NJVC
Driven by Your Mission®

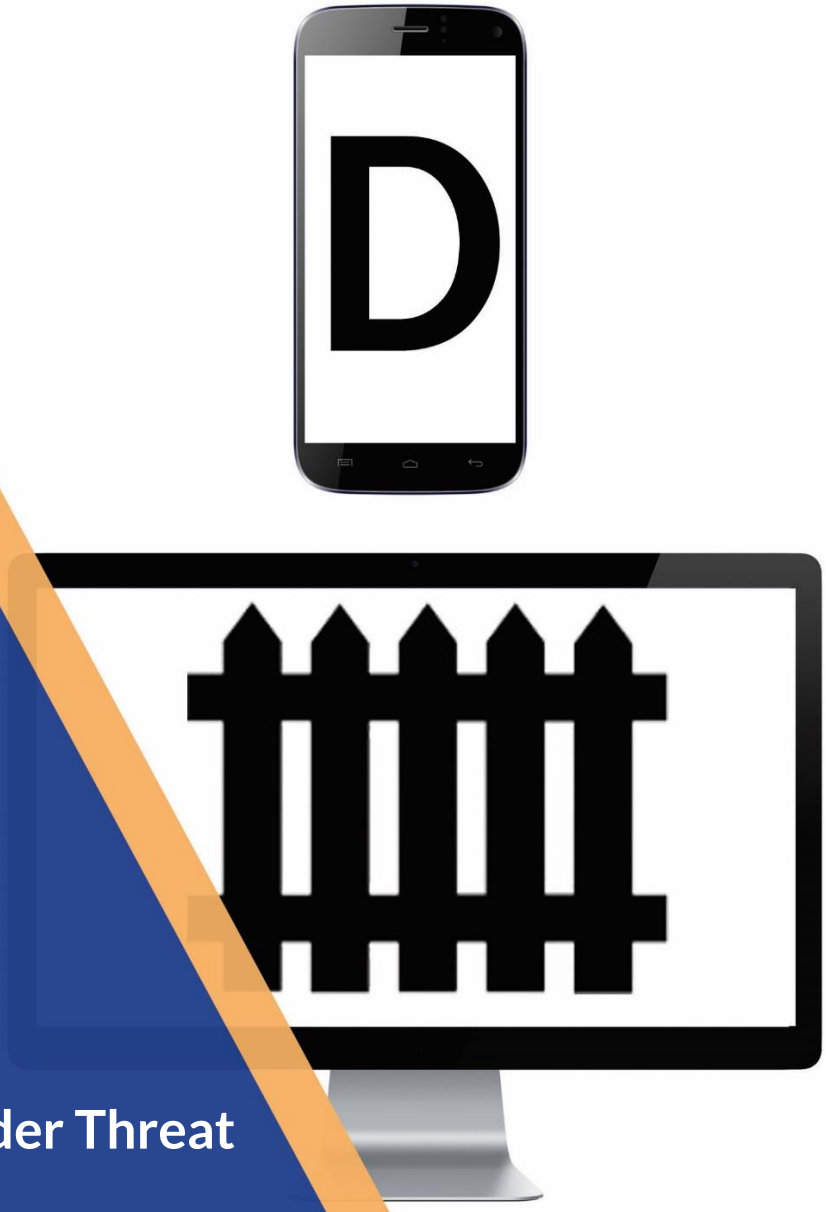

Detect and Deter

Playing Defense Against Insider Threat

Robert J. Michalsky
Principal, Cyber Security

September 2015

Learn more at NJVC.com

Contents

Defining the Threat

In cyber security, threats abound, system vulnerabilities are numerous and news of data breaches are as common as thunderstorms in summer. And yet, in this environment struggling to balance risk and privacy, the insider threat is particularly pernicious.

Insiders are individuals trusted to protect organizational secrets and intellectual property. As insiders, they are typically given privileged access and account privileges to carry out their position responsibilities. Since they have the greatest access, they also pose the greatest risk. Abuse of their privilege, most often for financial gain, can be as damaging to your organization as it is difficult to uncover.

What cuts to the bone in these situations is the violation of trust. Individuals have typically passed a background check or, over time, have proven themselves worthy of special trust; then that trust is tossed away.

As a general rule, people – not the underlying technology - are the biggest security problem organizations face. Individuals are human and therefore make errors with corporate data, forget security rules, overlook organizational policies and expose protected data. These actions can be either accidental or intentional. Both result in data exposures but malicious activity usually carries greater negative impact.

Two incidents highlight the extreme amount of damage that can be caused when insiders go rogue. In 2010, Pvt. Chelsea Manning leaked 251,000 classified and sensitive-but-unclassified diplomatic cables. These cables described in detail events which took place in 274 embassies over a 44-year period. Many unguarded conversations on nuclear disarmament, the war on terror and sensitive interactions with foreign countries were disclosed, causing harm and embarrassment.

There were also documents such as military logs and videos of military hardware. In total, the disclosure to WikiLeaks exceeded 720,000 documents. Manning received a sentence of 35 years in prison for his actions.

In 2013, NSA contractor Edward Snowden stole an unknown quantity of documents; more than 100,000 were leaked to journalists. The volume could be substantially larger since Snowden had access to over one million documents in the course of his duty. Whatever the actual number, it is dwarfed by the sensitivity of information he conveyed to foreign sources. This is the very definition of abuse of privilege conducted by a trusted insider.

Gen. Keith Alexander, former director of the National Security Agency and former commander of US Cyber Command, identified successfully mitigating insider threat as the No. 1 lesson to be taken from the Snowden incident.

So can insiders like these be stopped? Determined adversaries who have privileged access and understand the internal security controls in place will always be the most difficult cyber security challenge.

One approach is to redefine success and reframe the expected outcome. After all, a security incident is not a data breach until data actually leaves an organization. This point is often missed by decision makers. Simply because malware of some sort is discovered inside protected network boundary walls, does not mean organizational assets have been compromised. An investigation may be warranted, and a subsequent forensic analysis may be conducted, but not all security incidents lead to a data breach.

A successful mitigation approach requires attention to all three of these aspects: people, process and technology. Consider:

- Having a security policy – but not following it – may lead to data breaches
- Having great people – but not monitoring their actions – may lead to data breaches
- Having the latest cyber software vendor tools in place – but not regularly analyzing the resulting alerts – may lead to data breaches

Organizational focus is often centered on line items contained in a budget such as firewalls, Intrusion Detection Systems (IDS) and Intrusion Prevention Systems (IPS). Budget items can be simple to identify and quantify, but perimeter defense hardware and software systems such as these typically do not offer much protection against an adversary already inside network borders.

Background

There is no single master definition of the term insider threat. The Department of Homeland Security (DHS) provides a working definition as "a current or former employee, contractor, or other business partner who has, or had access to, an organization's network, system or data and intentionally misused that access to negatively affect the confidentiality, integrity or availability of the organizations information or information systems."

That may be a mouthful, but in essence, a trusted individual violates some measure of trust and impacts the holy triad of Information Security (InfoSec): Confidentiality,

Integrity or Availability (CIA) of data. This CIA approach is used by organizations such as the International Information Systems Security Certification Consortium ($ISC^{(2)}$) to organize and structure their respected cyber professional certification programs.

The Computer Emergency Response Team (CERT) at Carnegie Mellon University has performed research on data breaches since 2008 and has compiled a database of more than 700 cases. They further define insider activity into four types:

1. IT sabotage – Use of IT to direct specific harm at an organization or individual
2. Theft of IP – Use of IT to steal Intellectual Property (IP) from an organization. This also includes industrial espionage where outsiders may become involved.
3. Fraud – Use of IT for the unauthorized modification, addition or deletion of an organization's data for personal gain, theft of information that leads to identity crime (e.g., identity theft or credit card fraud)
4. Miscellaneous – All other insider threats

These groupings are not intended to be mutually exclusive as many real world recent data breaches fall into multiple areas, but the categorization provides an organized outline for detailed research. It also provides value to an organization by allowing a perspective beyond an often overwhelming amount of reported IT vulnerabilities.

In addition, thinking of these attacks based on what an organization has to lose, allows a prioritization to occur on security defense mechanisms. Organizations should prioritize their cyber mitigation implementation based on the projected costs of each breach type.

According to the Carnegie Mellon study, financial reasons were the primary motivation in 81 percent of insider cases. Consider corporate assets and how they can be monetized. Don't simply look at obvious high value items such as password files of Electronic Health Records (EHRs). Consider data items that will have value to others based on the groupings above.

Don't overlook non-financial situations. Twenty-seven percent of the cases involved sabotage either of the business or data or networks or operations as a motivating factor. These situations can often be even more disruptive to organizations if online systems are impacted or system outages incurred.

Organizational Concerns

Concerns about insider threat abound across every industry sector, all of which have individuals in positions of trust to operate and maintain their business, conduct their

operations and protect their data assets and intellectual property. Since IT is so integral to every business in the digital age, trusted individuals can bypass existing physical and logical controls intended to provide perimeter defenses against outsiders and external threats.

Organizational concerns typically fall into four categories: monetary, reputational, regulatory and legal.

Monetary

Financial loss is a result of almost any security incident. Forensic analysis may be required to ensure any infections are eradicated and provide insights to prevent future incidents. Customers may have to be notified and reimbursed. Additionally, there may be industry fines and costs to third-party firms to conduct audits or investigations.

These costs show up as expenses for publicly traded companies. Public companies may also see their stock price fall as a result. However, large ongoing concerns such as Sony and Target, who both experienced substantial data breaches, saw their stock prices recover after news of the negative event was ingested and overcome.

Still, data breach expenses can become substantial. Target cited more than $252 million in total breach costs as a result of their 2013 incident. Whether the amounts become material to their stock price, these are expenses no one wants to incur.

Reputation

Reputational loss is often cited by Chief Information Security Officers (CISOs) as their primary organizational concern. Loss of brand value is intangible and notoriously difficult to calculate or even estimate. Loss of customer loyalty can be a slowly eroding situation with both short-and-long-term implications.

In addition, in any competitive industry, corporate espionage is a valid concern. Once lost, data cannot easily be recovered as with a physical asset. Espionage can take many forms. Governments and military worry about the loss of classified information through nation states, terrorists and other sophisticated attacks. Healthcare providers worry about the loss of EHRs, which hold private patient medical data.

Also, it is unfortunate, but all too often with cyber crime it is the victim that takes the heat, particularly when the crime is perpetrated by an insider. Questions emerge about techniques in place, products, people and so on. Were sufficient protections in place? What is “sufficient?” Why was this allowed to happen?

Regulations

Virtually every industry is bound by regulatory concerns. Non-compliance can lead to fines, increased governmental oversight and legal actions from disgruntled suppliers, customers or both.

Customer notification has costs. It has almost become standard practice to offer identity or credit protection to individuals impacted by the breach.

Legal

One thing is certain, legal fees will not diminish after a security incident or data breach. Moreover, some existing legal agreements are likely compromised with a data breach.

New litigation costs may be incurred through combatting lawsuits. Target faced more than 40 lawsuits filed in courts all across the country in the months after their October 2013 breach. Often, these suits point to a company not implementing sufficient data protections and contending that is what led to the breach. Attorney fees accumulate even if the lawsuits do not end up in court.

From a legal standpoint, alteration of data - even if it remains on a local system - can be as disruptive as data breaches where data has been transmitted outside controlled network domains. Organizations have reported having financial bonuses altered, dates being modified and even data labels being removed, allowing easier data exfiltration because controls on key data items have been removed. Any unauthorized modification or transmission of protected data is most probably illegal.

Even if all these costs are not a material impact to an organization, money spent on breach cleanup is money that would have been better spent in breach prevention. No organization wants to make headlines as an example of how not to conduct cyber security operations.

Understanding the Landscape: Current Survey Findings

Surveys of IT personnel routinely point out either a strong bias towards optimism or a shocking disregard for reality. The SANS Insider Threat Survey[1], conducted from December 2014 to January 2015 had 772 respondents, a substantive sample size. Only 2-3 percent of respondents stated it took seven-to-12 months to detect an attack. Yet the

[1] https://www.sans.org/reading-room/whitepapers/analyst/insider-threats-fast-directed-response-35892

median number from the 2014 Verizon Data Breach Investigations Report (DBIR)[2], the largest sample size of actual customer data, was 229 days.

Thus, while very few IT personnel admitted it took more than half a year to uncover a compromise, the Verizon data, based on their real world case studies, indicated that fully half of their investigations took that long. Further, the Verizon results are based on a sample size of over 79,000 security incidents in just the 2015 report[3] alone. Cold water, therefore, needs to be thrown around.

Surveys must be read with a discerning eye given the very human tendency for optimism – or perhaps in this case overly wishful thinking. Some situational reality is present, however, as only 25 percent of respondents said they detected an attack in under eight hours which is at least some acknowledgement of the difficulty involved.

The risks however are borne out in the fact that Verizon also reports that 60 percent of attackers are able to compromise an organization within minutes. The trend is not encouraging as their data indicates that even for those compromises discovered within days, the average compromise occurs faster.

A March 2015 BeyondTrust Software white paper[4] indicated that fully 82 percent of respondent companies were aware that a data breach is an organizational issue, not simply IT, and looked at the necessary controls being cross-functional in nature and requiring support from business units as well as the IT group.

In that same survey, 47 percent of organizations acknowledged that there were individuals with elevated privileges beyond those required for their role. This is a serious cyber security risk. It violates the principle of "least privilege," which says individuals should be given the minimum set of privileges for the role they are performing. BeyondTrust reports 33 percent of companies have no privileged account management

Insider Threat Statistics

60% of attackers compromise anorganization within minutes.

- Verizon 2015 DBIR

47% of organizations acknowledge insiders have privileges beyond those required to their roles.

- BeyondTrust Software

70% of insider threats exploited vulnerabilities in business procedures, not software.

- BeyondTrust Software

85% of insider threat cases had at least one other person who was aware of the planned activities of the insider.

- Software Engineering Institute of Carnegie Mellon

[2] http://www.verizonenterprise.com/DBIR/2014/reports/rp_Verizon-DBIR-2014_en_xg.pdf
[3] http://www.verizonenterprise.com/DBIR/2015/
[4] http://www.beyondtrust.com/Content/whitepapers/Privilege-Gone-Wild-2015.pdf

policies. Since every organization that has an IT structure has privileged accounts by definition. That is a serious liability when it comes to handling any threats, let alone difficult to detect situations such as an insider threat.

The Software Engineering Institute (SEI) of Carnegie Mellon has looked at this issue in some detail. A 2012 Department of Homeland Security sponsored study by the Carnegie Mellon SEI titled "Common Sense Guide to Mitigating Insider Threats"[5] found 70 percent of insider threat cases exploited vulnerabilities not in software products, but in business procedures or processes. These typically required no special technical skills, but rather an exploitation of user role or software applications where attackers already had authorized access.

81 percent of the incidents were planned in advance – indicating actions being attributed beyond a simple target of opportunity (such as a server log in session left open or passwords posted in full view on a monitor screen). In fully 85 percent of the cases Carnegie Mellon SEI studied, someone other than the perpetrator was aware of the planned activities of the insider. Having anonymous reporting mechanisms can help cyber analysts get ahead of impending threats. Extending that thought, employees should also be encouraged to report others soliciting their passwords – as in a phishing attack.

Other findings included the criticality of "trigger" events. These include access of protected password files or excessive application logons. These type of actions are not data breaches – yet – but may lead to data breaches. Each organization will have their own unique triggers. Focusing in on user behaviors that serve as an "Indication & Warning" can be invaluable. Constructing and implementing triggers unique to each organization should become a best practice and based on the SEI research, this could be a valuable deterrent given the high percentage of security incidents that could be investigated prior to an incident becoming an actual data breach.

Managers and supervisors can play an active role by reporting disruptive employee behavior, violations of company security policies or instances where employees may decide to embark on retribution following dismissal or reprimand. Each of these activities are workflow focused and can be augmented with technical IT controls.

The Carnegie Mellon case studies found a wide range of personnel alerted organizations on the security incidents. In 61 percent of cases studied, non-security persons first reported insider threats. This indicates the potential value of having security awareness

[5] http://resources.sei.cmu.edu/asset_files/TechnicalReport/2012_005_001_34033.pdf

programs and anonymous reporting avenues available to all employees, suppliers and even customers.

In 74 percent of cases, the identity of the insider threat was detected through parsing various system logs. This indicates the value of generating and archiving system logs for future forensic usage. In 30 percent, formal forensic exams were needed to identify a perpetrator, often with crime enforcement support via obtaining home computers or obtaining information from ISP (Internet Service Providers).[6]

For insider threats, most incidents were uncovered by non-automated means. Discovery methods included individuals reporting on suspicious user behavior or simple concerns about an insider. Since insiders already have trusted access, it is difficult to build automated methods of detection. Instead, alerts can be used to focus human investigators. This is where constructing unique indicators and warnings can prove beneficial to create more automated methods of detection.

Auditing and monitoring detected just 20 percent of the cases, hence those depending on compliance and audits to uncover insider threats, may need to augment their defenses. Insiders generally know when audits are scheduled, what type of transactions are being examined, and will seek to move their activities to other times in the calendar. Audits are about compliance not detecting insider threat activity.

83 percent of the attacks took place from within the physical location of the organization and 70 percent of cases occurred during normal working hours. In this series of cases at least, insiders appeared to be doing their best to blend their activities with normal business operations.

The remainder of the cases involved remote access. That remote access constituted just 17 percent of locations for insider threat may be an indication that organizations are serious about their remote access, inherently understanding the higher degree of risk and implementing additional security protections and controls such as frequent transaction auditing and implementing multi-factor authentication schemes.

Mitigating the Insider Threat

Getting in front of the insider threat requires implementing a mix of active controls and a strong organizational culture of security awareness. While it is helpful for IT to have a

[6] Given there were multiple attributions for each case that was studied, the totals exceed 100 percent

strong mindset and implement a targeted grouping of security controls, that in itself is not sufficient.

Active controls are processes put in place to detect all forms of malicious insider activity. These triggers can lead to the implementation of additional controls or create alerts for further situational investigations by incident response staff.

Acknowledging the need to address the insider threat can be the first step to implementing adequate controls. Organizations often push back against implementation of stringent insider controls. Organizations trust their employees and do not want to put fine degrees of control on user activities.

However, a primary objective of malicious software is to infiltrate networks and conduct operations exactly as an approved and authorized user would. Evading detection is a key objective to allow adversarial operations to continue. Detecting and interrupting those activities is what controlling the insider threat is all about.

"Stringent internal controls aren't an expression of distrust with your employees, but a matter of cyber security principle to identify unusual behavior regardless of source."

As such, monitoring operations are intended to identify both the unexpected outsider penetrating network defenses as well as the determined or disgruntled internal individual. Collecting appropriate security metrics brings focus to monitoring and helps to define normal operations, and makes anomalous activities easier to spot. Therefore, stringent internal controls aren't an expression of distrust with your employees, but a matter of cyber security principle to identify unusual behavior regardless of source.

Controlling privileged accounts is one of the key aspects for addressing these concerns, yet is not strictly an IT issue. IT can only implement business rules and policies that support organizational objectives.

In order to thwart malicious insider activities, detection is the first step.

Detect and Deter

Effectively defending against insider threat requires a combination of two approaches: Detection of threats and deterrence, that is, discouraging attempts at insider threat.

When successful, these two practices will serve as mutual catalysts for insider threat defense.

Detection

Stopping data exfiltration caused by insiders requires illicit activities to be uncovered. This requires monitoring user behavior and looking for activities outside normal patterns.

Detection can occur through a mix of software-based monitoring of IT asset use along with human process implementation, such as conducting audits where investigators look for key events and activities.

All detection starts with a sufficiently detailed Security Policy that describes and characterizes acceptable actions by user role. Some organizations codify organizational security policy in a business-focused documents and the allow IT to translate business needs to specific IT actions.

For example, an organizational security policy may state – "we trust all our privileged account users; however we verify all activities". An IT policy statement translation would be "All privileged account usage shall be monitored and activities written into system logs."

Another item may be "All system changes within our protected network domains should generate a log record that can be reconstituted to construct a user activity timeline." In this manner, the organizational security policy states the ***why*** of an action, while the IT version decomposes it into the ***how*** enabling an audit function to be performed by a third party.

Detection methods include:

- Internal audits
- User behavior based security metrics
- RBAC (Role Based Access Control) violations
- Frequent account credential resets
- Utilization of unauthorized cloud hosting services
- Unauthorized elevation of privileges
- Connecting to unauthorized or suspicious external IP addresses
- Employee anonymous reporting
- DLP (Data Loss Prevention) triggers
- Cyber dashboard alerts
- Violations of separation of duties
- Violations of data movement policy based on volume

- Data transfers occurring outside specific time boundaries
- Utilization of unauthorized USB storage devices
- Excessive password resets

It is virtually impossible to move around in IT systems without leaving some type of record. Logs are generated – or can be – with virtually every activity performed. The challenge is to generate logs for future forensic analysis (when needed) without getting in the way of normal business operations. All access to key data items should be logged and archived.

While it may be impractical to implement all the controls listed above, enabling just a portion of them can provide the necessary alerts that can focus a human cyber analyst that a situation has emerged that requires further investigation.

Deterrence Techniques

Implementing Security Controls

Deterrence is achieved by putting roadblocks in the path of the adversary. A primary deterrent is security controls. These are countermeasures to vulnerabilities that reduce risk in IT systems across the enterprise. They slow down or stop an adversary from completing intended activities.

One of the most difficult aspects of insider threat mitigation is walking the tightrope of alerting and blocking. This means an alert is generated, but should an autonomous block be enacted, or should it require human intervention? Performing too many automated blocks can impede internal users in performing their jobs which can have an adverse impact on an organization.

Many real-world examples exist to indicate simply generating huge volumes of alerts is counterproductive. The intent is to build security rules over time which provide fine grained control, without overtly blocking normal business activity. An ongoing automation improvement process provides analytical insight and fine-tuning of business rules.

Some of these protections can be enabled automatically, but others will require human analysis and intervention. Thus, alerts are required. The intent is to prioritize risk and select and implement those security controls which provide the best potential to reduce the risk to the organization.

There are many lists of security controls, but the most widely accepted is the National Institute of Standards and Technology's NIST SP 800-53[7]. While this list is intended for U.S. Federal Information Systems, it has transcended its original purpose. Many organizations now use it as the basis for a smaller set of prioritized security controls, such as the SANS Institutes' Critical Security Controls list.[8]

While the list is intended to better secure the entire IT infrastructure of an organization, some of the controls are specifically important to controlling an insider threat, including: having an inventory of all approved software and hardware, monitoring user activities and controlling ports and devices that can be used to export ("steal") protected data.

The Software Engineering Institute at Carnegie Mellon also created a list of 19 best practices, based on the insider threat study previously discussed. Their list provides organizational guidance and focuses on processes that have to be in place to augment the technical security controls an IT group would enable. Successfully combating the insider threat requires controls and protections outside the typical IT domain across the entire organization structure.

Operational Support

Organizational Best Practices – Human Factors

Organizations with superior insider threat mitigation protections understand the need for the various business units and groups to also provide active support. Simply implementing a set of technical IT security controls is not sufficient. In addition, strong cyber security requires support from related business functions such as recruiting and Human Resources.

Sophisticated organizations take a long view and look at the lifecycle of an employee. In this manner, security protections can be put in place prior to hire (background investigation), during onboarding (signing acknowledgements of data disclosure and protections, assigning specific IT account role) and also upon employee exit (disabling accounts and access). At each step, checklists can be monitored and also help support security incident response, should it be needed.

Categorize threats and identify threat vectors – mobile, social media, suppliers, cloud providers, business units, subsidiaries and the like. While threat categories may span

[7] http://csrc.nist.gov/publications/drafts/800-53-rev4/sp800-53-rev4-ipd.pdf
[8] The full list (Version 5) can be found at https://www.sans.org/critical-security-controls/

organizations in a market sector, an individual organization will want to take a hard look at what makes them unique. International organizations may be exposed to civil unrest and geographic concerns.

The objective is to have sufficiently strong employee interactions in place to uncover morale issues which may lead to future malicious activities. A disgruntled employee does not have to be a skilled hacker; instead, they can simply be an insider conduit to valuable items such as account IDs and passwords. Simply providing equipment lists and network diagrams can be enough to lessen the amount of effort an external attacker has to render in order to breach network defenses.

Evaluating the trust of an employee begins before they are hired and given any type of system access. Many organizations conduct employee background checks and even drug tests. The leading proponents of this approach are the U.S. Government via the Department of Defense and Intelligence Community background checks. Each and every person granted access to protected systems has an extensive background investigation performed by a third-party agency. Once approved, individuals are granted specific levels of access where their activities are monitored and logged. Commercial software products support such monitoring and logging in common examples like the Microsoft Global Access List and system logs.

"The objective is to have sufficiently strong employee interactions in place to uncover morale issues which may lead to future malicious activities."

Once an employee is hired, onboarding procedures can be used to clarify organizational intellectual property and expectations of how data they are entrusted to protect should be handled. Specifying expectations at the time of hire minimizes later confusion or misunderstandings. Onboarding is the perfect time to indoctrinate appropriate data protection methods.

Define roles and responsibilities for each position. (Consider using a risk rating based on employee background or the role of the position.) Roles can define expected behaviors. For example, administrators will be expected to adjust user privileges, run protected software and apply operating system patches. Printing of unencrypted user account information, such as passwords, should not be performed. Uploading of information to non-approved personal cloud storage services such as Dropbox can be prevented or restricted.

Over time, user behavior can be used to create an activity profile that can set alert threshold levels. For instance, a list of normal IP addresses can be created based on conducting typical business activity. Volumetric measures can monitor file movements. Specific servers can be prevented from transferring data to external sites. Access to protected data such as personal health information should have all activity logged to prevent privacy breaches and to maintain regulatory compliance, such as with HIPAA.

Regular employee performance reviews can also directly support enhanced cyber security. Implemented correctly, this process should give a manager the chance to speak with each individual and determine if any employees are disgruntled enough to conduct some type of malicious activity.

Malicious insider activity rarely erupts out of the blue. Some portion of security incidents are 'accidents' but malicious activity is usually pre-meditated and planned to some extent. Instead, the more likely course of action is for resentment to build, leading an individual to harbor increasing hostility to an organization. Having solid conflict identification and resolution processes adds an additional layer of protection against insider threat.

Employee departure is also a valuable time for insider threat analysis. Exit interviews at the time of separation are the final chance to identify a disgruntled employee that may be considering revenge tactics or who may have already planted software to accomplish organizational disruption objectives. Mature organizations must have a formal exit procedure in place to minimize post-separation insider threat issues.

Security Awareness and Training

There is no way to sugarcoat the issue, people are the largest single source of security incidents. Yes, those same individuals that an organization has so carefully recruited and hired represents their largest threat vector. This issue has entered the popular and social media world as the acronym PEBKAC. What is PEBKAC? Problem Exists Between Keyboard and Chair.

While admittedly it may be overly derisive of the end user, it's useful shorthand to remind us that we are always dealing with the human factor, and not all insider threats are necessarily intents to cause harm, some are simply human error.

Verizon, in its 2015 Data Breach Investigations Report, states that it takes a phishing campaign of as little as 10 fraudulent e-mails to achieve a 90 percent probability that one of those messages will lead to a compromise. Users may not intend to cause organizational

harm, but compromising an employee account leads to network infiltration and a host of damaging actions subsequently.

Since a large percentage of security incidents can be traced back to people, it is imperative to have some type of security awareness process in place that regularly exposes employees to relevant cyber security risks to them and the organization.

The SANS Institute's Insider Threat survey points to "lack of training" as the single largest factor limiting an organizations ability to prevent/deter insider attacks

While training breeds awareness, it is crucial to make the training engaging and relevant to the users. Make employees aware that their IT actions will be logged. While technical users may be able to get around such controls by bypassing or turning off logging activity, knowing that an audit trail is being generated may be enough to dispel the more casual insider threat.

Make sure insiders know the organization is prepared to take cyber incidents seriously and will prosecute individuals, not simply sweep items under the rug for fear of organizational embarrassment. Employee association with a cyber incident may be a deterrent to future employees, making good cyber practice all the more important to the end user.

Make sure insiders know about the consequences others have suffered as a result of insider attack and how it has impacted their lives. In cases such as these, fear can be a positive motivator to not conduct illicit cyber activities.

Raise awareness of most prevalent current attack methods and results. Case studies can bring situations to life and make abstract events seem real.

Adopt a CDM approach – Continuous Diagnostics and Mitigation for training – make it continuous, organizational and/or industry focused, and related to what employees can and should actually do. Creating Training Measures of Effectiveness (MOE) can be useful to make sure training is actually improving over time and has not simply become a checklist item. Periodicity is important since security awareness is not a one-and-done activity. As new threats emerge, new breaches occur, and new organizational vulnerabilities must be addressed.

Provide detailed descriptive case studies which relate to the role an individual plays in the organization. Scenario-based training situations leave a longer lasting impact than rote classroom instruction.

Training does not have to occur strictly in a classroom. Hiring a third party to conduct a phishing campaign is effective and increases employee resistance to such attack vectors. Also, don't forget suppliers and the supply chain when deciding who should attend training.

For an organization of sufficient size or with sufficient assets to protect, it can be helpful to actually construct a formal Insider Threat Program that can intimidate less technical individuals from even attempting malicious activities.

Vulnerability Assessments

Conduct regular organizational vulnerability assessments.

This goes beyond simply IT network scanning for malicious software. This also means actively looking at various processes for vulnerabilities or gaps that can enable malicious insider activities.

A checklist of questions can be used to begin an organizational analysis. For instance:

- Is an incident response plan in place? Is that plan periodically tested?
- Are data breach exercises/scenarios conducted?
- Is an incident response team in place and trained?
- Are industry data breaches analyzed to assess new types of organizational risk?
- Are gaps in control structures identified and corrected?
- Are organizational cyber security best practices continually matured?
- Does the organization take advantage of new cyber protection innovations?
- Are security metrics in use evaluated for effectiveness?
- Are special rules for privileged accounts and users in place and utilized?
- Are cyber security software tools previously purchased in actual use?
- Are user endpoints monitored for unusual behavior?
- Are default passwords changed for all newly purchased devices?

An organization vulnerability assessment can be conducted in house or by an external third party. The most important factor is that it is regularly conducted and that the findings are used to enable additional controls and process changes.

Summary

Insiders, by definition, have a privileged view into the structure of an organization. Should they decide to abuse that position, they are in a position to cause an inordinate amount of

harm. Thwarting insider attacks requires a mix of people, process and technology controls. The organization needs to foster a culture of "trust, but verify" on all user activities.

In spite of the continual daily announcements of data breaches, they are not inevitable. Anticipating that your privileged users may become motivated by financial gain does not mean you do not trust your workforce. Instead, it is a reflection of the reality of a digital world where handshakes are often virtual. Additionally, external threat actors and malware seek to impersonate an insider and must be protected against.

Security principles such as "least privilege," "role-based access control" and "defense in depth" are well established principles because they are highly effective. Most data breaches can be traced to a violation one or more of these principles.

While no set of activities or controls will guarantee ***not*** having a security incident caused by an insider, deploying multiple defenses remains the best way to protect and secure all organizational data assets.

Some control items discussed within this paper include:

Identify and monitor insiders. Know who has privileged access to key organizational assets and make sure appropriate protection mechanisms are in place based on their level of access.

Monitor current events. Understand what is happening in the cyber industry at large and what steps organizations are taking to reduce the insider threat risk.

Mitigation requires an approach spanning people, process and technology. All aspects are required for a comprehensive defense.

Utilize industry and cyber vendor research. Look for specific actions that can be taken to enhance vulnerability mitigation.

Data breaches have substantial costs – both direct and indirect. Investigate security incidents before they escalate to an actual data breach.

Have a rigorous security policy in place. Follow it across the organization, update it regularly and build a culture of security awareness.

Utilize an appropriate security framework and controls. NIST has a robust model which can be tailored to individual organizations.

Build a culture of security incident reporting. Employees are human and should be encouraged to report security incidents so an organization can learn from previous errors and accidents.

Detect. Deter. Defeat.

About the Author

Robert J. Michalsky has served government and commercial customers for more than 30 years providing a wide range of IT and analytical services. As NJVC Principal, Cyber Security, he quantifies and pursues new business opportunities in cyber security. Mr. Michalsky is a strong advocate for protecting user data through technology enablement and enhancing business processes through modeling and analysis.