Managing Insider Risk: Malicious Insiders and Compromised Credentials

A practical guide for developing a mature insider risk program

Insider risk involves harmful or unintended actions originating from people who already have legitimate access to systems and data. These individuals include employees, contractors, partners, and service accounts trusted to carry out daily work. As organizations adopt automation, AI agents now perform tasks with authenticated access, and they should be included alongside human identities when monitoring activity and defining governance controls.

Most insider incidents fall into two familiar categories. Malicious insiders deliberately misuse their access to steal data, disrupt systems, or pursue personal or financial gain. Compromised users create a different challenge: Attackers operate under the guise of legitimate credentials, making their activity harder to distinguish from normal behavior. In both cases, the actions often look routine until they accumulate into meaningful harm. A mature insider risk program helps security teams identify these patterns early, assess their impact, and intervene before the situation escalates.

Malicious Insider	Compromised User	AI or Automated Agent
A person with authorized access who intentionally misuses their privileges to steal information, disrupt operations, or pursue personal gain. Their actions are deliberate and often planned, making early behavioral visibility important.	An otherwise legitimate account that has been taken over through phishing, malware, credential theft, or misuse of tokens or keys. Attackers operating under this identity typically blend into normal activity, making unusual patterns the key signal to investigate.	A system or agent that performs tasks using authenticated access to applications or data. These agents usually follow defined workflows, but misconfiguration, unexpected inputs, or compromised credentials can cause actions outside their intended scope. They should be included in access reviews and activity monitoring.

Table 1. The three categories of insiders to account for in modern risk programs: malicious insiders, compromised users, and automated or AI agents that operate with authenticated access.

Building a Strong Insider Risk Program

A strong insider risk program begins with clarity. Your security team needs to know who is responsible for insider risk activities, how decisions are made, and which processes guide investigations or remediation work. This foundational structure ensures insider threats are managed consistently, even as roles evolve or new technologies like AI-driven automation are introduced.

Establishing ownership gives the program shape. Someone must guide how access is approved, how activity is monitored, and how findings are communicated. Documentation reinforces this clarity by outlining expectations for user onboarding, access changes, entitlement reviews, and the handling of service accounts, including any AI agents that authenticate on behalf of business workflows.

Measurement brings the program forward. Tracking investigation volume, access exceptions, and changes in user behavior helps you observe how risk evolves over time. Reviewing how automated systems use their permissions ensures these systems remain aligned with intended tasks and don't accumulate unnecessary access. Over time, these insights guide program improvements and allow your security team to scale their efforts as the organization grows.

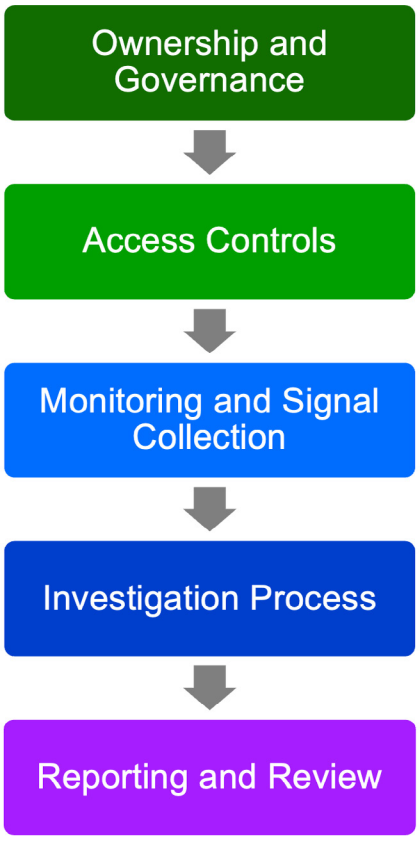

Figure 1. Core components of an insider risk program, from establishing ownership to reviewing findings and refining controls.

Activity Type	Examples of Normal	Indicators That Require Review
Human user	Typical login patterns, routine file access, scheduled work tasks, and activity that aligns with the person's role, responsibilities, and working hours	Logins at unusual times, attempts to access systems outside the individual's role, sudden spikes in data movement, or activity that doesn't align with their usual responsibilities.
Automated or AI agent	Consistent, predictable execution of defined workflows, regular data retrieval or updates, and actions that occur on a set schedule or follow a documented process.	Interaction with systems not associated with the agent's intended workflow, unexpected increases in volume or frequency, configuration changes not tied to planned updates, or actions triggered by typical inputs.

Table 2. Comparing expected activity patterns for human users and automated or AI agents helps identify anomalies that may signal risk.

Preventing Malicious Insider Activity

Prevention begins before an insider ever logs in. Organizations that screen personnel carefully and define access needs early significantly reduce the chance that unnecessary permissions linger in the environment. The same principle applies to automated systems. When an AI agent is created to support a workflow, its access should reflect the smallest set of tasks it must perform. This avoids situations where agents inherit broad or outdated permissions that introduce risk.

Once access is granted, visibility becomes essential. Changes in behavior can signal emerging issues: unusual login times, unexpected interest in sensitive files, or sudden increases in data movement. These signs appear in both routine user accounts and the service accounts tied to automation. Even when AI agents follow predictable routines, periodic reviews help ensure their activity remains aligned with their intended purpose.

A disciplined approach to permissions protects the environment over time. Roles shift, projects end, and people change responsibilities. Without regular reviews, access can drift. Removing unused accounts, adjusting entitlements when roles change, and retiring outdated automations help you maintain a predictable and justifiable access footprint across your organization.

Detecting Insider Activity Across Users and Systems

Detection depends on the ability to observe patterns, not just isolated events. When organizations understand what normal activity looks like, they are better equipped to detect deviations that may signal risk. These deviations might include attempts to access unfamiliar systems, changes in file access patterns, or sequences of commands that don't fit a user's typical workflow.

Automated systems can contribute valuable context here. For example, if an agent suddenly interacts with systems outside its assigned workflow or performs actions at volumes far above normal, those signals are meaningful. Even though these agents follow structured logic, they can still act in unexpected ways if their configuration changes or if compromised credentials direct them to new tasks.

Context enriches the detection process. A resignation notice, a major project deadline, or a role transition can all influence behavior. Monitoring solutions that combine activity patterns with contextual understanding help security teams distinguish between unusual but legitimate work and activity that warrants deeper investigation.

Investigating Insider Activity

When activity appears risky, a structured investigation helps teams determine what happened and why. Investigations often start with a timeline: a sequence of events that includes logins, data access, system changes, and any actions taken by automated processes. This timeline becomes the backbone of the inquiry, allowing analysts to understand whether an action was intentional, accidental, or triggered through compromised access.

During this process, it's helpful to consider whether automated systems or AI-driven workflows contributed to the event. Sometimes an agent may carry out a task based on a prompt or input from a compromised account, or its configuration may have changed unintentionally. Including these systems in the analysis ensures investigators see the full picture instead of focusing only on human activity.

Ultimately, findings should be documented in clear, actionable terms. A complete record explains what occurred, the impact on systems or data, and which controls may need adjustment. These insights support future improvements across access policies, monitoring rules, and user education.

Figure 2. A typical progression of actions taken by intruders after acquiring valid credentials.

Compromised Credentials: A Common Path for Intruders

Many intrusions begin with compromised credentials. Attackers frequently rely on passwords, tokens, or session cookies that belong to legitimate users. Once inside, they behave in ways that resemble real activity, making it difficult to identify them through simple rule-based monitoring.

Compromised credentials can come from phishing, malware, weak password practices, or unattended devices. Service accounts and AI-related automations aren't immune. These accounts often hold broader permissions and rely on API keys or tokens that, if exposed, provide attackers with consistent entry points.

The danger lies in how these intruders behave once they gain access. They may move laterally, gather information about users, create new accounts, or escalate their privileges quietly. Detecting these actions requires monitoring that highlights unusual patterns, especially when a normally predictable account begins operating differently.

To reduce the impact, organizations benefit from regular entitlement reviews, credential rotation, and monitoring that highlights unexpected authentication attempts or shifts in activity patterns. Investigating anomalies early prevents attackers from maintaining long-term access under the disguise of valid accounts.

Training and Awareness

Training plays a meaningful role in reducing insider risk. When employees understand the types of actions that create exposure, they're better equipped to avoid mistakes and report concerning activity. Scenario-based examples, such as improper file sharing or suspicious email behavior, make risks tangible and easier to recognize.

As more workflows incorporate AI tools, training should also help teams understand how these systems work with data, including what they can access and how unexpected inputs or misconfigurations may create risk. Simple guidance on safe usage, data handling, and reporting concerns strengthens your organization's overall posture.

Awareness becomes an ongoing discipline rather than a one-time requirement. Frequent updates on common risks and emerging patterns encourage teams to maintain good habits as the environment changes.

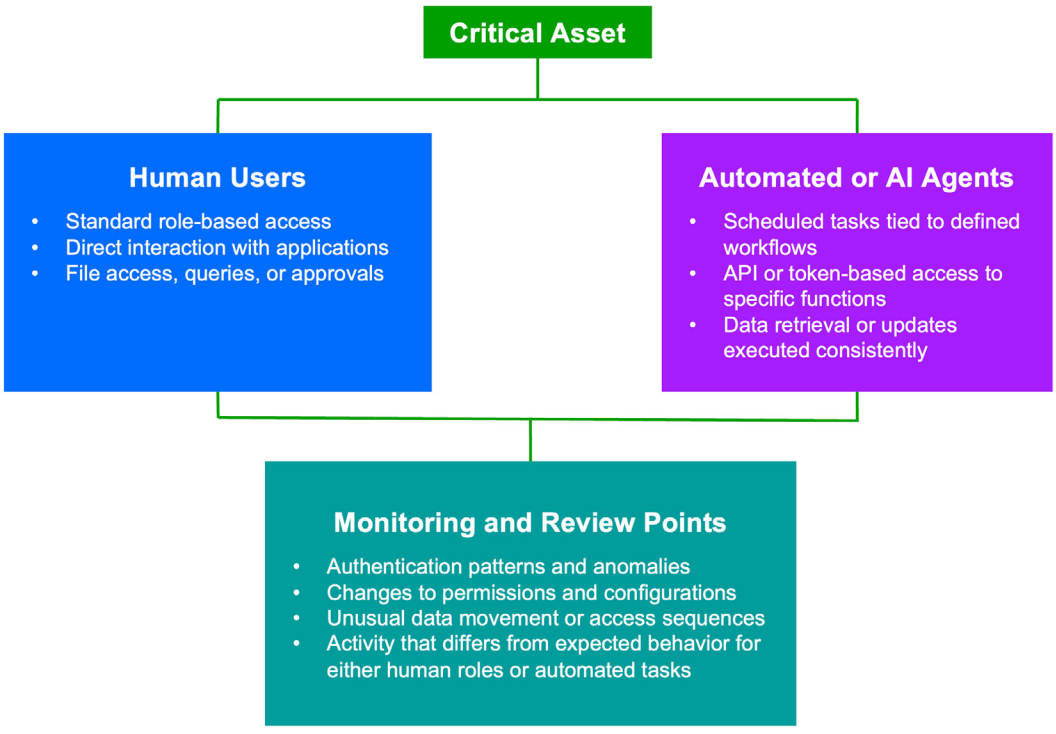

Figure 3. How people and automated systems interact with critical assets and the monitoring points that help identify unusual or high-risk activity.

Protecting Critical Assets

To focus resources effectively, your organization needs a clear understanding of which assets carry the greatest value or pose the highest impact if misused. Critical assets might include financial systems, customer information, proprietary research, or operational platforms.

Mapping how humans and automated systems interact with these assets helps reveal where risk concentrates. Some systems see frequent human interaction while others rely primarily on service accounts or AI-powered processes. Understanding these patterns allows teams to adjust monitoring, strengthen access controls, and apply additional safeguards where needed.

Over time, these maps help identify gaps, such as assets that accumulate new users without corresponding reviews or automations that continue to run after their original purpose has expired. Addressing these issues promptly prevents needless exposure.

Program Review and Reporting

As insider risk programs mature, regular reviews ensure the program keeps pace with changes in technology, staffing, and organizational priorities. These reviews assess how well controls are working, whether the monitoring strategy remains aligned with current risks, and how often incidents require deeper intervention.

Reporting provides visibility beyond your security team. When leadership receives consistent updates on trends, emerging issues, and the performance of controls, it becomes easier to support investments and process improvements. These reports also provide an opportunity to highlight new developments, such as the introduction of additional AI-based workflows that require monitoring or revised access expectations.

A program that evolves in this way is durable. It adapts to new tools, new access patterns, and new forms of risk while maintaining a consistent approach to managing insider activity.

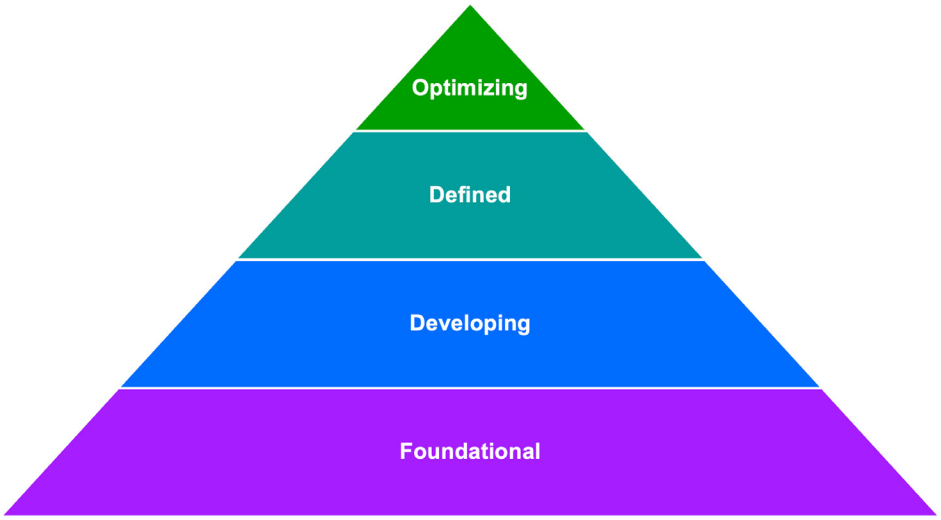

Figure 4. Stages of insider risk program maturity showing how capabilities progress from foundational practices to an optimized model.

About Exabeam

Exabeam is a leader in intelligence and automation that powers security operations for the world's smartest companies. As a global cybersecurity innovator, Exabeam provides industry-proven, security-focused, and flexible solutions for faster, more accurate threat detection, investigation, and response (TDIR).

exabeam

Learn more at www.exabeam.com →