TERAMIND

Best Practices for Building an Effective Insider Risk Program

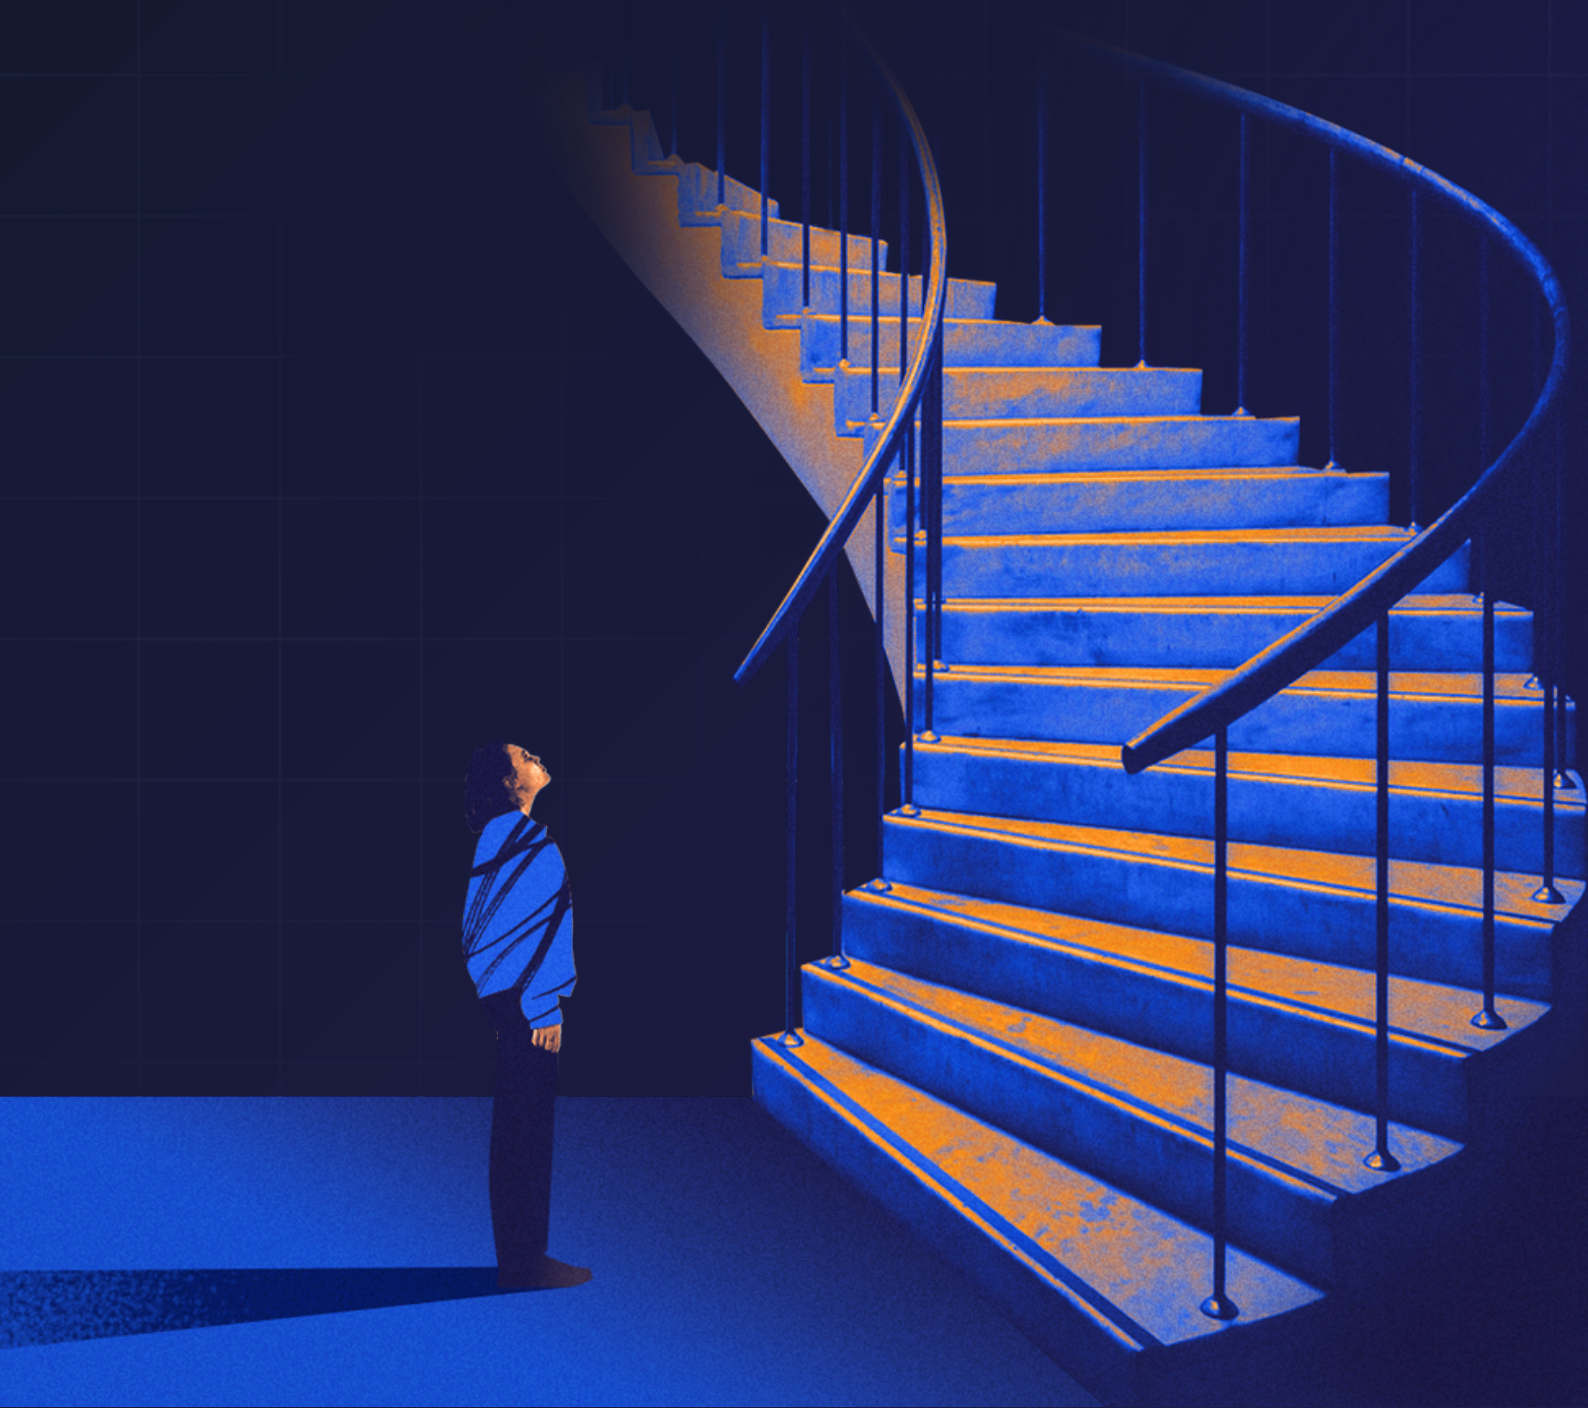

Staying ahead of malicious insiders demands a unique combination of **preparation, technology, and expertise.**

Excellent insider risk programs build a bridge between employees, leadership, human resources, and security. They enable analysts to contextualize security events in unique and valuable ways.

Understanding risky versus normal behaviors **requires baselines.**

Enterprises have many different business units, locations, and network segments. An enterprise that grows through acquisitions may even have whole subsidiaries inside its environment.

Yet, mitigating insider risk is a team sport. Detecting deviations from baseline behaviors requires a deep, contextual understanding of what those baselines look like.

Quantifying insider risk in this context is challenging. A malicious insider may already have all the privileges and authorizations they need to carry out an attack. Distinguishing between malicious behavior and false positives demands visibility and context.

Establishing baselines is a crucial part of User and Entity Behavior Analytics (UEBA), but obtaining accurate baselines is not just a technological problem. It extends into the organization's security architecture, company culture, and more.

Challenges to Obtaining Accurate Behavioral Insights:

Behaviors are not static, but baselines often are.

Normal baseline behavior changes from one employee to the next and from one team to the next. Individuals may change their behavior over time as well.

Behavioral models need continuous refinement and customization. Both organizations and people change constantly, and behavioral models must be maintained to keep up. Log data is a great place to start, but collaborating with other internal stakeholders allows analysts to better understand the context behind baseline.

Behavioral norms vary greatly across roles.

Actions that are completely normal for one department would be a compliance violation in another, and vice versa. As people change roles within the organization, monitoring rules also need to shift.

The ability to configure monitoring on a team / role-based basis, or even an individual basis, supports the ability to ensure that behavioral baseline data is not only accurate, but minimizes false positives. Granular settings are essential in this process – ensuring analysts see all of what they need to, and don't waste time on irrelevant data.

Identifying malicious behavior often relies on context.

A single data point is a starting place, but it takes additional telemetry and behavioral data collected over time to enhance context for determining whether a particular action is malicious or not.

Screen recordings can provide specific context for a particular action. This enables insider risk teams to pinpoint user intention more accurately. This crucial piece of context can mean the difference between identifying an intentional malicious act and an honest mistake.

Security controls can vary between locations.

Corporate campus controls may be different from branch office controls. For example, corporate employees may authenticate with a badge while branch office employees use a PIN code. Different network segments may have entirely different security tools and policies.

These systems may generate log data in many different types and formats. Some may be easier to normalize than others, and some may be practically unusable from an insider risk perspective. Insider risk teams will need to collaborate with employees and department leaders to gain visibility and context into those activities.

Accurately quantifying insider risk requires visibility, which doesn't always come down to log data.

Insider risk programs depend on visibility. In many cases, visibility is generated by accurate and well-documented logs analyzed through a Security Information and Event Management (SIEM) platform. But collecting more logs doesn't always lead to improved visibility.

Redundant logs offer a good example of why this is the case. A financial institution may not need to collect vault access logs if other tools already offer visibility into who accesses its vaults. Most institutions already have strict measures like dual control in place for this purpose.

Feeding all log data into the organization's SIEM presents additional problems. SIEM log storage costs do not scale well, and they only increase when organizations start moving data between different cloud providers and regions.

Insider risk teams that rely entirely on interpreting log data in a SIEM often face the following obstacles:

SIEM analysts feel pressure to delete old logs. This can impede investigations. If an insider threat has been acting on the network for months at a time, the security team will need that log data.

User intention becomes harder to pinpoint. It's not always possible to distinguish between malicious activity and an honest mistake strictly using log data.

Baseline behavioral models deteriorate over time. Log data accessibility can have a significant impact on the accuracy of baseline behavioral models, making them grow less accurate over time.

Insider risk teams can overcome these challenges in several ways:

- ✔ Cutting redundant logs and null values brings immediate value to insider risk operations.
- ✔ Routing logs to low-cost storage is a good solution, but it's not a feature native to most SIEM platforms.
- ✔ Augmenting log-based UEBA with on-demand visibility and context reduces the team's overall reliance on log data.
- ✔ Collaborating proactively with department leaders to identify and accommodate changes that will impact baseline behavioral models.

The best insider risk programs are **proactive, not reactive.**

New insider risk programs are often structured around alert-based tasks. This makes sense when the insider risk team is an offshoot from the Security Operations Center or Incident Response team. But if these teams expand their view, it can directly correlate to increased value in broader insider risk management.

Alert-based tasks are reactive in nature. In most cases, this causes the insider risk workflow to overlap with existing incident response processes. Good management will help ensure each team hands off tasks to the others efficiently, but there's more to insider risk than responding to alerts.

Proactive, referral-based tasks add value to security processes.

Insider risk teams can significantly change the organization's overall security posture by opening themselves up to referrals. These referrals can come from the security operations team, Human Resources, or any other department.

The referral process simply requests context and background on a particular activity or individual. Someone is taking initiative to flag unusual behavior without necessarily launching a formal investigation.

Example: Requesting a 360-degree View of a Particular Insider

A SOC analyst suspects that an employee may be misusing their privileges for personal gain. They reach out to the insider risk team requesting a report on their activity.

The insider risk team pulls records of that employee's position history in the company, their web searches on company devices, application usage, chat history, and more. They produce Sankey charts showing file movement flow, indicating what network assets the employee interacts with.

The team aggregates this data through its UEBA platform and feeds it into the SIEM and analyzes the user's behavior. It can then report its findings back to the SOC.

Collaborate with risk management teams in every business unit.

Successful insider risk teams operate in collaboration with incident response, security operations, and risk management. They help contextualize individual activities and draw conclusions about the scope of insider risk.

This is critical for building models with accurate baselines. From this starting point, the team can begin eliminating false positives and fine-tuning its capabilities.

When insider risk teams collaborate effectively with risk management team members in other business units, they help identify the scope of insider risks and policy violations.

The best insider risk teams help answer a critical question — *Is this particular problem an individual problem? A team problem? A company-wide problem?*

Insider Threat Examples: Incidents and Challenges

A proactive, referral-based insider threat team can work alongside incident response and security operations personnel to support investigations and improve event outcomes across the enterprise. We've picked a few real-world examples from our research to demonstrate this.

1. An employee secretly outsources their work to third-party contractors.

This potentially widespread practice introduces completely unknown security risks while obscuring visibility and bypassing controls. It can be difficult to detect using traditional security tools because the employee is not accessing network assets outside their usual routine.

This kind of activity might only be detected when someone refers that employee to insider risk analysis. SIEM-based behavioral analysis alone may not flag it, but the addition of screen recording software and other endpoint behavioral monitoring solutions will.

Mitigating this kind of incident may rely on establishing privileged access accounts and managing them with an Identity and Access Management (IAM) solution. Deeper integration and automation can help deter future instances of secret outsourcing by showing employees when their screen is being recorded.

Challenge:
Many roles include delegating tasks to third-party contractors. Distinguishing between fraudulent outsourcing and policy-compliant delegation requires insight that goes beyond log data.

Solution:
Screen recordings can provide crucial evidence proving that intentional fraud has occurred.

2. Client-facing team members are sending sensitive customer data over unsecured email.

This scenario happens frequently in industries that rely on the exchange of Personally Identifiable Information (PPI). For example, real estate lending teams at financial institutions have to solicit protected information from customers, including their names, addresses, and Social Security Numbers.

In this scenario, the practice may be so widespread that incident response is unable to address it efficiently. If a business unit doesn't have robust controls against sharing protected data in unsecured communications, it's likely that the entire team — or the entire company — will simply see it as a normal way of doing things.

Detecting and responding to individual policy violations is a job for insider risk teams.

Ensuring policy compliance on a large scale is a company culture issue. In this case, leadership will need to address the root cause of the problem and craft new policies to address it.

Challenge:
If the behavioral baseline model learns that sending protected data through unsecured email is normal, the activity will not trigger any alerts.

Solution:
Accurate baseline modeling is not just a snapshot of how things are done — it's a description of how things should be done. Compliance enforcement at the user level can stemthis without adding alerts to the SOC team's work load, but can still feed notifications to let them know how often attempts are still being made.

3. An employee tries to mass-delete files on their last day of work.

Protecting organizations against disgruntled employees is a common theme among insider risk technology vendors. Alarmingly, cybercriminal groups have learned to reach out to these individuals and co-opt them into their attacks as well [*].

Monitoring the behavior of authenticated users is a pillar of Zero Trust. Any employee who tries to delete, encrypt, or move a large number of files should be subject to investigation. In the best case scenario, investigators are equipped to distinguish between accidental negligence and intentional malicious behavior.

Preparation is key in this case. The insider risk team can only detect this activity if it has taken time to configure security tools to automatically flag and prioritize suspicious behaviors beforehand. It must also understand the employee's history and social disposition in a way that is hard to express as a datapoint.

Challenge:
From a log-centric perspective, the disgruntled employee looks very similar to the negligent one.

Solution:
Assessing that employee's intention can mean the difference between investing in better training or taking legal action.

Your insider risk program *will* impact workplace culture — **make it a positive change.**

It can be tempting to try establishing an insider risk team without impacting company culture. Security leaders might try to accomplish this by making insider risk a low profile unit with few regular connections to the rest of the organization.

This approach often has the opposite effect. The attempt to shield company culture from the impact of insider risk programs can actually cause that culture to deteriorate.

When employees find out that an internal security team is monitoring their activity — which they will — they will conclude that their employers don't trust them. Distrust can breed resentment that makes malicious insider threats an eventual certainty.

Our research has identified several best practices for implementing insider risk programs in a culture-positive way:

Insider risk team members should be friendly, personable, and professional.

Insider risk professionals will often find themselves in the position to educate employees about security policies and workflows. Not all of these employees are malicious insiders. The insider risk program must leave room for honest mistakes.

Explaining policies in a hostile or judgemental manner can push employees away. Employees who have had a bad experience won't reach out to the insider risk team when they notice unusual behavior in the future.

Discouraging referral-based investigations makes it much harder for the insider risk team to work effectively. Having the team work courteously with employees and encourage them to be candid significantly improves security event outcomes and enables early detection.

Building good relationships with other business departments is important.

Your insider risk program uses a SIEM to collect and analyze data from every corner of the organization. Most of these data sources will be embedded in other business departments, meaning you need cooperation and buy-in from department leaders and their teams.

Log data alone may not give you the context needed to accurately investigate a potential insider threat. To do that, insider risk teams need to collaborate with people who work in external departments. Their input helps reduce the amount of time and effort wasted investigating false positives.

External cooperation enables insider risk teams to build comprehensive custom dashboards using data from other business units. This is especially true for departments that frequently interact with the insider risk team, like Security Operations and Human Resources.

Don't expect your insider risk program to become a secret police force.

The temptation to keep the insider risk program a secret is understandable, but not optimal. Secretive teams quickly run into a variety of problems:

- **Subjective interpretations become the norm.** If the team only looks at log data, it can't quantify or communicate risk in a contextually meaningful way.
- **Budget requests are difficult to justify.** If there is no visibility into insider risk activity or outcomes, there is little reason to dedicate more funding to the program.
- **Oversight is limited or non-existent.** When insider risk actions lead to negative outcomes, the process that led to that outcome is rarely examined. This makes improvement almost impossible.
- **Contextualizing baseline behavioral change is impossible.** Since secretive insider risk teams don't collaborate with other departments and leaders, they have no idea how and why normal behavior may change over time.

Instead, security leaders may position insider risk programs as a measure for improving employee safety and security. Employees should know who the insider risk team is, and be encouraged to reach out to them whenever something makes them feel unsafe. This small adjustment makes insider risk management a source of comfort for employees, not a source of anxiety.

Insider risk teams can and should provide ongoing support for leadership.

Leaders tend to get attached to their employees, and discovering a breach of trust can be a painful and disruptive experience. The responsibility for supporting leaders through that experience falls on the insider risk team.

That support comes in the form of justifying and confirming actions taken in response to insider threats. It comes in the form of visibility into insider risk and improved policies for mitigating those risks.

Insider risk support also helps leaders appreciate risks associated with third-party contractors. The insider risk team is responsible for ensuring contractor activities are monitored through UEBA and reporting on contractor-related security risks.

Providing this kind of support helps insider risk team leaders avoid having to justify the value of insider threat detection and response in terms of ROI. Quantifying returns on insider risk management can lead decision-makers to overlook the value and nuance of insider risk operations.

Invest in solutions that **enhance security decision-making and performance.**

Effective insider risk programs require expert personnel and a robust tech stack. The best programs unify these two important elements to produce meaningful long-term results.

Security and usability must remain in balance.

As with all cybersecurity investments, the balance between security and usability is key. An excellent insider risk program knows how to manage these two priorities according to real-world context.

For example, imagine the insider risk team discovers widespread abuse of remote access vulnerabilities. The immediate recommendation may be to shut down all remote access. However, this can have unintended consequences.

If technical employees can't get support through remote access, downtime may increase. This can lead to significant operational costs and reputational damage.

Scalable, automation-friendly technologies enhance the capabilities of small teams.

Insider risk teams are rarely large. Even in a large organization in a high-risk sector like finance, the team may only consist of two or three full-time employees.

This may remain the status quo even if the enterprise grows significantly or makes large acquisitions. Scaling insider risk operations is vital to maintaining predictable performance over time.

Designing insider risk operations around scalability often means investing in automation. Reduce the number of times analysts have to pivot between tools and improvise solutions to gain visibility.

For example, many messaging platforms provide a unique file for every individual message sent on the platform. Analyzing an employee's chat history might involve uploading thousands of these files into the SIEM. This is a time-consuming and error-prone process that holds back productivity for the insider risk team.

Prioritize tools that optimize end-to-end tracking of user activity and real-time visibility.

Many SIEM solutions include integrations for behavioral risk scoring, visualization, and other features that improve the user experience. At the same time, insider risk programs tend to be underfunded, making up less than 8% of the IT budget.

Technological solutions like endpoint user behavioral monitoring can augment native SIEM capabilities and provide meaningful context into insider activities. However, executive decision-makers aren't always convinced that investing in additional capabilities will pay off.

Insider risk teams that open up their processes to oversight and grant visibility into their activities have a much better chance of getting the features and toolsets they need.

Stay ahead of challenges that can impact insider threat detection and response.

Most executive decision-makers already agree that controlling insider risk is a high-priority issue. However, organizations often run into obstacles transforming that issue into an initiative that demonstrates shareholder value.

Creating enforceable policies remains a major challenge.

Without comprehensive and well-documented policies, insider risk teams may face pushback from employees under investigation.

Employees who aren't guided by clear policy may feel threatened by the insider risk program. They may hold up investigations by claiming they didn't know they were violating policy, or by pointing out that there are no policies that apply to their actions specifically.

At the same time, the enforceability of the organization's policies are directly related to its technological capabilities. If the insider risk team can't tell when employees are secretly outsourcing their jobs to third-party freelancers, telling employees not to do it is an empty demand.

Maintaining custom rulesets demands visibility and expertise.

Default SIEM configurations don't generally provide the level of detail that insider risk programs need. Every organization is unique, and its exposure to insider threat risks will not align perfectly with an off-the-shelf SIEM implementation. However, customizing your rules and enhancing telemetry can not only save your analysts valuable time, but make the SIEM you have invested in even more effective.

Custom rules apply labels to certain activities and correlate to dynamic risk scores that can trigger immediate investigations. Risk-based alerting and similar customized triggers can help reduce false positives and streamline insider risk management overall.

Achieving this level of performance requires in-depth visibility and specialist expertise. Not every organization is equipped to implement this level of security using in-house resources and technologies. Most will need to augment their SIEM with technology that improves visibility and makes existing talent even more effective, allowing them to expand their roles to also cover insider risk at the outset of the program.

Conclusion

The number of organizations experiencing between 21 and 40 insider security incidents per year rose by 67% between 2021 and 2023 [*]. **The average cost of malicious insider incidents is more than $700,000, not counting long-term reputational damage.**

As long as insider threats continue to pose the greatest danger to operational security, insider risk management must remain a top priority for business leaders in every industry.

Building an effective insider risk program is the first step towards managing these risks. Implementing technologies and policies that improve the program help guarantee meaningful event outcomes and ensure long-term operational security excellence.

TERAMIND

Request Your Custom Demo Now

Get Demo