CHAPTER 2

Big Data: Exposing the Risks from Within

Erica J. Briscoe

Erica J. Briscoe
Chief Scientist ATAS Laboratory, Georgia Tech Research Institute

A critical element in any institution is the existence of a trusting environment, which allows people to interact with one another without fear of adverse effects either on their professional or personal lives. Preservation of trust, however, is challenging. The rising number of threats to cybersecurity, fueled by an increasing reliance on data-driven devices, is coupled with a growing unease about the power that overseers tasked with ensuring that security (both corporate and government) possess as a result of their access. When taken in context with several high-profile cases of espionage, intellectual property (IP) theft, and workplace violence, both the private and public sectors are faced with a common challenge: How can institutions leverage technology to decrease their risks, especially those that involve malicious human behavior (such as insider threats)? This question cannot be answered without a careful consideration of how technology solutions affect those involved. How can these institutions minimize their vulnerability to threats, while maintaining an ethical, legal, and privacy-respecting environment? While there are no easy answers to these questions, recent research and security programs have shed some light on how a balance may be achieved, through a combination of technology and policy-driven solutions. Regardless of the responses devised to suffice today, given our increasingly automated world, institutions and the public will likely need to revisit this question continuously, ideally informed by both shared experiences and evolving research into human behavior.

Trust in Public and Private Sectors

The general concept of trust is not only complex, but its manifestation and characterization depend highly on the participating parties and the specific context in which trust exists. Whether considering individuals, governments, or machines (and all combinations thereof), there are several critical components[73] of trust. The first is that trust is made necessary

[73] Christel Lane and Reinhard Bachmann, eds., *Trust within and between organizations: Conceptual issues and empirical applications* (New York: Oxford University Press, 1998).

when one party's actions are consequential or require cooperation with another. The second is that relationships require risk (e.g., that a vendor will fulfill an order on time), which trust is used to mitigate. The third is that working together requires parties to become vulnerable, where trust ensures that one party does not take advantage of the other's vulnerability. Though these aspects are usually unavoidable, trust does not mean that an organization or entity must necessarily give their partners unrestricted access to information and sensitive resources; rather, successful institutional trust usually resides in a (sometimes delicate) balance between adequate security controls and acceptable risk. This balance is not static or well-defined, but requires comprehensive approaches that allow an organization to dynamically perform identity management and access controls, as well as flexible governance coupled with education and empowerment.

Though it is widely accepted that organizations require trust, each may engender different types, either intentionally or inadvertently. Lewicki and Bunker[74] outline three types of trust that are commonly found in work environments. *Deterrence-based* trust, as it uses reprisal to deter undesired behavior, is the most explicit and fitting for new institutional relationships or for those in an environment with low levels of information control. This type is often imposed through government agency or corporate policies, where the consequences for violations are clear and able to be imposed.

Knowledge-based trust requires that the involved parties have enough familiarity to be able to predict one another's behavior. This predictability reinforces the trust over time. Interestingly, even if one party is consistently untrustworthy (e.g., an employee often fails to clock in on time though there is an explicit policy that employees must be on time), the predictability of this behavior substantiates trust (in the belief that he will always be late). This type of trust may be relevant to organizational security in many aspects. Certain violations (such as being late to work) may serve as poor indicators of a person's malicious character (or lack of trustworthiness) if that behavior is consistently inconsistent (as later discussed relevant to detecting insider threats that behave anomalously). Changes in predictability (where behavior is increasingly anomalous) is a potential red flag for diminishing trustworthiness.

The last type, *identification-based* trust, involves one party acting as an agent for the other, serving as a substitute for that entity in interpersonal transactions. Trust of this type takes time and effort to build and often results in the most surprising and devastating responses when broken. Something akin to this type of trust is found in the relationships between the federal government and its contractors, who are often seen as acting on behalf of the government; however, rather than having that bond build through time and dedication, the trust is derived from intensive security screens and usually coupled with deterrence-based methods (which are questionably reliable given the recent high-profile security breaches, for example).[75]

Building a Trusted Environment

Early in 2013, President Barack Obama issued an executive order titled "Improving Critical Infrastructure Cybersecurity"[76] describing the need for the development of a voluntary cybersecurity framework to manage cybersecurity risks associated with critical infrastructure services. This order was the federal government's acknowledgement of the extreme vulnerability of many of the country's critical systems, as well as a call for organizations to develop and instantiate processes that effectively maximize and maintain trust within and between organizations.

The president's acknowledgement of cybersecurity risks coincides with a seemingly universal interest in harnessing the power of big data, that is, the ability to derive insights from the huge amount of information generated by the many computing devices that are used every day. Though the threats to information systems take familiar forms, including common criminals, disgruntled employees, terrorists, and dishonest business partners, potential indicators of these threats may be increasingly determined by recent developments in high-performance computing, machine learning, and new analytic techniques that leverage this large-scale data collection. This utilization, in addition to the increasing sophistication of potential threats, is feeding a common realization that traditional reliance on information technology (IT) specialists alone cannot protect an enterprise from malicious behavior. Organizations must focus not only on common technological solutions (such as password change policies), but also by leveraging advances in computationally driven methods that benefit

[74] Roy Lewicki and Barbara Bunker, "Developing and maintaining trust in work relationships," in Roderick Kramer and Tom Tyler, eds., *Trust in organizations: Frontiers of theory and research* (Newbury Park, CA: SAGE Publications, 1995), 114–139.

[75] Ellen Nakashima, Matt Zapotosky, and John Woodrow Cox, "NSA contractor charged with stealing top secret data," *Washington Post,* October 5, 2016, https://www.washingtonpost.com/world/national-security/government-contractor-arrested-for-stealing-top-secret-data/2016/10/05/99eeb62a-8b19-11e6-875e-2c1bfe943b66_story.html.

[76] White House, "Executive Order no. 13636, Improving Critical Infrastructure Cybersecurity, DCPD-201300091," February 12, 2013, http://www.gpo.gov/fdsys/pkg/FR-2013-02-19/pdf/2013-03915.pdf.

Passengers watch a television screen broadcasting news on Edward Snowden, a contractor at the National Security Agency (NSA), on a train in Hong Kong June 14, 2013. *Photo credit:* Reuters/Bobby Yip.

from the wealth of information that is produced by modern computing systems, both at the individual and network level. Additionally, most security experts agree that a comprehensive approach that integrates best practices across policy, technology, and people while building secure, transparent relationships is a necessary and effective security strategy.

Policies and Privacy

The extent to which an employer may monitor employees is dependent on a number of factors, including the ownership of the information systems, "what the state's laws and employer's policies are, what the employee's objective expectations of privacy are," where the employee is physically located, and "whether the employer has a legitimate interest in viewing the communication."[77] Protection of employee privacy has become a popular topic, which can be broadly classified into three types: statutes restricting unauthorized access or monitoring of data; health-related information (the Genetic Information Nondiscrimination Act, the Americans with Disabilities Act, the Family and Medical Leave Act, the Health Insurance Portability and Accountability Act); and statutes protecting personally identifiable information (PII), such as identity theft statutes, the Fair and Accurate Credit Transactions Act, and state data breach laws.[78] With the blending of work and personal lives (such as on social media) and increasing efforts to improve employee home and work life balance (e.g., by allowing employees to work from home), these issues are becoming more complex and salient.[79]

[77] "The Generation Gap...Tell me about it!" The Creative Network, Inc., accessed April 4, 2017, http://creativenetworkinc.com/blog/blog1.php.

[78] Karen McGinnis, "The Ever Expanding Scope of Employee Privacy Protections," *ACC Charlotte Chapter Q4 2014 Newsletter,* December 2014, http://www.mvalaw.com/news-publications-373.html.

[79] "The Generation Gap...Tell me about it!".

Table 2.1: Identified Insider Threat Types and Their Associated Behavior and Related Indicators.

Threat Behavior	Associated Activities	Behavioral Indicators
Espionage	Contact with foreigners Security violations Mishandling of sensitive information	Email, texting, social media Unauthorized access attempts, sharing passwords Unauthorized copying/downloading
Fraud	Theft of financial information Modification of sensitive information	Unauthorized copying/downloading Stress indicators, e.g., from financial hardship
Sabotage	Destruction or modification of sensitive information or software that will have detrimental results	Unauthorized access Communications exhibiting unprofessional behavior or grievances Stress indicators, such as from anger/resentment
IP Theft	Transmission of sensitive information Unjustified access to IP	Unauthorized copying/downloading Unauthorized access attempts Foreign or competitor contacts

Expectations on the type or level of trust and privacy may be set or influenced by explicitly stated policies (at the government agency or corporate level) and laws (at the state and federal level). Often these policies run up against privacy issues, where data collected on employees meant to ensure cybersecurity, for example, may not coincide with an individual's expectations of privacy. These issues are becoming more and more relevant as the world sees an explosion of "smart" devices. The prevalence of these devices allows for a much greater ability to see into the lives and behaviors of citizens and employees. At the extreme, the situation has become a case of big brother meeting big data, where, for example, China's use of the "Sesame Credit" scoring system means that all aspects of a citizen's life may be evaluated to determine his or her trustworthiness by keeping track of individuals' financial and consumer data.[80] Additionally, formal government agency or corporate policies that require employees to sign consent to monitoring as a condition of employment may set the tone of an environment of mistrust from the beginning of an employee's tenure at an organization. This impression, along with the anxiety that arises from an employee being aware that he is under constant surveillance, may be a catalyst for subversive and malicious behavior.

Insider Threats

Perhaps the most devastating case of a breakdown in trust occurs when an individual, who is part of an organization, uses his or her access for activities that are detrimental to that organization. These insider threats are often described as current or former employees or trusted partners within an organization that abuse (or have the potential to abuse) their authorized access to the organization's system.[81] As found in a recent survey conducted by *CSO* magazine, the US Secret Service, PricewaterhouseCoopers, and the Software Engineering Institute CERT, around 30 percent of electronic attacks on both public and private organizations came from the inside.[82]

[80] Celia Hatton, "China 'social credit,' Beijing sets up huge system," *BBC News*, October 2015, http://www.bbc.com/news/world-asia-china-34592186.

[81] Jeffrey Hunker and Christian W. Probst, "Insiders and Insider Threats-An Overview of Definitions and Mitigation Techniques," *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications*, 2011.

[82] Roger Parloff, "Spy Tech That Reads Your Mind," *Fortune*, June 30, 2016, fortune.com/insider-threats-email-scout.

Table 2.2: Example of Notable Insider Threat Cases

Insider Threat Case	Case Description	Incident Type	Threat Indicators
Xiang Dong Yu	In 2006, Yu, a product engineer for the Ford Motor Company with access to Ford trade secrets, accepted a new job at a Beijing-based automotive company that was a direct competitor of Ford. Before resigning, Yu copied 4,000 system design documents onto an external hard drive, which he later copied onto his new employer's computer.[a]	IP Theft	Email, texting, social media Unauthorized access attempts, sharing passwords Unauthorized copying/downloading
Tim Lloyd	In 1996, after being told he was fired, Lloyd planted a software "time bomb" in a server at Omega Engineering's Bridgeport, NJ, manufacturing plant. "The software destroyed the programs that ran the company's manufacturing machines, costing Omega more than $10 million in losses."[b]	Sabotage	Unauthorized copying/downloading Stress indicators, e.g., from financial hardship
William Sullivan	Discovered in 2007, Sullivan stole 2.3 million bank and credit card records from his employer, Certegy, a check processing company, including names, addresses, phone numbers, birth dates, and bank account information to sell.[c]	Fraud	Unauthorized access Communications exhibiting unprofessional behavior or grievances Stress indicators, such as from anger/resentment
Edward Snowden	Snowden worked as a US National Security Agency contractor who, in 2013, leaked a trove of documents about top-secret surveillance programs. He has been charged "in the United States with theft of government property, unauthorized communication of national defense information, and willful communication of classified [communications] intelligence."[d]	Espionage	Unauthorized copying, downloading

a. US Attorney's Office, Eastern District of Michigan, "Chinese national sentenced for stealing ford trade secrets," April 12, 2011, https://archives.fbi.gov/archives/detroit/press-releases/2011/de041211.htm.

b. Sharon Gaudin, "Computer sabotage verdict set aside," *Computer World*, July 12, 2000, http://www.computerworld.com/article/2596062/networking/computer-sabotage-verdict-set-aside.html.

c. Reuters, "Guilty plea in fidelity Nat'l data theft case," November 29, 2007, http://www.reuters.com/article/certegy-theft-idUSN2933291420071129.

d. Peter Finn and Sari Horwitz, "U.S. charges Snowden with espionage," *Washington Post*, June 21, 2014, https://www.washingtonpost.com/world/national-security/us-charges-snowden-with-espionage/2013/06/21/507497d8-dab1-11e2-a016-92547bf094cc_story.html.

The difficulties in preventing, detecting, and countering insider threats are an increasingly major task for information security professionals, highlighted most prominently in the United States by Edward Snowden's actions involved with leaking National Security Agency data. With the collection and analysis of big data, especially through corporate insider threat programs, it is likely that the prevention and detection of malicious activities are much more feasible than previously possible; however, with this potential, there remain many questions and areas for further research.[83] Advancement in this area is also met by multiple challenges, many arising from the difficulty in balancing expectations of privacy while maintaining a trust-maximizing environment.

Types of Insider Threats and Behavior

Based on the analysis of historical cases, several descriptive taxonomies have been developed to describe insider malicious activities. For example,

[83] Carly L. Huth, David W. Chadwick, William R. Claycomb, and Ilsun You, "Guest editorial: A brief overview of data leakage and insider threats," *Information Systems Frontiers 15*, 2013.

Phyo and Furnell's taxonomy[84] is based on the level(s) of information systems in which each type of incident may be detected or monitored. Internet-based activities are classified at the network level, while theft of sensitive information occurs at the operating system level. Nefarious interactions between users exist at the application level. This type of breakdown may be useful for creating a security strategy that applies to each level. Table 2.1 presents an overview of the most common types of insider threat behavior and the associated activities and indicators with each.

> **"The motivations behind insider threat behavior differ according to the specific individuals and their particular circumstances."**

While much attention has been given to prominent insider threat cases (see table 2.2), these individuals exemplify the rarest type of threat, that which results from intentional, directed malicious behavior. These malicious insiders possess the greatest potential to cause significant harm to an organization, especially because they are likely to try to hide or cover up their behavior, making them more difficult to detect. Exploited insiders are those who may be vulnerable to the influence of outside parties, such as through social engineering (the intentional social manipulation of individuals by adversarial actors to acquire confidential or personal information) or targeted spear phishing campaigns. Careless insiders are irresponsible with regard to security, and their accidental behavior may have detrimental consequences.[85]

Motivation and Indications for Insider Threats

Careless and exploited insiders are not malicious; rather, their actions result from lack of awareness, naivety, or lax security precautions. Malicious insiders are a much more thoroughly researched group, as they pose the greatest danger to organizations and often have complicated factors contributing to their behavior. Of course, the infrequency of these events makes it difficult to develop scientific studies into the variety of motivations for such behavior; however, case studies[86] show that analyzing individual psycho-social motivations and the developmental histories of formerly trusted insiders can lead to better insight into security vulnerabilities and preventative strategies.

Based on historical cases, Shaw et al.[87] suggest six personal qualities that may contribute to malicious insider behavior:

- "False sense of entitlement" or a "lack of acknowledgement" causing a "desire for revenge"
- "Personal and social frustrations, anger, alienation, dislike of authority and an inclination for revenge"
- Computer-focused, aggressive loners, intrinsically rewarded by exploring networks, code breaking, and hacking
- "Ethical flexibility lacking moral inhibitions that would normally prevent malicious" behavior
- "Reduced loyalty identifying more with their" job or tasks than with their employer
- "Lack of empathy or inability to appreciate the impact" of behavior on others

The motivations behind insider threat behavior differ according to the specific individuals and their particular circumstances. For example, the motivation for committing fraud may be more commonly due to financial reasons,[88] while espionage may be committed for ideological or narcissistic reasons. A common pattern for insider activity is that "attacks are typically preceded by high rates of stressful events including work-related and personal events," such as following employment suspension or termination.[89] Despite known patterns, many insider activities are discovered but never made public, in order for organizations to avoid any detrimental effect on their reputational or perceived security practices.

[84] William Cheswick, Steven M. Bellovin, and Aviel D. Rubin, *Firewalls and Internet Security: Repelling the Wily Hacker* (Boston: Addison-Wesley Longman Publishing Co., 2003).

[85] Russell Miller and Merritt Maxim, "I have to Trust someone…Don't I?," CA Technologies, 2015.

[86] Stephen Band, Dawn M. Cappelli, Lynn F. Fischer, Andrew P. Moore, Eric D. Shaw, and Randall F. Trzeciak, "Comparing insider IT sabotage and espionage: A model-based analysis," (Pittsburgh, PA: Carnegie Mellon University, 2005).

[87] Eric D. Shaw, Jerrold M. Post, and Kevin G. Ruby, "Inside the Mind of the Insider," *Security Management 43*, 1999.

[88] Adam Cummings, Todd Lewellen, David McIntire, Andrew P. Moore, and Randall Trzeciak," Insider threat study: Illicit cyber activity involving fraud in the US financial services sector," (Pittsburgh, PA: Carnegie Mellon University, 2005).

[89] Stephen Band, Dawn M. Cappelli, Lynn F. Fischer, Andrew P. Moore, Eric D. Shaw, and Randall F. Trzeciak. "Comparing insider IT sabotage and espionage: A model-based analysis" No. CMU/SEI-2006-TR-026, Carnegie-Mellon University, Software Engineering Inst, 2006; Andrew P. Moore, Dawn M. Cappelli, and Randall F. Trzeciak, "The "big picture" of insider IT sabotage across US critical infrastructures," In *Insider Attack and Cyber Security*, (Santa Clara, CA: Springer-Verlag TELOS, 2008), 17-52.

Insider Threat Detection and Prevention

Security measures, such as data-loss prevention software, database activity, and network traffic monitoring programs, as well as security information event management systems, provide organizations with basic defenses, but do not much help to identify and prevent damage from insider threats. Although enterprise-wide defenses are becoming more sophisticated, the human aspect of security remains a weak link. A study of insider threat cases by the Computer Emergency Response Team (CERT) Insider Threat Center, a federally funded research and development entity at Carnegie Mellon University, found that 27 percent of insiders who became threats had drawn the attention of a co-worker because of his/her behavior prior to the incident.[90] These reports provide good support for the development of methods and systems that monitor individuals' *behavior* to detect and alert security professionals when their behavior first becomes detrimental or otherwise abnormal.

The benefit of focusing on user behavior has recently resulted in the incorporation of user behavior-focused methods as a critical component of many current enterprise systems that work to maximize cybersecurity. This often involves applications that monitor user behavior across multiple networks.[91] For example, users' computers may run an application that collects behavioral traces, which are then batched and sent to a central server to be processed at specified intervals (usually daily). The central server will also correlate and fuse information to create risk scores, which are more easily visualized and communicated to non-expert users, such as the managers who must assess the threat on a personal level.

Technical approaches for the continuous monitoring of insider behavior vary. The most straightforward method involves the direct identification of malicious activity, using what is referred to as rule-based detection, where observed events are matched against known models of threatening behavior. For example, a known threatening behavior may be the activities associated with a user accessing files that are outside of his security clearance level. While these approaches are likely to result in accurate detections, they require precise identification of the behaviors, which means that only previously *known* types of attacks will be detected.

Another clever approach that is relatively straightforward is through the use of *honeypots*. A honeypot is some type of digital asset (such as a file) that is put on a network specifically so that it can be monitored. Because the honeypot has been created to test for malicious behavior, no users should have a legitimate use for it (though it is often made to look attractive to would-be threats). This means that any interaction with the honeypot, such as a rogue user accessing it, is, by definition, suspect.

A group of much more computationally sophisticated methods use anomaly detection, which focuses on discovering rare activities within a large corpus of observation data. When considered from the perspective of an organization, the vast majority of user activities are normal and the insider threat actions are outliers.[92] Within the outlier set, insider threat activities represent an even smaller set of actions; the task is then identifying this subset of outlier actions.[93] At best, a successful insider threat detection capability would result in the identification of the actions that correspond to truly threatening behavior, but given the inherent ambiguity in determining threatening behavior, an intermediate success is the paring down of the dataset so that a human may reasonably comprehend it.[94] A successfully implemented system would allow, for example, security personnel to produce a report that would show which employees in the organization were the most anomalous or even disgruntled,[95] which may, in turn, provide an opportunity for early intervention or an increase in security measures.

Anomaly detection approaches usually require three components. First, information that represents "normal" behavior must be collected and stored. This could be employees' daily logs on activity or file accesses, for example. This information becomes the training data on which behavioral norms are modeled using a variety of machine-learning approaches, such as Markov models, support vector machines, or neural networks. Once these models of normal behavior are created (and, ideally, frequently updated), each individual's regular activity is monitored and compared against the model to determine if significant deviation occurs, which

[90] Marisa Randazzo, Michelle Keeney, Eileen Kowalski, Dawn Cappelli, and Andrew Moore, Insider threat study: Illicit cyber activity in the banking and finance sector, No. CMU/SEI-2004-TR-021, Carnegie-Mellon University, Software Engineering Institute, 2005.

[91] Splunk, "Machine Learning Reveals Insider Threats," last accessed March 20, 2017, https://www.splunk.com/en_us/products/premium-solutions/user-behavior-analytics/insider-threats.html.

[92] David B. Skillicorn, "Computational approaches to suspicion in adversarial settings," *Information Systems Frontiers 13*, 2011.

[93] Rudolph L. Mappus and Erica Briscoe, "Layered behavioral trace modeling for threat detection," International Conference on Intelligence and Security Informatics, 2013.

[94] Scott Shane and David E. Sanger, "N.S.A. suspect is a hoarder. But a leaker? Investigators aren't sure," *New York Times*, October 6, 2016, http://www.nytimes.com/2016/10/07/us/politics/nsa-suspect-is-a-hoarder-but-a-leaker-investigators-arent-sure.html.

[95] Roger Parloff, "Spy tech that reads your mind."

Table 2.3: Example Anomaly Detection Methods with Associated Elements

Method Elements	Method 1	Method 2
Method Type	Cross-sectional	Temporal
Entity Comparison	Individual user	Users
Baseline Population	All users / groups	Users
Baseline Feature(s)	Number of emails per day	URLs visited each day
Baseline Feature(s) Distribution	Normal (mu, sigma)	Vector of URL counts
Baseline Time Period	N/A	Last six months
Degree of Difference	Number of standard deviations from mean	Vector distance

may trigger an alert, for example, to signal a human supervisor for further investigation. Table 2.3 outlines two examples of anomaly detection methods and their distinguishing elements. Method one determines the difference in email volume between an individual user and his or her peers at one point in time compared to their average behavior over the past year. Method two compares the previous Internet activity (by creating lists of websites visited) of each user with more recent activity of that user. The primary difference between the two methods is that method one determines anomalies by comparing users to other users, while method two evaluates how a particular user changes his or her behavior over time. Comprehensive approaches that include this type of variability in methods is necessary for catching the variety of potentially malicious anomalies that may occur.

Though these detection methods usually focus on detecting deviations in normal computer usage activity, early detection methods may also concentrate on finding more subtle changes in user behavior that arise from either personal stress (which may be the motivation for becoming a threat) or the stress associated with a user knowingly committing an illegal act. The variability in a person's response to stress depends on various factors, including individual differences and the situation in which that response takes place. The effect of stress on performance can be seen as a continuum, ranging from no effect to a significant degradation in performance (e.g., the person makes errors or inadequately slow responses). This resulting change in behavior due to stress is another potential source for anomaly detection methods. Additionally, though most anomaly detection systems currently concentrate on passive detection of these types of indicators or "tells," new government research is evaluating whether these passive detectors can be combined with active indicators—those that arise from specific, intentional stimuli.[96]

While corporations are usually limited to user data collected while their employees are on corporate-owned devices, recent government employee insider threat incidents have emphasized the need to incorporate external data sources as well. This need is exacerbated by the potential detrimental effects that these employees can have with their access to highly classified information. While these workers are required to undergo fairly intensive background checks of both their financial and private lives, notable recent cases, such as that of Aaron Alexis, a former Navy reservist and military contractor who killed twelve people at the Washington Navy Yard in 2013, highlight the potential inadequacy of traditional background checks and lack of agency coordination. Former Director of National Intelligence James Clapper told Congress that what is needed is a "system of continuous evaluation where when someone is in the system and they are cleared initially, then we have a way of monitoring their behavior, both their electronic behavior on the job as well as off the job."[97] This type of employee monitoring systems might access multiple data sources in an attempt to discover patterns of suspicious behavior not caught by traditional background checks, which may be appropriate given the potential vulnerabilities for national security but seem much too invasive for ordinary citizens and employees. Examples of external sources include "private credit agencies, law enforcement databases and threat lists, military and other government records, licenses, data services

[96] GCN Staff, "IARPA preps insider threat monitoring projects," *GCN*, March 19, 2015, https://gcn.com/articles/2015/03/19/iarpa-scite-insider-threat.aspx.

[97] Stephen Braun, "U.S. intelligence officials to monitor federal employees with security clearances," *PBS News Hour*, March 10, 2014, http://www.pbs.org/newshour/rundown/us-intelligence-officials-monitor-federal-employees-security-clearances/.

and public record repositories,"[98] and social media, in addition to potential electronic surveillance.

Challenges

Regardless of the type or number of sources used, there are several challenges to using analytic methods to detect insider threats.[99] Of course, most malicious insiders do not wish to be detected; therefore, they try to hide their detrimental actions by concealing them within legitimate activity. This concealment makes detection much more difficult even for advanced anomaly detectors. Most algorithmic approaches also require training data, which consist of labeled cases of both known "normal" and nefarious behavior; however, the collection of these sets is difficult due to the rarity of cases and the reluctance of government agencies and companies to share information regarding their identified vulnerabilities. In application, the ratio of "bad" to "good" users in an organization is extremely low, which makes for few opportunities to test the effectiveness of implemented approaches. Given a large number of employees and multiple data sources, reducing a mass amount of information down to simplistic measures, such as risk scores, may still result in too much information for a person to process, making continuous monitoring ineffective.

Preventative Measures

While insider threat detection programs are growing more sophisticated, so should approaches that concentrate on the individual *before* he or she starts down the criminal path. These techniques probably best address the *careless* and *exploited* threat types, but may also deter *malicious* insiders by increasing the visibility of an organization's security presence. Increasingly, effort is invested in the development of security awareness and risk communication programs to raise computer users' awareness about practicing safe habits and recognizing security threats. Communications usually take five forms: warning dialogues, notices, status indicators, training, and corporate policies. These programs may also be informed by massive data analytics, usually through large-scale testing and analysis that helps to pinpoint who the most vulnerable users are.

Because malicious attacks can take many forms, so must preventative training. A growing body of research shows that there are several useful factors to a successful security awareness campaign. As one example, studies show that highly self-referencing messaging, such as those using wording that focuses on the specific individual or their personal data, is more effective than appealing to the community or corporation. A message "Protect your personal data by changing your password every month" is likely to be more effective than "IT policy requires a password change to increase cybersecurity." Also, research demonstrates that perceived threat severity can have a negative impact on self-efficacy, which is the belief that one is capable of taking effective actions to avoid the threat. These findings suggest that security messages should include references to the user and information to increase self-efficacy beliefs.

> **". . . Insider threat detection programs are growing more sophisticated. . ."**

Evaluation of the effectiveness of security awareness campaigns often take the form of mimicked attacks initiated by security management. Subsequent security awareness messages after these "tests" are likely to be particularly effective, as users are immediately made aware of their risky behavior. With precise test construction, it is possible to ascertain exactly what attack methods are likely to result in security breaches.[100] This information, along with observed user responses, can then be used to target future messaging, campaigns, and/or training. This is more nuanced than merely understanding what types of threats people are more likely to succumb to, but which characteristics of those threats influence the users' perceptions and actions. For example, "normal" security indicators, such as a padlock icon, often go unnoticed and, therefore, serve little purpose.

Looking to the Future: Trust in an Increasingly Automated World

Traditional sources of institutional trust are usually found in the relationships that exist between employers and employees or citizens and their government, but as humans become more technology-reliant, socio-technical trust, which results from the complicated interactions between people and technology,[101] is a significant aspect in everyday life. Given the recent advances of and attention to autonomous systems, the topic

[98] Stephen Braun, "U.S. intelligence officials to monitor federal employees with security clearances."

[99] Amos Azaria, Ariella Richardson, Sarit Kraus, and V. S. Subrahmanian, "Behavioral analysis of insider threat: a survey and bootstrapped prediction in imbalanced data," *IEEE Transactions on Computational Social Systems* 1, No. 2, 2014.

[100] Ronald C. Dodge, Curtis Carver, and Aaron J. Ferguson, "Phishing for user security awareness," *Computers & Security 26*, February 2007.

[101] Albert Bandura, "Social cognitive theory: An agentic perspective," *Annual review of psychology 52*, 2001.

NCCIC Classification Level:
UNCLASSIFIED
U.S. DEPARTMENT OF HOMELAND SECURITY
NCCIC
Akamai
Akamai global internet observer
Ping Latency
BGP Churn
Darknet

US Department of Homeland Security employees work in front of US threat level displays inside the National Cybersecurity and Communications Integration Center. *Photo credit:* Reuters/Kevin Lamarque.

of human-machine trust has risen to prominence in recent years[102] and will continue to increase as automation becomes more ubiquitous, requires less human involvement, and is increasingly relied upon throughout society.

Although there is an abundance of research that suggests that trust is the appropriate concept for describing human-machine interaction, there are several notable differences between that and what is understood about human-to-human trust. The most notable is that machines (even with their increasing personalization, e.g., Amazon's Echo) lack intentionality, which is a necessary component for other trust-inducing characteristics, such as loyalty, benevolence, and value congruence.[103] The asymmetry between humans and machines negates typical social cues and expectations, which in turn causes people to trust and react to machines in a dissimilar manner than they do to other humans. The facilitation of trust between humans and machines is currently most focused on the appropriate design of interfaces; however, with the increasing complexity of artificial intelligence, interface design alone is still insufficient to establish the trust that is necessary for humans to put their faith in automation. This is leading to research into how to open up the "black box," where transparency in the computational reasoning behind a machine's behavior is expected to increase the human's trust in it.[104] This transparency may be difficult in many cases, especially when the machine's reasoning mechanisms utilize representations that are not

[102] Lee Hutchinson, "Four hundred miles with Tesla's autopilot forced me to trust the machine," May 22, 2016, http://arstechnica.com/cars/2016/05/four-hundred-miles-with-teslas-autopilot-forced-me-to-trust-the-machine/.

[103] John Lee and Katrina A. See, "Trust in automation: Designing for appropriate reliance," *Human Factors: The Journal of the Human Factors and Ergonomics Society 46,* 2004.

[104] Davide Castelvecchi, "Can we open the blackbox of AI?," Nature *538*, 2016.

human interpretable (such as deep learning networks).[105]

As fallible and risky as human behavior is, it is certainly not a given that machines are (or will be) much better. Their risks are similar to those assumed with humans, in that detrimental behavior may arise from both intentional and unintentional actions (where software bugs or hacks may cause a machine to behave unpredictably or maliciously). As technology improves, machines will become "smarter" and more social, able to communicate among themselves (creating the so-called Internet of Things),[106] and therefore less likely to require "humans in the loop." These decentralized systems, those that are not monitored by a single executive function and that have no prior knowledge of one another (but are flexible and scalable), are potentially ripe for malicious behavior. Recent approaches for managing the inherent risk within these types of systems have been inspired by other human-based techniques, such as the use of reputation.[107]

Research has found that while consumers are aware that their data are being collected on a continuous basis, they do not necessarily understand the specifics or motivations behind that collection. This lack of understanding is a source of anxiety.[108] Studies on consumer-based data have found that transparency about the use and protection of consumers' data reinforces trust, but that this trust varies across the identities of the collectors.[109] Internet-based finance firms, such as PayPal, are generally perceived to be the most highly trusted, followed by e-commerce companies, consumer electronics makers, banks and insurance companies, telecommunications carriers, large Internet companies (e.g., Google), and the government. Interestingly, retailers and entertainment-focused companies were the lowest trusted organizations, rated above only social networking sites, such as Facebook. These findings point to the fact that both government and private institutions should aspire to increase their levels of transparency in order to counteract feelings of mistrust and anxiety that may accompany necessary cybersecurity programs.

Recommendations

The following actions are recommended to create a secure yet trust-respecting environment.

- Efforts and policies toward protecting personally identifiable information may assuage some of the fears that collected information could be used to negatively affect employees (a legitimate fear given recent corporate and government data breaches). PII protection policies may include encrypting employee PII, maintaining adequate firewalls and anti-virus software, avoiding use of employee social security numbers as means of employee identification, running an adequate record retention program, and employing measures that ensure business partners who access data also employ similar processes.
- Tools to manage access to data and personal information require the right balance of permissiveness and monitoring, achieved through fostering accountability, continuous training, security procedures (such as user monitoring), and control mechanisms. No matter what the strategy, communicating the intent of both security and privacy-respecting processes will provide people with more confidence in their employers and government. Balanced programs involve monitoring both known threats and user behavior concurrently, so as to quickly inform users to new threats and to augment methods used to assess those threats. This approach will pave the way for a unified approach (both human- and enterprise-focused) to information security.
- Institutions need to foster a cybersecurity mindset that is capable of continually adapting to counter changing threats. This mindset is likely best attained through a leadership-driven cybersecurity culture throughout the enterprise that results in shared "digital trust." Therefore, the responsibility for maintaining this trust not only lies with those in an organization tasked with monitoring information systems (such as that found in a security operations center—SOC—a group within an organization whose mission is to continuously monitor and improve an organization's security posture while preventing, detecting, analyzing, and responding to cybersecurity incidents with the aid of both technology and well-defined

[105] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, "Deep learning," *Nature 521*, 2015.

[106] Jayavardhana Gubbi, Rajkumar Buyya, Slaven Marusic, and Marimuthu Palaniswami. "Internet of Things (IoT): A vision, architectural elements, and future directions," *Future Generation Computer Systems 29*, 2013.

[107] Euijin Choo, Jianchun Jiang, and Ting Yu, "COMPARS: toward an empirical approach for comparing the resilience of reputation systems," *Proceedings of the 4th ACM conference on Data and application security and privacy*, March 3–5, 2014.

[108] Timothy Morey, Theodore Theo Forbath, and Allison Schoop, "Customer data: Designing for transparency and trust," *Harvard Business Review 93*, 2015, https://hbr.org/2015/05/customer-data-designing-for-transparency-and-trust.

[109] Ibid.

> processes and procedures), but extends from the top leadership to the entire workforce.[110]

Governments and other organizations that implement insider threat programs should be transparent and make clear to their workforces what types of personnel data and activities they monitor to help identify insider threats with the intent of protecting the workforce, sensitive information, and the viability of the organization itself.

Conclusion

Organizations must place trust in each employee that accesses sensitive data or systems; however, a well-trusting environment does not mean that users have unrestricted access to information or that an institution must accept massive amounts of risk. By analyzing employees' cyber footprints as well as non-IT–based behavioral indicators, organizations may have a more complete picture of potential risks. To ensure a healthy and trusting environment requires that institutions facilitate a cultural norm around security; one that includes high levels of transparency and standardization and that is capable of adapting to evolving threats, including non-human ones.

[110] Pierluigi Paganini, "What Is a SOC (Security Operations Center)?" *Security Affairs*, May 24, 2016, http://securityaffairs.co/wordpress/47631/breaking-news/soc-security-operations-center.html.