

Insider threat: The human element of cyber risk

Cyber programs often miss the significant portion of risk generated by employees, and current tools are blunt instruments. A new method can yield better results.

Tucker Bailey, Brian Kolo, Karthik Rajagopalan, and David Ware



Insider threat via a company's own employees (and contractors and vendors) is one of the largest unsolved issues in cybersecurity. It's present in 50 percent of breaches reported in a recent study. Companies are certainly aware of the problem, but they rarely dedicate the resources or executive attention required to solve it. Most prevention programs fall short either by focusing exclusively on monitoring behavior or by failing to consider cultural and privacy norms.

Some leading companies are now testing a microsegmentation approach that can target potential problems more precisely. Others are adopting in-depth cultural change and predictive analytics. These new approaches can yield more accurate results than traditional monitoring and can also help companies navigate the tricky business of safeguarding assets while also respecting employees' rights.

Understanding the threat

Organizations sometimes struggle to clearly define insider threat. In this article, we use the term to mean the cyberrisk posed to an organization due to the behavior of its employees, rather than other kinds of insider threat, such as harassment, workplace violence, or misconduct. For these purposes, contractors and vendors are also considered employees; many of the largest cases in recent memory have trusted third parties at their center.

Briefly, inside threats arise from two kinds of employees: those who are negligent and those with malicious intent (see sidebar, "Double trouble"). Negligent or co-opted insiders are easy for companies to understand; through poor training, middling morale, or pure carelessness, usually reliable workers can expose the company to external risks. However, organizations often misunderstand malicious insiders in two ways.

Double trouble

Two types of workers can create cyberrisk:

Malicious insiders are those who purposefully seek to benefit themselves at the organization's expense or to harm the organization directly. They might steal valuable data, commit fraud for financial gain, publicly expose sensitive information to attract attention, or sabotage IT systems in disgruntlement. Most organizations focus their attention on malicious insiders, using activity-monitoring software and small investigative teams.

Negligent or error-prone insiders may not harm an organization intentionally but expose the organization

to risk through their mistakes or carelessness. This can happen in two ways. First, an employee can carelessly create a vulnerability, which can be exploited by attackers directly. For example, a developer might misconfigure a company's Simple Storage Service (S3) buckets, or someone might lose a hard drive carrying sensitive data. Employees can also make themselves personally vulnerable to attack and co-option. For example, by sharing too much personal information online, employees may make themselves easy targets for spear-phishing attacks, in which attackers co-opt a user's account and use it to conduct further nefarious activities.

First, malicious insiders do not always seek to harm the organization. Often, they are motivated by self-interest. For example, an employee might use client information to commit fraud or identity theft, but the motive is self-enrichment rather than harm to the employer. In other cases, employees may be seeking attention, or have a “hero complex” that leads them to divulge confidential information. They might even think they are acting for the public good, but in reality they are acting for their own benefit. Understanding motive can help companies shape their mitigation strategy.

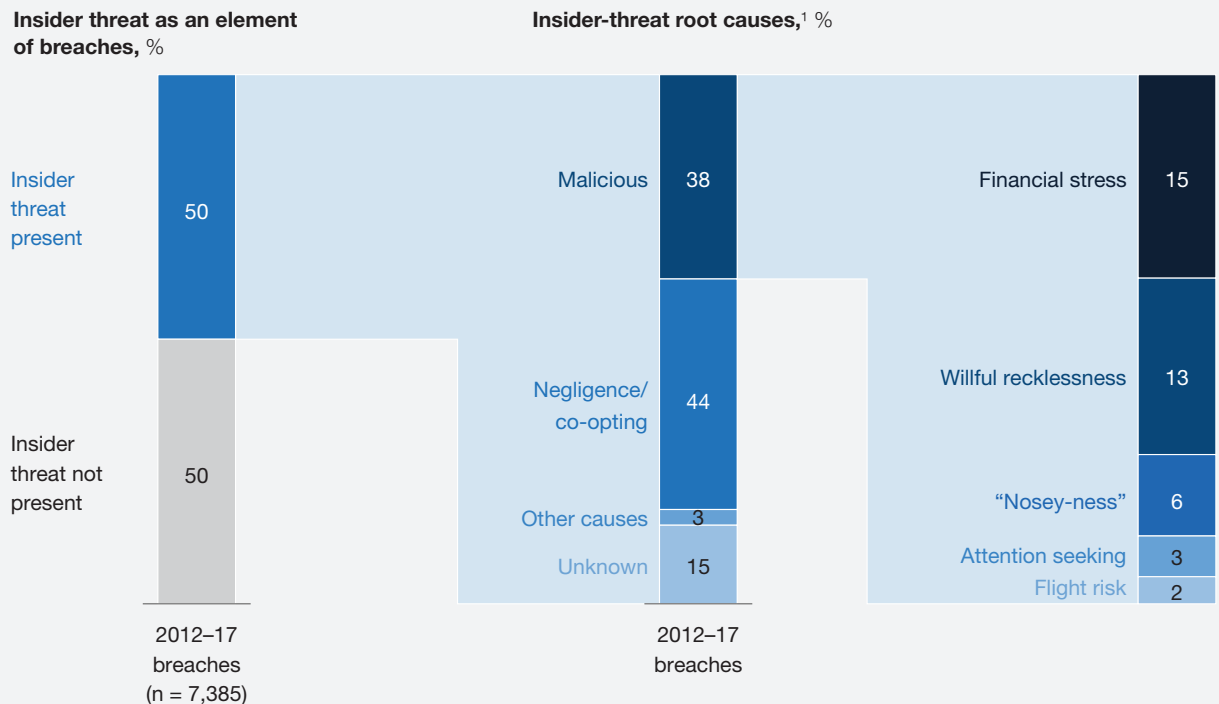
Second, malicious insiders rarely develop overnight or join the company intending to do it harm. In most recent examples of malicious insider events, normal employees became malicious insiders gradually, with months or years of warning signs leading up to a culminating insider event.

How big an issue is it, really?

In a world of competing cyber-priorities, where needs always seem to outpace budgets, it can be tempting to underinvest in combating insider threat. The risk is not well understood, and the solution feels less tangible than in other cyber areas. Many executives have asked us, “Is this actually an important issue? How much risk does it represent?”

We recently reviewed the VERIS Community Database, which contains about 7,800 publicly reported breaches from 2012 to 2017, to identify the prevalence of insider threat as a core element of cyberattacks. We found that 50 percent of the breaches we studied had a substantial insider component (Exhibit 1). What’s more, it was not mostly malicious behavior, the focus of so many companies’ mitigation efforts. Negligence and co-opting

Exhibit 1 Insider threat is present in 50 percent of cyberbreaches.



¹ Figures are approximate and may not sum, because of rounding.

Source: VERIS Community Database

accounted for 44 percent of insider-related breaches, making these issues all the more important.

In addition to being frequent, insider-threat breaches often create substantial damage. We have seen high-value events in which customer information was stolen by both negligent and malicious insiders in financial services, healthcare, retail, and telecom in recent years. Some companies lost hundreds of millions of dollars. Pharmaceutical and medical-products companies, as well as governments, have seen a significant rise in intellectual-property theft by malicious insiders.

Why current solutions fall short





To combat the risks of malicious insiders, most companies rely on user-behavior monitoring software (Exhibit 2). These rules-based or machine-learning-

based applications ingest troves of data about employee actions, especially their use of IT systems. Generally, they attempt to identify divergence from what is considered “normal” behavior for that employee. When the software spots an anomaly, a small team investigates.

While this method can be helpful, we find that it usually falls short, for four reasons:

- By the time negative behaviors are detected, the breach has often already occurred. The organization is already at a disadvantage, and it cannot deploy an active defense.
- Monitoring for “divergence from normal behavior” creates a huge number of false positives, wasting much of the investigation team’s time.

Exhibit 2 Current methods of management fall short.

				
Typical approach	Prevention and monitoring	Event detection: Behavior-variability analysis	Investigation	HR/business-unit action
	<ul style="list-style-type: none"> • “Dragnet” monitoring of all employee actions, all the time • General controls and preventions 	<ul style="list-style-type: none"> • Analyze divergence from “normal” behavior 	<ul style="list-style-type: none"> • Manually investigate numerous cases 	<ul style="list-style-type: none"> • Take actions on a case-by-case basis
Pain points/risks	<ul style="list-style-type: none"> • Massive number of signals • High risk of misuse of data • Perception of privacy invasion • Preventions not customized to risks, actors, and actions 	<ul style="list-style-type: none"> • Bad behaviors can be built into baseline • Huge volume of false positives (>30% in some cases) 	<ul style="list-style-type: none"> • Often a long backlog of cases • Little ability to prioritize investigations 	<ul style="list-style-type: none"> • Uncertainty about how to manage between investigation and action • Actions not well defined or tailored to individual incidents

- Serial bad actors may not be caught; malicious activity may be built into the baseline of “normal” activity.
- Collecting massive amounts of employee data creates privacy concerns and significant potential for abuse.

Beyond these issues, some organizations take this type of monitoring to an extreme, deploying military-grade software and conducting full-blown intelligence operations against their employees. Several recent news stories have highlighted the risks of overstepping the organization’s cultural and privacy norms. Best practices and necessary precautions in the defense industry may be seen as invasive at a bank or insurer.

Finally, to the extent that companies pursue insider threat, they often focus on malicious actors. While most cyber organizations know that negligence is an issue, many start and end their prevention efforts with half-hearted employee education and anti-phishing campaigns.

A better way

Some leading cybersecurity teams are using a different approach, built on three pillars:

- **Microsegmentation** allows the organization to home in on the “hot spots” of risk and take a targeted rather than blanket approach to threat monitoring and mitigation.
- **Culture change** makes malicious, co-opted, or negligent risk events less likely, and puts the company in a preventive rather than reactive posture.
- **Prediction** allows an organization to identify and disrupt insider activities much earlier in the threat life cycle.

Microsegmentation

Rather than going immediately to wholesale monitoring, we believe that organizations should take a much more nuanced approach, tailored to their information assets, potential risk impacts, and workforce. The key to this approach is microsegmentation, which identifies particular groups of employees that are capable of doing the most damage, and then develops focused interventions specific to those groups.

To create a microsegmentation, the first step is to understand the business capabilities or information most important to protect. Next, companies can use identity-and-access-management (IAM) records, as well as organizational and HR information, to determine which groups and individual employees have access to those assets. These groups form the microsegments that are most important for the program to focus on. For each segment, a company can then determine which types of insider threats are most likely to cause damage, and it can create differentiated strategies to monitor and mitigate insider events.

Imagine that a pharmaceutical company wants to protect the intellectual property created in new drug development. An analysis of IAM and HR data reveals that specific portions of its product-development and its R&D organizations represent the highest risk. The company knows that sabotage of this kind of IP is relatively rare (other researchers would easily catch mistakes), but that flight risks—scientists who take IP with them when hired by competitors—are very probable. The company designs strategies to identify flight risks in the R&D team (such as people who missed promotions, poor workforce satisfaction, and low pay relative to peers), and then monitors the group for these characteristics. The company could then design interventions, such as retention programs, specifically for its flight risks.

Microsegmentation offers three key benefits. First, it creates a clearer understanding of risk; not all insider-threat events are created equal. Second, it allows organizations to identify a clear set of remediation actions, tailored to a particular group of employees. This helps them to move from reacting to insider-threat events to preventing them. Finally, the analysis allows the organization to monitor groups rather than individuals, using metrics such as employee attrition and workforce satisfaction of a team rather than individual behaviors. This provides significant privacy benefits.

Exhibit 3 shows an illustrative microsegmentation analysis for several kinds of information assets.

Culture change

While many programs focus on catching and responding to negative behaviors, it's also vitally important to directly and assertively address the cultural issues that drive negligence and malicious behavior.

To combat negligence and co-opting, companies often conduct rudimentary cybersecurity trainings, as well as phishing testing. Too often these focus only on behavior—educating employees on proper cyber-procedures—and miss the attitudes-and-beliefs part of the equation. Targeted interventions (such as periodic communications on cyber-impact) help employees see and feel the importance of “cyber-hygiene,” and

Exhibit 3 Microsegmentation can reveal groups at risk, the actions they might commit, and their likely personas.

Threat assessment, illustrative example

■ Very likely ■ Somewhat likely ■ Not likely

Top assets	Employee populations with access	Insider-threat actions they might take			Likely personas involved
		Fraud/theft	Exposure	Destruction	
Intellectual property for new products	<ul style="list-style-type: none"> R&D team Business-unit (BU) exec 	Very likely	Not likely	Somewhat likely	<ul style="list-style-type: none"> Flight risk Disgruntled
Financial forecasts	<ul style="list-style-type: none"> Finance/investor-relations team BU execs 	Very likely	Somewhat likely	Not likely	<ul style="list-style-type: none"> Financially stressed Negligent
PII/PHI ¹	<ul style="list-style-type: none"> HR team Sales team 	Somewhat likely	Very likely	Very likely	<ul style="list-style-type: none"> Negligent Reckless Snooper
High-net-worth customer information	<ul style="list-style-type: none"> High-net-worth sales and delivery team 	Somewhat likely	Not likely	Very likely	<ul style="list-style-type: none"> Flight risk Financially stressed
Core financial platform	<ul style="list-style-type: none"> IT team BU execs 	Somewhat likely	Not likely	Somewhat likely	<ul style="list-style-type: none"> Saboteur Disgruntled
Records of corporate conduct	<ul style="list-style-type: none"> HR/legal 	Not likely	Somewhat likely	Very likely	<ul style="list-style-type: none"> Attention seeker

¹ PII = personally identifiable information, PHI = protected health information.

purposeful reinforcement from senior executives is critical to achieving workforce buy-in. Best-in-class organizations rigorously measure both behaviors and attitudes and develop comprehensive change plans to beat cyber-negligence, complete with targets and clear owners within the organization.

Addressing the drivers of malicious behavior is an even more personal task. The drivers vary for each organization, and often for each microsegment. For instance, they might include personal financial stress, disgruntlement over lack of promotion, or flight risk due to poor management. Organizations that successfully address drivers of malicious behavior often begin by analyzing workforce trends (using satisfaction surveys, for example) to determine potential hot spots. They then design changes in process, governance, hiring, compensation, and so on, specific to the identified risk areas aligned to their microsegmentation strategy. For example, if an employee group has a high prevalence of “flight risks” due to disgruntlement over a manager, the

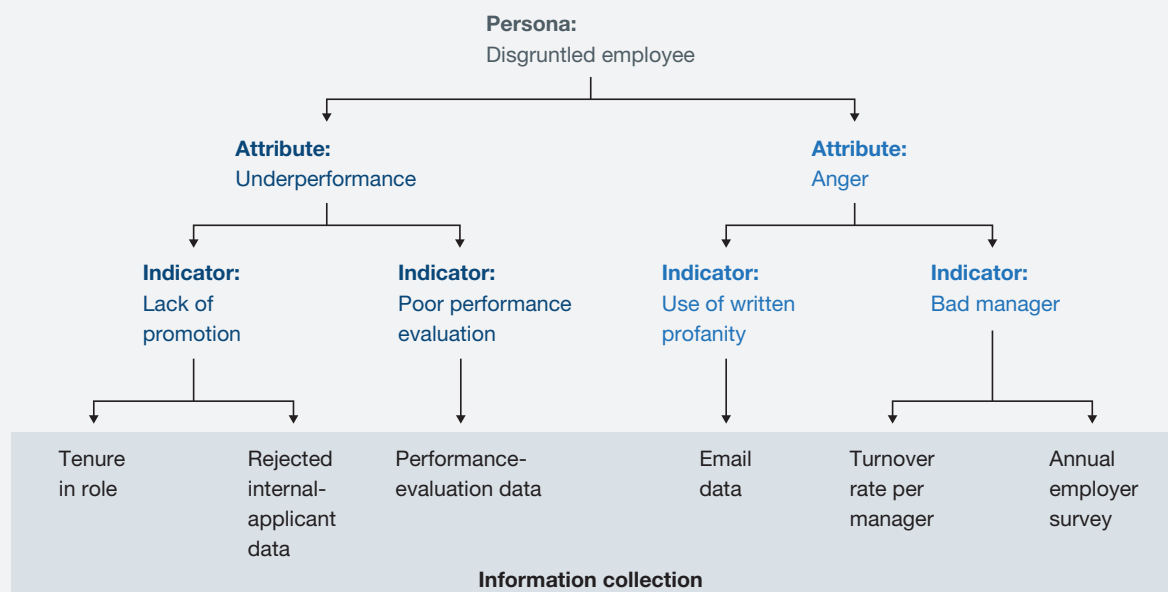
organization may require leadership coaching or even rotating the manager out of the group. If financial stress seems to be an issue, the organization may choose to provide free financial-planning help or to reevaluate its compensation model.

Prediction

Advanced organizations are taking one further step to identify groups or individuals early in the threat life cycle: predictive insider-persona analytics. The main personas that present a risk are well established and have been studied at length. High-performing organizations have identified the markers of these personas and actively monitor these markers for specific personas, rather than looking for divergence from normal. This analysis can identify a group or individual likely to represent a threat well before the event takes place; companies can then take steps to mitigate the threat. Exhibit 4 outlines the predictive analysis for identifying disgruntled employees, one of the established personas.

Exhibit 4 The markers of risky personas can give companies a head start on intervention.

Example risky persona



While powerful, these analytics require careful consideration about their use in the context of an organization's culture, its privacy norms, and the evolving standards of privacy in society at large. Failing to think it through often results in employee complaints about invasion of privacy.

A few words on privacy

Privacy is an inherently personal and intangible subject—its meaning and importance varies by geography, by industry, by company, and often by individual. Many individuals are fiercely protective of their privacy, even when at work and even in their use of corporate assets. This is never more true than when it comes to monitoring their use of communications systems such as email—even corporate email. As standards on individual and corporate privacy rights evolve (for example, through the European Union's General Data Protection Regulation), organizations need to design their insider-threat programs based on what will work within their own cultural and regulatory environments. In all cases, organizations need to tailor their insider-threat program by respecting what data may be gathered, how it may be collected and used lawfully, and how best to create awareness of the program, both generally and specifically, with potentially affected staff.

While each organization must make its own trade-offs between privacy and risk, we believe our approach will make such trade-offs easier to navigate than traditional programs. First, the microsegmentation approach does not require a baseline of individual activity (by which traditional programs judge “normalcy”), which some organizations could perceive as a privacy concern. Second, microsegmentation presents natural groups of employees for analysis,

which improves the anonymity of the analysis.

Microsegmented groups can be analyzed for potential threats with reasonable precision of results. Investigations of specific individuals can be conducted only when there is reasonable suspicion of a threat and must be done in line with applicable law.



Insider threat is one of the largest problems in cybersecurity, representing a massive share of attacks and financial damages. Monitoring technologies have their place in organizations' cyber-arsenal. But their effectiveness increases significantly when combined with more nuanced approaches, like microsegmentation, prediction, and direct cultural engagement. ■

Tucker Bailey is a partner in McKinsey's Washington, DC, office, where **Brian Kolo** is a digital expert and **David Ware** is an associate partner; **Karthik Rajagopalan** is a consultant in the Dallas office.

Copyright © 2018 McKinsey & Company.
All rights reserved.