

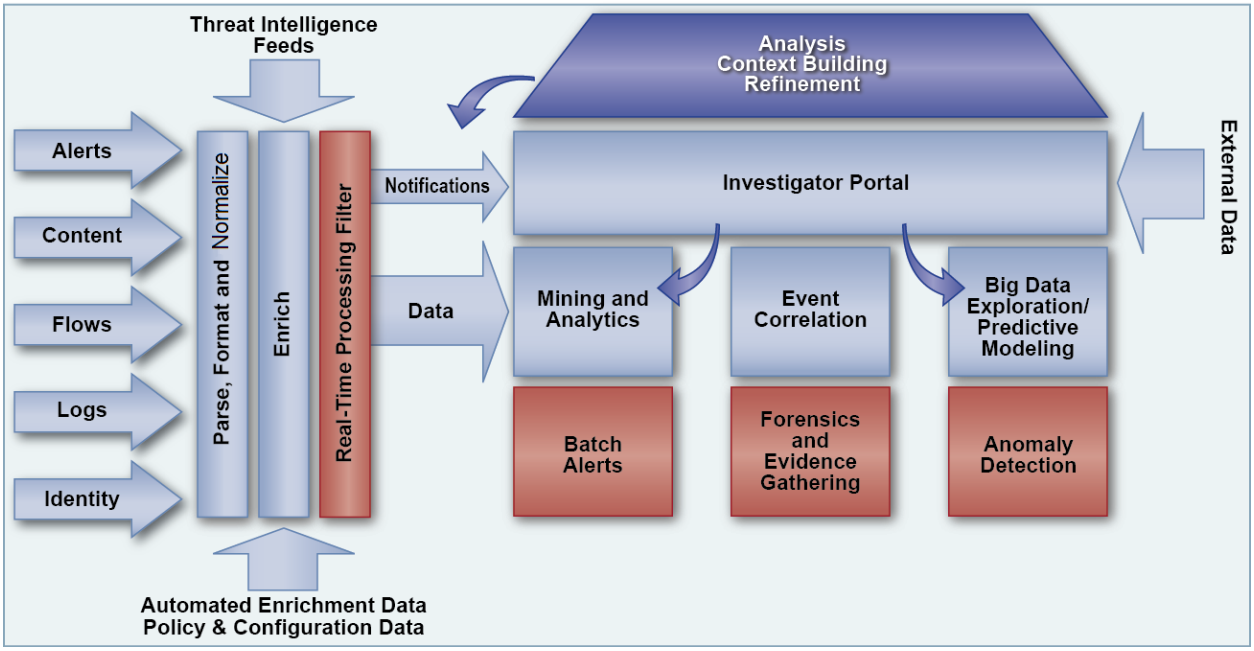

Figure 6: *Security Data Processing Overview*

From the left side of **Figure 6,** data is generated from the various IT and security components and delivered to a set of processes which parse, format, and normalize the data into units suitable for follow-on analysis. This step is usually followed with an 'enrichment' step that adds enterprise-specific information (and possibly privacy anonymization) based on the data type and source. This enriched information is then passed through a real-time processing filter which evaluates the data flow for threshold events that need to trigger notifications to an incident responder through the analytic portal. (For example, a large data flow destined to a suspicious external Internet address might generate such an alert.) Following the real-time alert filter, the data is then stored in a usable form for follow-on batch analysis that correlates activity across multiple data sources looking for indications of suspicious activity that would warrant further investigation. In another form of batch analytics, user and network activity is compared to typical activity baselines to detect anomalous events. To support specific investigations, authorized investigators can then use ad-hoc query capabilities provided by the portal to rapidly reconstruct events and gather forensic evidence. Finally, the external feeds shown in **Figure 6** represent sources of information that help determine alerting thresholds and updated threat vectors.

Security-relevant data comes in two basic forms —events and content. Events are usually captured as flows, log files, and alerts and may be forwarded to an analytic server for correlation and analysis. Communications metadata that describe network connections and data flows between network elements are a special kind of event that may be gathered directly from network components instead of from infrastructure or application log files. Event data is useful in answering questions about "what", "when" and "who". Content, such as data files, email, and chat, is usually stored on a service server and may be accessible from this storage point for content-scanning analysis. User content can also be gathered directly from IP networks as packets and re-assembled into sessions for analysis (e.g.,

web/database sessions). Content data is best suited to answer specific questions on "what" and "who", and can possibly reveal user intent or sentiment.

The left side of **Figure 6** includes an identity block to emphasize that insider detection based on cyber telemetry is critically dependent on identifying the user. For some data types, investigators may need to correlate multiple identifiers to positively identify the agent. For example, the "user" in network-based data will typically be in terms of IP addresses. In most enterprises, IP addresses are often dynamically assigned (using DHCP or other) when devices connect to the network. So, in order to associate this network data to a person, the IP address must be correlated to a specific machine (either a machine name or MAC address) using DHCP logs and/or an asset management system. Then the machine must be associated to a user account and the user account to a person using login logs (SSO, AD, AAA, etc.). An inventory system may also help identify the location, the organization responsible for the component, and perhaps even the unique user assigned to the component. Some enterprises may consider the "user" information to be privacy-sensitive and may need to anonymize or mask the identity information in a way that supports consistency and valid use by authorized investigators.

In addition to identity, many insider detection techniques require the temporal correlation of events such as login times, facility entry times, data creation and movement times, etc. To support this type of correlation, all the data sources (logs, events, content applications, etc.) should be configured with and connected to a universally unambiguous time source that can be used for time stamping across the enterprise.

With the growth in mobility and wireless connectivity, user location information can generally be gathered routinely within the enterprise. This data, when correlated with user events, can provide useful clues to detect potentially risky or negligent insider behavior. In environments that support Bring Your Own Device (BYOD), this may be one of the few methods available to locate unmanaged endpoints potentially deployed by an insider.

3. HOW THE DATA RELATES TO ANALYTICS

Important analytic indicators will require access to data from various sources. In general, it is possible to identify the types of data required for each analytic, even though the specific data sources will depend on the organization's system architecture, security architecture and data processing choices. **Figure 7** shows the data types required for each of the analytic indicators described above. In some cases, identity information may not be readily available from the typical data source and may need to be correlated from additional sources. In these cases, the analytics requirement for identity information is indicated as either part of the real-time processing of the analytic (R) or as a forensic capability (F).

Analytic Indicator Category	Indicator	Data Types				
		Alerts	Content	Flows	Logs	Identity
Activity-Based Analytics						
System	Authentication and Authorization Failure	X			X	X
	Changes in Data Access Patterns				X	X
	Access Inconsistent With User Class				X	X
	Changes in Network Patterns			X		F
	Network Patterns Inconsistent with User Class			X		R
	Data Exfiltration	X	X	X	X	X
	Unauthorized Data Access Methods	X			X	F
	Privilege Change	X		X	X	F
	Erroneous Defensive Posture Changes	X		X	X	F
	Improper Command Usage	X			X	F
	Knowledge Access		X	X	X	X
	Audit Log Modification	X			X	F
Facility	Time of Access Pattern Changes	X		X	X	X
	Locality of Access Pattern Changes	X		X	X	X
Business Capabilities	Failure Correlation		X		X	X
	Malware Deployment	X			X	X
	Deletion or Modification of Data or Infrastructure	X			X	F
	Analysis of Competitor		X			X
	Analysis of Public Media		X			X
	Attribution of Disclosure		X		X	F
	Retrieval			X	X	X
Content-Based Analytics						
Social	Disregard		X			R
	Personal Inflexibility		X			R
	Unusual Business Travel		X			X
	Unusual Personal Travel		X			F
	Unauthorized or Inappropriate Associations		X	X		F
	Withdrawal		X	X		R
	Workplace Events		X			R
	Workplace Satisfaction		X			F
Health	Mental instability		X			R
	Impulse Control		X			R
Human Resources	Major Life Event		X			R
	Complaints Against the User		X			R
	Negative Reviews		X			R

Analytic Indicator Category	Indicator	Data Types				
		Alerts	Content	Flows	Logs	Identity
Inferential Analytics						
Financial	Observed Temporal Change in Means		X			R
	Observed Change in Means Relative to Peers		X			R
	Financial Reporting		X			R
Security	Change in Violation Patterns	X			X	R
	Duration and Regularity of Security Events	X			X	R
	Unauthorized or Inappropriate Use of Tools	X			X	F
Criminal	Restraining Orders		X		X	R
	Wage Garnishments, etc.		X		X	R
	Violence Outside Workplace		X		X	R
	Recent Increase in Criminal Events	X	X		X	R

Figure 7: *Data Requirements for Analytic Indicators*

4. DATA PROCESSING REQUIREMENTS AND CHALLENGES

There are a myriad of big data infrastructure technology and tooling choices when it comes to handling the flood of telemetry needed for SOC operations and insider detection as shown in **Figure 7**. The primary requirements for this SOC-based infrastructure include the ability: to scale processing and storage capacity; to handle structured and unstructured data; and to support streaming and ad hoc queries. Unfortunately, there is no single obvious technology choice that can meet all of these requirements. In most cases, an enterprise will choose a cost-performant mix of commercial or open-source technologies that combine some sort of "big data" infrastructure with a SIEM data aggregation and user interface tool which carry out the SOC functions shown in **Figure 7**. The factors to keep in mind when making these technology choices include: availability of existing security analytics and applications; flexible tool interfaces; standard data import/export capability; ability to anonymize or limit access to privacy-sensitive data; and overall lifecycle support costs.

Data quality is a key requirement for successful security operations and insider detection. While much of the infrastructure can be instrumented to provide information, careful thought should be given to how and where the information will be generated; how the information will be analyzed; who has access to it; and how long it will need to be retained. For example, super-user logs that are resident on endpoint hosts can also be manipulated by any user that gains super-user privileges—including an insider. Network flow data that is merely sampled or doesn't cover all the egress points in the enterprise may not reveal suspicious flows. In some cases, critical business content may be held by a SaaS or cloud provider and flows of relevant access logs and content information to the SOC needs to be

arranged. Special data retention, robustness, and quality requirements may need to be addressed for data gathered subject to legal chain-of-custody specifications.

In all cases, legal and policy consideration must be given to where and what specific kinds of instrumentation can be used to track user activity. For example, certain forms of endpoint instrumentation may be considered a privacy risk and may not be allowed due to legal or policy reasons. This can have a significant impact on basic enterprise security operations as well as potential insider detection. In addition, there may be restrictions on scanning user-generated content for insider-relevant indicators. There may also be restrictions on how long the security data may be retained. Some recent reports suggest a minimum of four months retention of security data is needed to be able to effectively spot a potential insider.

Encryption can and should be used routinely to protect key business information and communications. However, encrypted content and network sessions pose a challenge to traditional plaintext-based scanning and inspection techniques, especially those hosted in the network. A multi-faceted approach that combines robust endpoint inspection with network activity monitoring can help mitigate this risk.

Cost may be a considered a challenge to implementing an effective data collection and processing infrastructure for insider threat. However, as described above, most enterprises should be able to leverage much of its existing security components and SOC operations to cover the cost of insider threat. Additional capital costs may be incurred for increased telemetry, storage, or data retention. There may also be operational costs for insider investigator training and for additional analytic software development and support. When evaluating the additional costs to fund an insider-specific extension, it is important to calculate the business risk posed by a potential insider. This can help identify both the critical assets to be protected from insider risk (information, facilities, people, etc.) as well as where to invest for the greatest effect.

G. RECOMMENDATIONS

Implement an Insider Threat Program: Organizations are strongly encouraged to develop an insider threat program to address one of the largest risks to their enterprise. Such a program provides an organizationally supported and integrated approach to addressing insider-based risks as well as setting the appropriate legal, ethical, and policy framework for effective, privacy-preserving implementation. In determining the scope of an insider threat program, the enterprise should weigh the costs needed to develop an early detection capability against the risk of potential losses that could result from the various types of insider attacks. For the losses that can be sustained, the cost of the recovery should be included in the analysis. As part of the program, all relevant attack types should be considered and all critical resources should be identified. See sections B and C for further discussion and references.

Deploy a Continuous Assessment Capability: As part of a well-governed and securely-operated insider threat program that is consistent with its legal, ethical and policy framework, an organization should implement a robust and flexible continuous assessment

capability. This capability should leverage tools that typically already exist within the organization's security infrastructure to detect malware and anomalous behavior. Additional automated tools and advanced analytics can be added to the existing security infrastructure to help designated insider threat security analysts detect potential threats and to distinguish between attack types posed by true insiders from external actor activity. See sections 0 and F for further discussion.

Deploy Analytics to Discover Potential Insider Threats: Organizations should identify their assets and assess their risk of damage from potential insider threat agents. This risk-cost analysis will help determine the type and deployment priority of analytics devoted to detecting the insider threat agents of concern. For each analytic deemed important for an organization, the data requirements should be determined and data feeds identified. If necessary, systems should be instrumented to securely provide the required data. Efficient data handling, processing, and storage capabilities should be established and separation mechanisms put in place to ensure privacy and civil liberties are maintained. Authorized individuals having access to sensitive data should be properly trained and their accesses audited to prevent misuse. See sections D, 0, and F for further discussion.

Provide Investigative Tools: Automated analytics can provide early indications of a potential insider threat. These indications must be correlated and investigated to understand the observed activity and determine if it is a false positive. To assist with this determination, organizations should consider implementing the analytics that are strong indicators for multiple insider attack types, such as those highlighted in **Figure 3**. They should also consider implementing as many analytics that are indicators for relevant attack types as feasible. Analysts should be able to collect additional data relative to their current understanding of a threat and tailor additional analytics to support further investigation. See sections D and 0 for further discussion.

Facilitate Attribution of Individuals: The most critical component in addressing insider threats is the ability to determine the specific individual associated with a potential attack activity. Facilitate attribution of individuals through a comprehensive identity management system for individuals. Consider using two-factor authentication with vetted identification e.g., Personal Identity Verification (PIV). Inventory management systems can also assist in associating network activities to specific device identifiers and the users or organizational elements responsible for those devices. The insider threat program will provide the governance and training necessary to determine appropriate courses of action for investigating positive indicators and determining attribution. This includes the analysis required to rule out potential external threats; methods for addressing potential false positives; access control, privacy preservation, and integrity protections for insider threat information; and investigative processes. See sections D, 0 and F for further discussion.

APPENDIX A: LEGAL CONSIDERATIONS FOR INSIDER THREAT MONITORING PROGRAMS

An effective insider threat program relies on information drawn from multiple sources. Companies, agencies, or other organizations designing such a program should consider restrictions under U.S. law (federal and state), company and government policies, and workplace and contractor agreements that may limit their use and disclosure of specific types of information used as insider threat analytic or telemetry. In addition, organizations conducting business outside the United States or handling the personal data of foreign customers and/or employees must be cognizant of the data privacy laws and directives that may govern how that information must be maintained and restrictions on its use and disclosure in the various countries in which they operate.

In many instances, if legal, policy and contractual requirements are met and applicable national laws are satisfied, protected data may be used in an insider threat program within certain parameters; however, a public or private organization must factor in necessary legal review as well as the time and processes it will require when planning and implementing the program. It is, therefore, prudent for an insider threat program to be developed in close coordination with legal counsel capable of furnishing advice about how protected data should be handled under relevant state and federal law. In addition, a multi-national company may need additional legal advice on running an insider threat program under the laws of other countries where it has operations.

1. INFORMATION PROTECTED UNDER STATE AND FEDERAL LAWS

There is no single, comprehensive federal law governing the collection and use of the type of data used in an insider threat program. Instead, a patchwork of federal and state laws and regulations dictate how employers may use and handle protected personal and confidential employee information. Adding to this complexity are the disparate national privacy laws with which a multi-national organization also may need to comply. Additionally, some employers may face contractual limitations on the use and collection of data due to agreements with unions or similar employee organizations.

The legal protections granted to employee information used in connection with an insider threat program are principally related to the type of employer, the nature and source of the data, and the manner in which the information is used. Implementers of an insider threat program should be particularly mindful of legal concerns associated with the circumstances and categories of information discussed in the following sections.

2. ELECTRONIC COMMUNICATIONS

Much of the analytic and telemetry used in connection with an insider threat program may be obtained through an employee's use of the employer's computer network. Gathering such information using electronic monitoring capabilities will likely implicate federal and state electronic surveillance laws and, in the case of government employers, the

Constitution. Such surveillance may also be subject to company-specific collective bargaining agreements, particularly when an insider threat program is newly instituted and subject to challenge as a substantial and material change in workplace conditions requiring collective bargaining.

The Wiretap Act (18 U.S.C. § 2510 et seq.), which prohibits the interception of oral, wire, and electronic communications unless conducted pursuant to a statutory exception or under a court order, may be implicated by content-based online monitoring such as payload-based network anomaly detection. Similarly, the Pen Register/Trap and Trace Statute (18 U.S.C. § 3121 et seq.) prohibits the use or installation of a device that captures, records, or decodes non-content electronic communications unless done under an exception or with a court order and may be triggered by the collection of non-content telemetry such as netflow. Most states have comparable applicable electronic surveillance laws. A government entity instituting an insider threat program will also need to ensure that its insider threat-related monitoring complies with the Fourth Amendment, which prohibits the government or state agents from conducting searches and seizures that are "unreasonable" and conducted without a warrant or under an applicable exception to the Fourth Amendment.

Notwithstanding these statutory and constitutional restrictions, an employer can lawfully monitor an employee's use of the employer's network pursuant to a statutory or constitutional exception, most often the consent exception. Consent to monitoring communications on an employer's computer system is typically obtained through a properly worded log-on banner, workplace policy, computer user agreement, training, or a combination of some or all of these methods. Other statutory and constitutional exceptions may apply as well.[6]

Employers should give careful consideration to the legal issues associated with any use or collection from employer-owned devices of information unrelated to the employee's work, including GPS information regarding an employee's location during non-work hours, personal contacts, or financial information stored on an employer-owned device. By ensuring that an employee's consent to monitoring of his or her use of the employer's network and equipment includes use and collection of work and non-work information, employers can minimize potential claims by employees of a reasonable expectation of privacy.

Further, employers should consider the legal implications of the use in an insider threat program of information collected from an employee-owned device used to access employer networks or data. Courts may find that employees have a broader expectation of privacy in devices that they purchase and for which they pay for internet and phone service. The

[6] More guidance on electronic surveillance issues can be found online in Department of Justice guidance. *See Searching and Seizing Computers and Obtaining Electronic Evidence in Criminal Investigations* (3d ed. 2009), available at http://www.justice.gov/criminal/cybercrime/docs/ssmanual2009.pdf; and Stephen G. Bradbury, *Legal Issues Relating to the Testing, Use, and Deployment of an Intrusion-Detection System to Protect Unclassified Computer Networks in the Executive Branch*, 33 Op. Off. Legal Counsel 1 (2009), available at http://www.justice.gov/sites/default/files/olc/opinions/2009/01/31/e2-issues.pdf.

extent to which an employer may monitor an employee's use of such a device to access employer networks or data likely will turn on the specific wording of company policies and the breadth and detail of employee consent to monitoring. This is an evolving area, impacted by varying federal and state laws regarding electronic communications and privacy and with different requirements applicable in other countries as well. Agreements with unions and other employee organizations also may impact any employer's use of information collected.

3. NON-TELEMETRY INFORMATION

This whitepaper focuses on insider threat programs that rely principally on electronic data derived from employees' use of an employer's computer network, equipment, and other facilities and as such, does not discuss the use of non-telemetry sources of information such as an employee's personnel file. However, legal considerations stemming from the use of non-telemetry sources are discussed here because an insider threat program may use such supporting sources of information to corroborate conclusions drawn by analytic tools.

Employers are required by law to keep certain information about employees, which is usually maintained in a personnel file e.g., employee history and performance appraisals. Information from a personnel file can provide helpful analytic data for use in an insider threat program.[7] However, personnel information is subject to federal and state laws intended to protect employees' privacy and civil liberties. Some states require employers to seek an employee's approval before employee records can be collected. In other states it may be illegal to disseminate personal information from a personnel file without an employee's consent.

There are specific types of employee information that receive special protections. A prime example is employee medical data. Federal laws, including the Health Insurance Portability and Accountability Act of 1996, the Americans with Disabilities Act (ADA), and the Family and Medical Leave Act,[8] require employers to treat medical records with confidentiality. This includes medical information, such as drug testing results, which may be germane to an insider threat program. Any record of medical history obtained from an employee health program and any other document relating to an employee's medical condition should be specially controlled, and its use for an insider threat program should be carefully considered by legal counsel. Under emerging state case law, employers who improperly handle medical information may be liable for the torts of intrusion and public disclosure of private data.

Many states have laws concerning employee personnel files; however, most focus on the rights of government employees rather than private sector employees or only impose restrictions on government employers. Similarly, federal law limits the type of information

[7] For example, a personnel file may contain employee attendance records; employee history, such as performance appraisals and corrective action write-ups; documents showing that the employee was informed of specific workplace policies or attended job-related training session; and employee separation information.
[8] This is not an exhaustive list of potentially applicable federal statutes protecting medical information from disclosure.

that federal agencies, the military and other government employers may keep on their workers but does not impose the same limitations on private sector employers.[9] State laws regulating private employers' handling of employee personnel files focus mostly on furnishing employees a right to review the content of a personnel file, correct erroneous material, or obtain copies of personnel documents. Similar to certain federal restrictions, some states also limit the content of information that an employer may keep in a personnel file. Federal or state law or employee agreements may also have restrictions on the use of employee information outside of a personnel file. In certain instances, these restrictions will not apply if an employee is adequately notified of and provides consent to the employer's use of his or her data; under some laws, such consent must be provided in writing. An organization may want to consider whether it can incorporate written notice to and consent by the employee during the onboarding process and refresh it periodically as necessary. Employers should review periodically whether their use of employees' information is consistent with the scope of employees' consent.

4. LAWS OF OTHER COUNTRIES

Multi-national organizations may face additional legal challenges implementing a company-wide insider threat program. There are significant differences between U.S. law and the laws of other countries concerning what constitutes personal information and how it may be collected and used in connection with an insider threat program. In some regions of the world, privacy and civil liberty laws are more restrictive and require precautions not required by U.S. law. In particular, the European Union's (EU) Data Privacy Directive and related regulations have a major impact on U.S. companies doing business in Europe. Other regions of the world have their own data privacy regimes with their own requirements, from mandatory data preservation to mandatory data destruction.

Such legal and privacy considerations may mean aspects of a multi-national organization's insider threat program will need to be implemented locally to satisfy different requirements that may exist under varying U.S. and the laws of other countries in which the organization does business. For instance, the disparate requirements for obtaining consent to monitoring under U.S. law and the EU's Directive may require localized administration of an insider threat program for a specific country. The best course for ensuring an organization's insider threat program satisfies the full range of legal concerns is to confer early and often with an organization's legal counsel and privacy officers through all phases of planning and implementing the program.

5. LEGAL RECOMMENDATIONS

An insider threat program can pose challenging legal issues that implicate U.S. federal and state laws, as well as the laws of other countries where an organization does business.

[9] The Privacy Act (5 U.S.C. § 552a) limits the type of personally identifiable information government entities may collect and disseminate.

These legal issues are generally surmountable, but because applicable laws differ from state to state and, in the case of multi-national organizations, from country to country, managing an insider threat program requires the active participation of an organization's general counsel's office and privacy officer to ensure that the collection, use, and disclosure of insider threat information comports with U.S. laws and, when necessary, the laws of other countries. The following recommendations will help an insider threat program stay on firm legal footing:

Provide employees with clear notice and obtain their consent for use, collection, and disclosure of personnel information in connection with an insider threat program, and insider threat monitoring of communications on the employer's network. Doing so will help address U.S. law issues including statutory restrictions related to the use of employee information and statutory and constitutional prohibitions on electronic monitoring. The employer should obtain consent in writing and maintain records of notice and consent in accordance with legal requirements.

Obtain advice on relevant laws of foreign jurisdictions prior to implementing insider threat programs outside the U.S.

Because improper use or disclosure of personal information can be detrimental to an employee and lead to potential lawsuits against the employer, institute appropriate safeguards to minimize the likelihood of legal claims and the risk of litigation. Such safeguards could include limiting an employee's supervisor's access to insider threat information that is not job related and restricting disclosure of employee information to anyone other than those who have a legitimate need to know the information.

If an insider threat program is contracted out, be familiar with the contractor's procedures for handling personal information and other legally protected data, and ensure that it is aligned with the organization's corporate policies and requirements.

APPENDIX B: ASSUMPTIONS

In order to focus on matters relevant to the scope of this whitepaper, various assumptions are made about the maturity of an organization's IT system. If these assumptions are not valid for a particular organization, it will be challenging to implement the recommendations without first addressing these gaps.

Those assumptions are:

- System security analysts have access to capabilities common to a mature Security Operations Center (SOC) that monitors enterprise security posture and alerts to common security incidents. This monitoring and alerting capability is usually managed by a Security Information and Event Management (SIEM) system.

- The organization has well-established inventory, configuration, policy, and identity management systems, and that an access control system with comprehensive logging of accesses to critical resources is available. This is an essential forensic tool to determine the sequence and origin of activities when an event is investigated.

For guidance on meeting these assumptions, please see the following:

The SANS website
https://www.sans.org/

NIST's Framework for Improving Critical Infrastructure Cybersecurity
http://www.nist.gov/cyberframework/upload/cybersecurity-framework-021214-final.pdf

The National Security Agency's Manageable Network Plan
https://www.nsa.gov/ia/_files/vtechrep/ManageableNetworkPlan.pdf

BIBLIOGRAPHY

Axelsson, S. (2000). The base-rate fallacy and the difficulty of intrusion detection. *ACM Transactions on Information and Systems Security, 3*(3), 186-205.

Cappelli, D. M., Moore, A. P., & Trzeciak, R. F. (2012). *The CERT Guide to Insider Threats: How to Prevent, Detect, and Respond to Information Technolgy Crimes.* Addison-Wesley Professional.

Casey, T. (2007, September). *Threat Agent Library Helps Identify Common Security Risks.* Retrieved from https://communities.intel.com/thread/49315: https://communities.intel.com/docs/DOC-1151

Casey, T. (2007). *Threat Agent Library Helps Identify Information Security Risks.* Intel.

Casey, T. (2015). *Insider Threat Field Guide.* Intel Corporation.

CERT Division. (n.d.). *Insider Threat*. Retrieved from CERT | Software Engineering Institute | Carnegie Mellon University: http://www.cert.org/insider-threat/

Guido, M. D., & Brooks, M. W. (2013). Insider Threat Program Best Practices. *2013 46th Hawaii International Conference on System Sciences* (pp. 1831-1839). Wailea, Maui: HICSS.

NIST. (2014). *Framework for Improving Critical Infrastructure Cybersecurity.*

Schulze, H. (2015). *Insider Threat Spotlight Report.* Crowd Research Partners.

Shey, H., Mak, K., Balaouras, S., & Luu, B. (2013). *Understand The State Of Data Security And Privacy: 2013 to 2014.* Forrester Research, Inc.

Silowash, G., Cappelli, D., Moore, A., Trzeciak, R., Shimeall, T. J., & Flynn, L. (2012). *Common Sense Guide to Mitigating Insider Threat, 4th Edition.* Software Engineering Institute. Software Engineering Institute. Retrieved from http://resources.sei.cmu.edu/library/asset-view.cfm?assetid=34017

GLOSSARY

Agent – a person who accidentally or maliciously takes steps to cause harm.

Analysis – the discovery, understanding, and communication of meaningful patterns or relationships; can include the use of analytics and manual discovery.

Analytic – automated process run against data to identify meaningful patterns or relationships in the data.

Analytic Indicator –an analytics' output that suggests the presence of an insider threat; may prompt decision making e.g., further analysis, analytic refinement, legal response.

External – originates from outside the organization.

Espionage – any effort to clandestinely obtain information from another party that would otherwise be kept confidential, in order to gain a competitive, economic, military, or political advantage.

Indicator – output of analysis that suggests the presence of an insider threat; may prompt decision making e.g., further analysis, analytic refinement, legal response.

Insider Threat – the potential for a current or former employee, contractor, or business partner to accidentally or maliciously misuse their trusted access to harm the organization's employees, customers, assets, reputation, or interests.

Insider Threat Program – as a concerted effort by an organization to detect insider threats and insider attacks. An insider threat program can be implemented via external, internal, or manual processes, or some combination thereof.

Internal – originates from within the organization.

Telemetry – data obtained from network components, security components, and other devices and applications (e.g., log files, network traffic) for the purpose of monitoring user activity.