Building a Behavior-Driven Insider Threat Program: A 10-Step Playbook

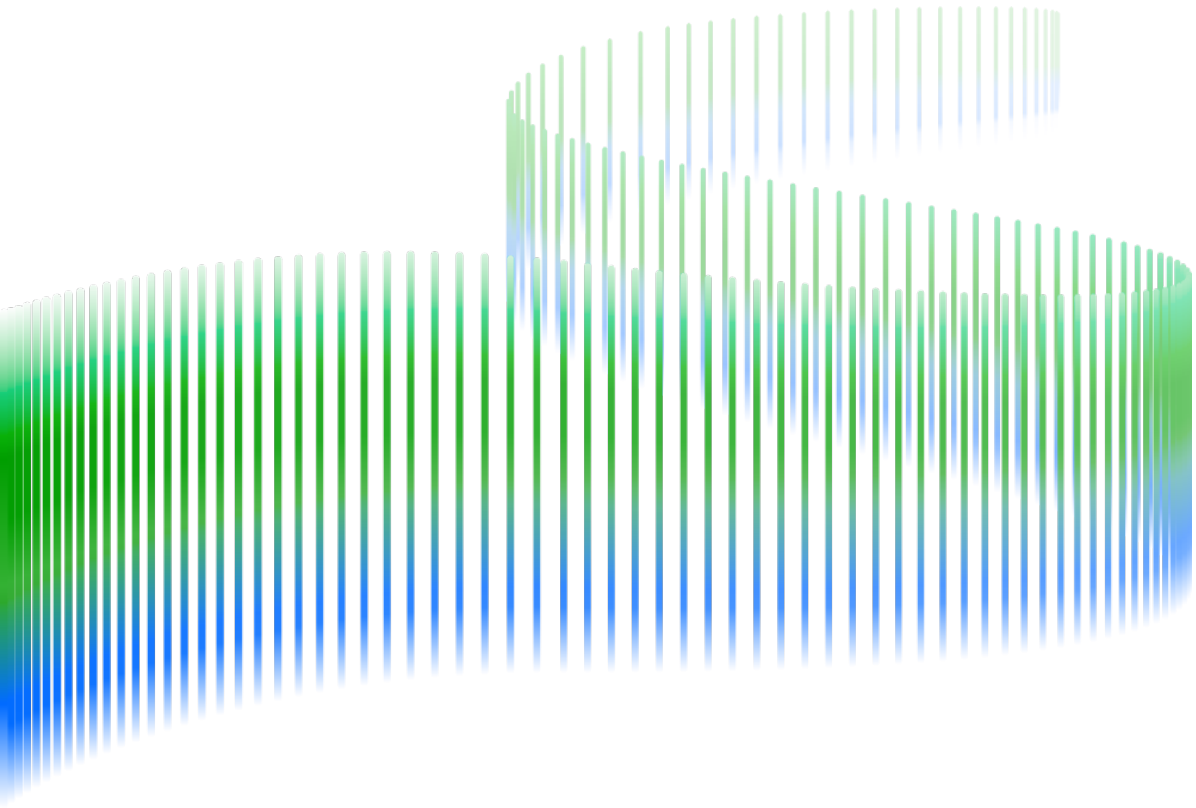

Table of Contents

Abstract

Insider threats are among the most complex and costly risks facing large organizations. The definition of "insider" now extends beyond employees to include contractors, partners, service accounts, automation, and AI agents. These identities can misuse legitimate access or become compromised by external attackers.

Risk spans malicious insiders, negligent users, and hijacked accounts. As AI-driven automation scales, non-human identities introduce insider risk at machine speed, often through activity that appears legitimate.

Security leaders need a structured approach to detect behavioral risk early, respond effectively, and continuously improve.

This playbook provides a prescriptive framework to design, implement, and mature an insider threat program based on behavior-driven risk management. It combines best practices from the CERT Insider Threat Center and CISA with Exabeam capabilities, including:

- User and entity behavior analytics (UEBA)
- Agent Behavior Analytics (ABA)
- Dynamic risk scoring
- Outcomes Navigator

This framework helps organizations govern human and non-human insiders, prioritize critical assets, and detect behavioral patterns that static controls miss.

Introduction

Insider threats are widespread, costly, and difficult to detect. Organizations continue to face high incident rates, long dwell times, and rising financial impact. Most incidents are not driven by malicious intent but by negligence, misuse, or compromised access.

These trends highlight a core challenge: Traditional, policy-based controls were not designed to detect how insider threats unfold.

Tools such as data loss prevention (DLP) and static correlation rules can identify clear violations. However, they often miss behavioral patterns like data staging, privilege misuse, and abnormal access over time.

This challenge is growing. Insider risk now includes service accounts, automation, and AI agents operating at scale. Many actions appear legitimate and only become risky when viewed in context.

The State of Insider Risk

Insider incidents are widespread

83%

83% of organizations experienced at least one insider incident in the past year[1] and 93% say insider threats are as difficult or harder to detect than external attacks[2].

Slow detection and containment

81 days

Organizations take an average of 81 days to contain an insider incident, allowing risk to persist without early detection.[3]

Soaring cost of insider risk

$17.4M

The average annual cost of insider incidents is $17.4 million per organization, reflecting the operational and financial impact of delayed response.[3]

Most incidents are not malicious

55%

More than half of insider incidents are caused by negligent employees, not malicious actors.[3]

Figure 1.

Insider risk is prevalent, costly, and difficult to contain. Most organizations experience insider incidents, face extended containment timelines, and incur significant cost, with many incidents driven by negligent or compromised behavior rather than malicious intent.

[1] Cybersecurity Insiders, 2024 Insider Risk Report

[2] Cybersecurity Insiders, 2025 Insider Risk Report

[3] Ponemon Institute, 2025 Cost of Insider Risks

To address this shift, organizations must adopt a behavior-driven approach that combines UEBA, ABA, multi-layered risk scoring, and cross-functional governance.

This playbook outlines how to move from reactive controls to a proactive insider threat program that detects risk early, prioritizes high-impact activity, and scales across human and non-human identities.

These trends highlight a core challenge: Traditional, policy-based controls were not designed to detect how insider threats unfold.

Tools such as data loss prevention (DLP) and static correlation rules can identify clear violations. However, they often miss behavioral patterns like data staging, privilege misuse, and abnormal access over time.

This challenge is growing. Insider risk now includes service accounts, automation, and AI agents operating at scale. Many actions appear legitimate and only become risky when viewed in context.

To address this shift, organizations must adopt a behavior-driven approach that combines UEBA, ABA, multi-layered risk scoring, and cross-functional governance.

This playbook outlines how to move from reactive controls to a proactive insider threat program that detects risk early, prioritizes high-impact activity, and scales across human and non-human identities.

Before You Begin: Identify Critical Assets

Start by defining what needs protecting.

Identify sensitive data, critical systems, and key business processes. Include traditional assets such as customer data, source code, financial systems, and identity infrastructure. Also account for AI-enabled workflows, autonomous agents, service accounts, and automated data pipelines.

Focus the program on scenarios that create real business impact instead of treating all policy violations as equal.

Build Your Insider Threat Program

Ten steps to detect, investigate, and contain insider risk across the agentic enterprise.

Step 1
Establish Governance and Cross-Functional Alignment
Assemble a multidisciplinary insider threat team with executive sponsorship and a clear program charter.

Step 2
Secure Identity Lifecycle Processes
Develop strong onboarding and offboarding procedures to reduce insider risk from joiners and leavers.

Step 3
Build an Identity-Centric Visibility and Data Foundation
Integrate key logs across identity, endpoint, cloud, network, and HR systems to monitor behavior and establish baselines.

Step 4
Move Beyond DLP to Behavior-Driven Detection
Go beyond static rules and DLP alone to detect contextual behavior patterns and anomalies that reveal real risk.

Step 5
Deploy UEBA and ABA for Proactive Detection
Use behavioral analytics to monitor users and non-human identities for anomalous activity and emerging risk.

Step 6
Use Dynamic Risk Scoring and Build Risk Narratives
Turn disjointed alerts into prioritized risk stories with dynamic risk scoring and linked event context.

Step 7
Develop Investigation Workflows and Response Playbooks
Predefine escalation paths, evidence requirements, and containment actions for insider threat response.

Step 8
Extend Coverage to the Agentic Enterprise
Expand the program to cover contractors, vendors, service accounts, and AI agents with tailored monitoring and controls.

Step 9
Measure Coverage and Continuously Improve
Map data sources and analytics content to insider threat use cases, identify gaps, and improve over time.

Step 10
Track Metrics and Demonstrate Business Value
Measure dwell time, coverage, incidents prevented, and business impact to show progress and ROI.

Figure 2.

This 10-step model moves from identifying critical assets to governance, identity controls, behavior-driven detection, agentic enterprise coverage, continuous improvement, and measurable business value. The following sections detail each step and show how Exabeam capabilities support implementation, along with an insider threat matrix and real-world examples.

This playbook outlines how to move from reactive controls to a proactive insider threat program that detects risk early, prioritizes high-impact activity, and scales across human and non-human identities.

These trends highlight a core challenge: Traditional, policy-based controls were not designed to detect how insider threats unfold.

Tools such as data loss prevention (DLP) and static correlation rules can identify clear violations. However, they often miss behavioral patterns like data staging, privilege misuse, and abnormal access over time.

This challenge is growing. Insider risk now includes service accounts, automation, and AI agents operating at scale. Many actions appear legitimate and only become risky when viewed in context.

To address this shift, organizations must adopt a behavior-driven approach that combines UEBA, ABA, multi-layered risk scoring, and cross-functional governance.

This playbook outlines how to move from reactive controls to a proactive insider threat program that detects risk early, prioritizes high-impact activity, and scales across human and non-human identities.

Step 1: Establish Governance and Cross-Functional Alignment

To build a strong governance foundation, establish a cross-functional insider risk steering committee and formal program charter from the outset. This team should include senior representation from cybersecurity, physical security HR, legal, compliance, privacy, risk, data, and AI governance, where applicable.

Executive sponsorship is critical. Support from the CIO, CISO, or CEO signals that insider risk is a business priority, not just a security operations concern.

Define the program's mission, scope, authorities, escalation model, privacy safeguards, and categories of behavior to monitor. These may include fraud, intellectual property theft, sabotage, negligence, account compromise, and AI or agent misuse.

Clarify roles and escalation paths in advance:

- When does HR get involved?
- When does legal intervene?
- Under what conditions is law enforcement engaged?

For AI-enabled processes and autonomous agents, assign explicit ownership:

- Who approves new agents
- Who maintains access
- Who reviews behavioral anomalies
- Who can disable an agent
- Who leads response efforts

Clear ownership prevents accountability gaps as insider risk extends beyond employee activity.

Exabeam integrations with HR, identity, cloud, endpoint, and AI data sources align monitoring with governance policies and provide the context needed to distinguish routine activity from risky behavior.

Step 2: Secure Identity Lifecycle Processes

Identifying lifecycle management is a primary control point for insider risk.

Secure Onboarding

Define ownership, purpose, and access scope for every identity, including service accounts and AI agents. Apply least-privilege access and enable logging from the start.

Work with HR to identify high-risk roles or individuals, such as privileged users or employees entering from competitors. Provide clear guidance on acceptable use and data handling.

Apply the same rigor to non-human identities. Service accounts, automation, copilots, and AI agents should have:

- A named business owner
- A defined purpose
- Documented access scope
- Logging enabled from day one

Exabeam integrates with HR and identity and access management (IAM) systems to automatically import identities and establish baseline risk profiles. Threat Center watchlists help monitor new identities during onboarding, when normal behavior is still being established.

Rigorous Offboarding

When employees leave, the risk window increases, especially in involuntary or sensitive exits.

Collaborate with HR and IT to create an offboarding checklist that includes:

- Immediate access revocation
- Exit interviews
- Reminders of ongoing obligations, such as NDAs
- Post-departure monitoring for residual or unauthorized account activity

Offboarding should extend to:

- Contractors and vendors
- Service accounts
- AI agents tied to users, projects, or workflows

Exabeam Nova can assign a high risk score to activity that occurs after an offboarding date and trigger alerts for investigation. Teams can also use Threat Center watchlists to monitor identities during resignation, termination, or layoff periods.

Strong lifecycle controls reduce exposure across human and non-human identities. In many real incidents, incomplete offboarding leaves accounts active longer than intended. In environments that include automation and AI, this risk extends to service accounts, API tokens, AI assistants, and automated workflows that continue running after the business need ends.

Step 3: Build an Identity-Centric Visibility and Data Foundation

With governance and identity processes in place, the next step is to establish comprehensive visibility into identity activity.

Insider threats manifest as patterns across systems over time. That means visibility must extend across endpoints, networks, applications, cloud services, and AI-enabled workflows, and all activity must be correlated by identity.

Start by integrating critical data sources into a centralized analytics platform, such as a SIEM with UEBA and ABA capabilities:

- **Authentication and identity logs:** Active Directory, SSO, VPN access, and badge entry systems to track who is accessing what and when
- **Endpoint and data logs:** EDR tools, DLP alerts, file access logs, USB usage, and cloud storage activity to detect file movement, data exfiltration, or privilege escalation
- **Network telemetry:** Proxy, DNS, firewall, and VPN data to identify connections, data transfers, and access to risky or unauthorized destinations
- **Communication and collaboration systems:** Email, messaging platforms, code repositories, ticketing systems, and file-sharing tools to identify unusual sharing, insider coordination risks, or unauthorized data movement
- **AI and agent activity:** Copilots, chatbots, LLM interactions, automation tools, and agent telemetry to monitor prompts, tool calls, agent actions, data access, and policy violations

All this data must be normalized and correlated at the identity level.

For human users, this means linking multiple accounts, devices, sessions, and applications to a single identity. For non-human identities, it means mapping service accounts, API tokens, bots, and AI agents to a business owner, application, or workflow.

New-Scale Fusion and the Exabeam Common Information Model (CIM) fuse these event streams into a unified identity timeline. This allows a wide range of activity, such as a Windows login, SaaS data export, VPN session, cloud configuration change, and AI agent tool call, to be analyzed together in context.

This identity-centric view is critical. Many insider behaviors appear normal within a single system but become risky when correlated across systems.

After data integration, establish baselines of normal behavior for users, entities, and agents. This includes:

- Typical login times and locations
- File access and data usage patterns
- Application and system usage
- Peer group behavior
- Expected activity for service accounts and AI agents

For non-human identities, baseline behavior includes the systems accessed, APIs called, tools used, frequency of activity, and volume of data processed.

Exabeam automatically builds these baselines and detects deviations over time. This enables the program to identify subtle behavioral drift that indicates misuse, compromise, or policy violations.

Step 4: Move Beyond DLP to Behavior-Driven Detection

With data pipelines in place, shift your detection approach from static rules and DLP toward behavior-driven detection.

Traditional tools monitor specific events or content and trigger alerts when predefined rules are met, such as detecting an email that contains sensitive data. While this remains useful, it does not fully address insider risk. Human insiders, compromised accounts, and AI-enabled workflows can all perform actions that appear legitimate in isolation.

Risk emerges from sequence, timing, context, and deviation from expected behavior.

The limitations of static approaches are evident in real-world insider incidents. In the Snowden case, the insider used valid credentials and legitimate backup processes to exfiltrate classified data. Without behavioral context, this type of activity is difficult to detect using DLP or static correlation rules alone.

This challenge applies to AI-enabled environments, as well. A user may prompt an approved AI assistant to summarize sensitive files. A service account or agent may invoke tools it does not typically use or access a new data source. Individually, these actions may appear normal. When evaluated together in context, they can indicate misuse, compromise, or policy evasion.

DLP and content controls remain important for detecting known violations and enforcing policy, but they cannot determine intent or identify behavioral drift. To address modern insider threats, organizations must adopt a behavior-driven approach that incorporates UEBA and ABA to:

- Learn baseline behavior for users, entities, and agents
- Identify anomalies and deviations from expected patterns
- Detect misuse of human and non-human identities before it results in data loss or operational impact

Step 5: Deploy UEBA and ABA for Proactive Detection

UEBA should be the core detection engine of your insider threat program.

UEBA applies machine learning to identify behavioral anomalies such as:

- Accessing significantly more data than usual
- Working outside established patterns or hours
- Deviations in network or system usage

Individually, these actions may not appear suspicious. When evaluated together, they can indicate emerging insider risk

Exabeam includes a broad range of behavioral models and rules that cover scenarios such as data staging, privilege misuse, data aggregation, and indicators of credential compromise. This allows teams to establish broad detection coverage without building custom logic from scratch.

Each anomaly contributes to a dynamic risk score assigned to the associated user or entity. As anomalies accumulate, the risk score increases. This lets the system prioritize activity based on overall risk rather than isolated alerts, reducing noise while ensuring high-risk patterns are surfaced.

ABA extends this capability to non-human identities, including service accounts, automation, bots, and AI agents. These identities often have broad access and operate at scale, making their behavior critical to monitor.

ABA profiles expected activity for these entities and detects deviations such as:

- Access to new systems or data sources
- Unusual tool execution
- Increased data movement
- Changes in workflow behavior

This approach enables early detection of misuse, compromise, or misconfiguration in both human and non-human identities.

By deploying UEBA and ABA together, organizations move from reacting to predefined alerts to continuously evaluating behavioral risk.

Step 6: Use Dynamic Risk Scoring and Build Risk Narratives

Security operations teams often face large volumes of alerts, many of which are low priority or lack sufficient context. The objective of an insider threat program is to surface meaningful risk, not generate more alerts.

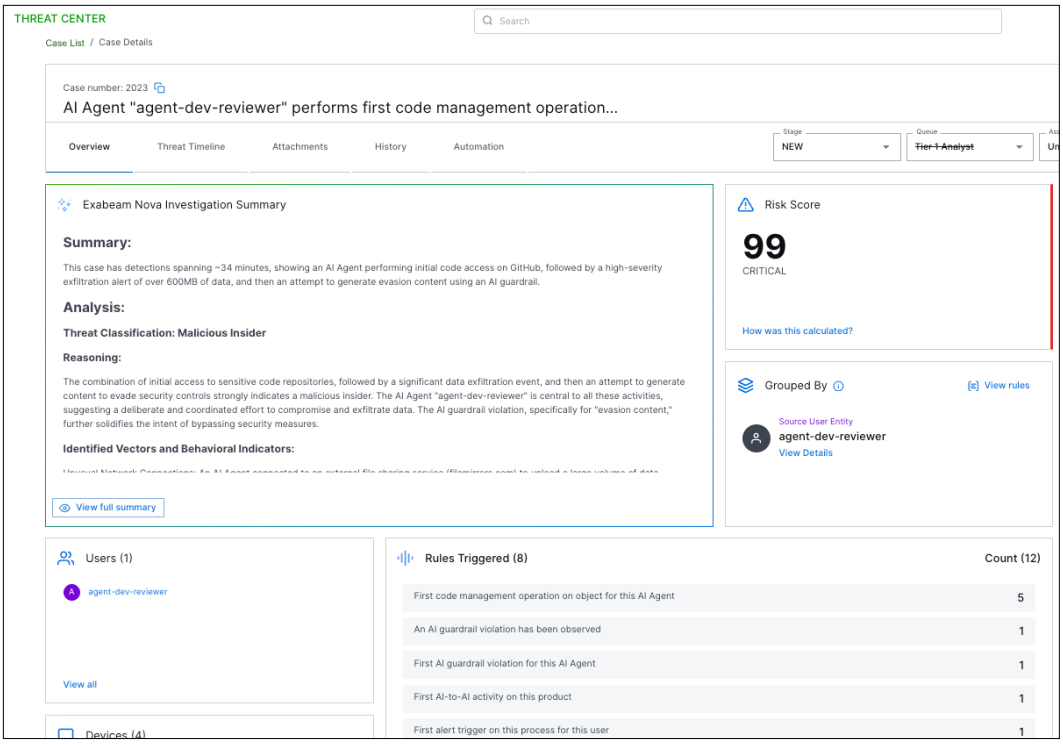

Figure 3.

Threat Center summarizes the agent-driven incident, assigns a critical risk score, groups activity by the source entity, and presents the triggered rules analysts need to investigate and coordinate response.

Dynamic risk scoring continuously updates the risk level for each identity based on observed behavior. Instead of generating multiple independent alerts, the system aggregates related anomalies and increases risk scores as suspicious activity accumulates.

This is especially important for behavioral drift, where a sequence of small changes signals increasing risk. For example:

- A first-time login from an unusual location
- Access to a system not previously used
- Increased volume of file access
- Activity outside normal working hours

Individually, these events may not trigger concern. Together, they form a meaningful pattern that should rise above the noise.

Risk Narratives and Timelines

Risk scoring is most useful when paired with investigation context. In New-Scale Fusion, timelines bring together relevant events, detections, and risk scores in chronological order. This transforms a vague alert such as "multiple anomalies detected" into a usable investigation path. For example, a timeline might show that a user logged in from a new location, accessed a sensitive database not normally used, and transferred large files to an external drive, resulting in a high risk score. That sequence gives analysts a structured place to start the investigation.

This is important because security teams are often overloaded with alerts, and many go uninvestigated because of volume and false positives. Dynamic risk scoring and narrative-based investigation help reduce that burden by focusing analyst time on a smaller set of high-risk identities rather than a large pool of disconnected, low-confidence alerts.

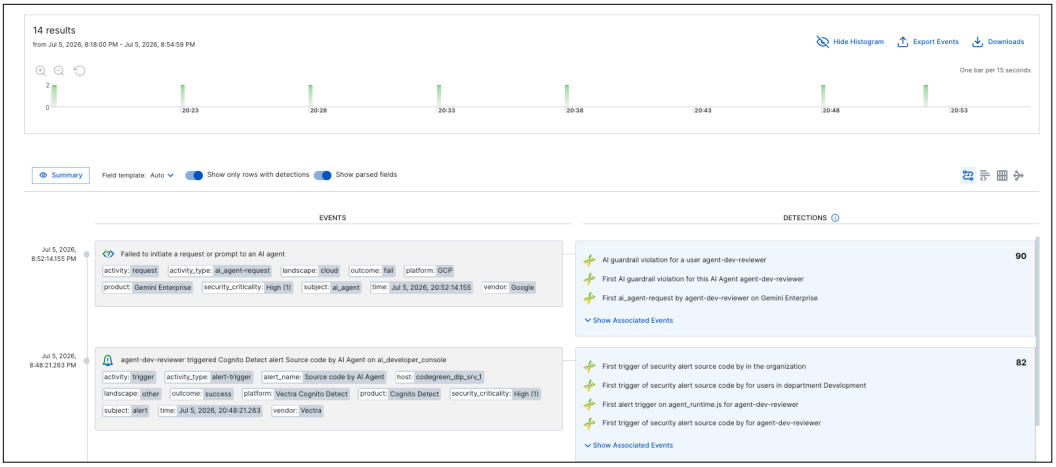

Figure 4.

Timelines correlate multiple AI agent activity events and detections, showing how an agent's code management activity, data access, and guardrail violations build into a higher-risk behavioral pattern.

Behavioral Risk Narrative Timeline

How separate signals become one prioritized insider risk story

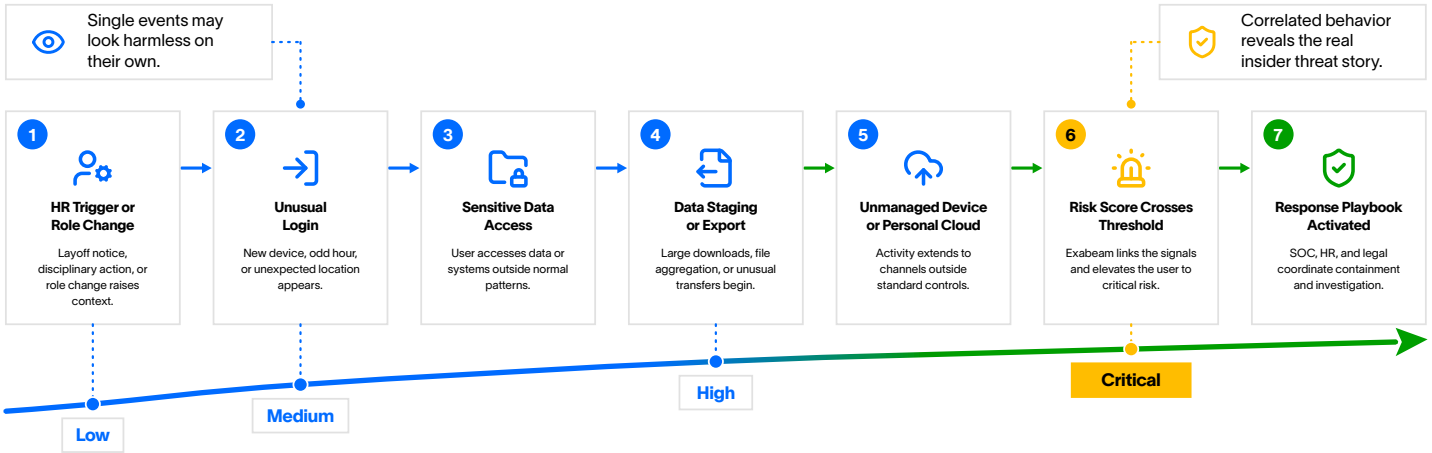

Behavioral analytics, dynamic risk scoring, and automated timelines help analysts move from isolated alerts to fast, confident response.

Figure 5.

Disconnected events, such as HR context, unusual login activity, sensitive data access, and data staging, become a prioritized risk story through behavioral analytics, dynamic risk scoring, and coordinated response playbooks.

Step 8: Extend Coverage to the Agentic Enterprise

Insider threat programs must extend beyond employees to cover all identities with access to systems and data. This includes:

- **Contractors and third-party vendors:** Often have access to internal systems and sensitive data, sometimes with less oversight than employees
- **Service and application accounts:** Can be hijacked or abused to perform insider-like actions
- **AI assistants and automated agents:** Can be manipulated by malicious users, compromised by external attackers, or misconfigured in ways that create risk.

Extend monitoring and policies to all these identities. Third-party accounts should be onboarded into monitoring systems, tagged appropriately, and scored using stricter context where warranted. ABA should be applied to service accounts and AI agents by profiling expected behavior, including which database tables they query, which APIs they call, which tools they invoke, and what data volumes are normal. This gives security operations teams visibility into human-to-agent, agent-to-agent, and agent-to-system workflows as the organization becomes more automated.

Your Step 8 planning should also incorporate third-party risk management. Coordinate with procurement, supply chain, and IT to define minimum security requirements for any vendors with internal access, such as background checks, security training, and contractual data handling clauses. Bringing these identities into scope helps ensure the program does not miss insider risk vectors outside the traditional employee population.

Step 9: Measure Coverage and Continuously Improve

Insider threat programs must continue to evolve as business processes change, threats shift, and AI adoption expands. Use Outcomes Navigator to:

- **Map detection coverage to use cases:** Map ingested data and active content, including rules and models, to insider threat use cases and relevant MITRE ATT&CK® techniques. It helps teams understand where they have strong coverage and where gaps remain, including areas such as data exfiltration, privilege abuse, fraud, account compromise, AI misuse, shadow AI discovery, or risky behavior by non-human identities.
- **Identify gaps in data sources or detection logic:** Receive guidance for improving coverage. Outcomes Navigator can recommend additional log sources or enable specific analytics content to cover an uncovered technique, accelerating tuning and improving your program's ability to detect insider threats without leaving data or content visibility gaps unresolved.
- **Validate data quality:** Track whether key log sources are functioning and properly parsed. If coverage is low because a source is missing or misconfigured, Outcomes Navigator highlights the problem so the team can correct ingestion or enrichment issues.

For AI and non-human identities, this review should also include telemetry sources, behavioral anomalies, policy violations, shadow AI usage, service account drift, and high-risk agent workflows.

Regular review of coverage scores and recommended changes helps demonstrate a continuous improvement cycle and makes the program easier to defend with leadership and internal stakeholders.

Step 10: Track Metrics and Demonstrate Business Value

As the program matures, measure performance and communicate results to stakeholders.

Define key metrics that reflect both operational performance and business impact, such as:

- Time to detect and contain insider incidents
- Number of incidents detected, contained, and prevented
- Reduction in false positives, duplicate alerts, and low-confidence investigations
- Coverage improvements across insider threat, compromised account activity, AI misuse, and non-human identity use cases
- Program adoption and usage metrics, such as data sources integrated, employees trained, service accounts onboarded, AI agents monitored, shadow AI discoveries, agent policy violations, and high-risk population trends

Translating technical metrics into business terms is essential. Reports to the CISO, risk committees, or the board should describe outcomes such as faster containment, better coverage, fewer unmanaged high-risk identities, stronger AI governance, and reduced exposure around critical assets. This business framing is central to showing program maturity and securing ongoing investment.

Outcomes Navigator and Exabeam Nova Advisor Agent can produce executive-ready reports that benchmark detection coverage and track progress. Exabeam Nova peer benchmarking can show leaders how insider threat coverage compares with industry peers, adding credibility to budget and staffing discussions.

Over time, using these metrics to communicate progress helps advance the program from foundational maturity to a more sophisticated operating model in which insider risk management becomes part of the organization's broader security culture.

Insider Threat Matrix

Types of Insiders vs. Detection and Response

The nature of insider incidents varies widely. This section maps key insider threat categories to typical indicators, response approaches, and how Exabeam supports detection and response for each type.

The matrix includes five primary categories:

1. Malicious insiders
2. Negligent insiders
3. Compromised accounts
4. Third-party insiders
5. AI-enabled or autonomous agent insiders

The AI-enabled insider category is separated to reflect how autonomous agents and AI-driven workflows introduce different visibility, ownership, and response considerations compared to traditional users and third parties.

Each category includes:

- Common behavioral indicators
- Recommended response actions
- Relevant Exabeam capabilities for detection and investigation

This structured view helps teams align detection and response strategies to specific insider behaviors rather than applying a one-size-fits-all approach.

Insider Threat Matrix

Types of insiders, signals, response, and how Exabeam helps

1 Malicious Insider

Who it is

Trusted employee or user acting with intent, such as personal gain, retaliation, fraud, or espionage

High-risk signals

- Behavior shifts after a trigger event
- Access outside role or time
- Data staging
- Attempts to cover tracks or bypass controls

Response approach

Run a covert investigation with SOC, HR, and legal on a need-to-know basis. Preserve evidence. Prepare lockout, device recovery, legal hold, and possible law enforcement engagement.

How Exabeam helps

Risk modeling and anomaly correlation connect HR context, data access spikes, unusual locations, and data movement. Timelines reveal multi-step patterns that single DLP rules miss.

2 Negligent Insider

Who it is

Well-meaning user with poor security judgment or repeated policy violations

High-risk signals

- Uses personal cloud or email for convenience
- Repeated minor DLP alerts
- Falls for phishing
- Weak security practices, attempts to work around controls

Response approach

Treat as coaching first. Involve the manager or HR. Improve controls and apply stricter controls if the pattern persists.

How Exabeam helps

Behavioral risk accumulation adds points for repeated low-severity issues. Contextual risk factors such as privilege level and access to sensitive data help prioritize review.

3 Compromised Account

Who it is

Legitimate credentials hijacked by an external threat actor

High-risk signals

- Impossible travel, new device or region
- Unusual login time
- Sudden access to large datasets or admin consoles
- Machine-speed activity that is abnormal for the user

Response approach

Treat as an external breach. Disable or lock the account, force password reset and MFA, scan devices, and investigate downstream access.

How Exabeam helps

UEBA identifies deviations in login, device, location, session, and access behavior. Automation can revoke access or trigger MFA when risk becomes critical.

4 Third-Party Insider

Who it is

Vendor, contractor, partner, or outsourced user with legitimate internal access

High-risk signals

- Access outside agreed scope or hours
- Activity from unusual locations
- Use of shared or stale accounts
- Data access inconsistent with contract or project role

Response approach

Disable or limit access while investigating. Coordinate with procurement, legal, and the business owner. Treat malicious misuse as an insider case and compromise as a breach.

How Exabeam helps

Identity monitoring and risk scoring can tag third parties, compare behavior to expected scope, and elevate anomalies from external identities that touch sensitive assets.

5 AI-Enabled or Autonomous AI Insider

Who it is

AI assistant, copilot, chatbot, service account, or autonomous agent acting with delegated access

High-risk signals

- New tool invocation, unexpected data access
- Prompt or workflow misuse, agent policy violations
- Shadow AI usage
- High-volume queries or actions outside normal workflow

Response approach

Disable or constrain the agent, revoke tokens, preserve prompts and tool-call evidence, involve the app owner and AI governance team, and adjust permissions, rate limits, and guardrails.

How Exabeam helps

ABA baselines normal agent behavior and detects behavioral drift across tools, data access, prompts, and actions. Risk scoring prioritizes high-risk agent workflows and non-human identity anomalies.

Insider risk now spans employees, third parties, and non-human identities.
Effective programs combine visibility, behavioral analytics, and response playbooks.

Figure 6.

Insider risk categories from malicious and negligent users to compromised accounts, third parties, and autonomous AI insiders. This matrix shows the signals, response approach, and Exabeam capabilities that help detect and contain each scenario.

Real-World Scenarios: Bringing the 10-Step Program to Life

The following scenarios illustrate how a behavior-driven insider threat program operates in practice. Each scenario shows how governance, identity controls, visibility, behavior analytics, risk scoring, and response playbooks work together to identify and contain risk.

Governance is included in every example because insider incidents typically require coordination by security, HR, legal, risk, IT, and, in AI-related scenarios, AI governance teams.

Scenario 1: RIF Offboarding Aftershock

During a large workforce reduction, a recently terminated engineer attempts to access the source code repository using VPN credentials one day after departure.

Because the organization followed structured offboarding processes, most access was revoked immediately. However, the individual successfully logs in through a previously overlooked cloud administrator account.

Detection

Exabeam identifies this post-termination activity as a high-severity anomaly. Behavioral analytics flags the login and assigns a critical risk score because the activity occurs after termination and falls outside expected behavior.

Response

A predefined response playbook is triggered. Security, HR, and legal teams coordinate to revoke access, contain the account, and preserve evidence.

Outcome

The activity is contained quickly with minimal impact. The organization updates its offboarding procedures to ensure that all accounts, including cloud administrative accounts, are deactivated without delay.

Related Steps

- Step 2: Offboarding
- Step 5: Behavioral Detection
- Step 6: Risk Scoring and Timeline Analysis
- Step 7: Response Playbooks

These steps work together to identify post-termination access, elevate the risk through behavioral context, and trigger a coordinated response across teams.

Scenario 2: The Negligent Data Leak

A business analyst transfers sensitive files to a personal email account and later uploads a customer dataset to a personal cloud storage platform to work remotely.

Detection

Exabeam correlates multiple events over time:

- DLP alerts for external email usage
- Creation of an unusual email forwarding rule
- Access from an unmanaged device or service
- After-hours activity

Dynamic risk scoring aggregates these behaviors into a single high-risk profile.

Investigation

Timeline analysis shows the sequence of events, allowing analysts to understand the full behavior pattern rather than reviewing isolated alerts.

Response

The issue is classified as negligent behavior. HR and management address the situation through training and policy reinforcement. Controls are updated to prevent similar activity.

Outcome

Risk is mitigated before significant data exposure occurs.

Related Steps

- Step 3: Visibility and Data Foundation
- Step 5: Behavioral Detection
- Step 6: Risk Scoring and Timeline Analysis
- Step 7: Response Playbooks

These steps combine visibility with behavioral correlation and risk scoring to surface low-and-slow policy violations before they escalate into data loss.

Scenario 3: Malicious Privileged Insider

A database administrator attempts to exfiltrate sensitive customer data and delete backups after being reassigned from a critical project.

Detection

UEBA detects abnormal database access patterns, including off-hours queries and access to data not previously used. A rule-based detection also flags attempts to delete system logs.

Dynamic risk scoring combines these signals into a high-priority alert.

Investigation

Analysts review the timeline to confirm coordinated suspicious activity. A controlled investigation begins involving security, HR, and legal.

Response

Access is revoked using automated response actions. The system is secured and the individual is removed from the environment.

Outcome

The attempted data theft and sabotage are prevented before escalation.

Related Steps

- Step 1: Governance and Cross-Functional Alignment
- Step 5: Behavioral Detection (UEBA)
- Step 6: Risk Scoring and Timeline Analysis
- Step 7: Response Playbooks

These steps connect governance, behavioral analytics, and coordinated response to detect and contain malicious insider activity before damage occurs.

Scenario 4: Compromised Employee Account

A finance employee's credentials are stolen through a phishing attack. An external actor uses the credentials to access financial systems and begin extracting data.

Detection

Behavioral analytics detects:

- Login from an unusual geographic region
- Activity outside normal working hours
- Increased data access inconsistent with the user's role

Risk scoring escalates quickly as anomalous behavior continues.

Response

Automated actions revoke access and force credential resets. The endpoint is investigated and malicious software is identified and removed.

Outcome

The incident is contained before significant data loss occurs, demonstrating the importance of behavioral detection beyond static controls.

Related Steps

- Step 4: Beyond DLP and Static Detection
- Step 5: Behavioral Detection (UEBA)
- Step 6: Risk Scoring and Timeline
- Step 7: Response Playbooks

These steps detect credential misuse through behavioral deviation, escalate risk dynamically, and enable rapid containment of compromised accounts.

Scenario 5: Outsider Infiltrator and Malicious AI Agent

An organization faces two simultaneous risks:

- A contractor accessing data outside of assigned responsibilities
- An AI-powered support bot manipulated to extract sensitive information

Detection

Contractor

Behavioral analytics identifies abnormal data access patterns and unusual access locations. The contractor is flagged as high risk through dynamic risk scoring and watchlist monitoring.

AI Agent

ABA detects abnormal activity from the support bot, including large-scale data queries and behavior outside its normal workflow.

Response

Contractor

Security, HR, and legal initiate a controlled investigation.

AI Agent

Security operations coordinates with IT and AI governance to disable the agent, preserve evidence, and adjust permissions.

Outcome

Both threats are contained. The organization validates detection coverage using Outcomes Navigator and strengthens controls for third-party access and AI systems.

Related Steps

- Step 2: Identity and Access Management (Onboarding and Access Control)
- Step 5: Behavioral Detection (UEBA and ABA)
- Step 6: Risk Scoring and Timeline Analysis
- Step 7: Response Playbooks
- Step 8: Agentic Enterprise Coverage

These steps extend insider threat detection and response to contractors and AI-driven systems, enabling consistent monitoring and control for all identity types.

Conclusion

A behavior-driven insider threat program integrates people, process, and technology to detect and respond to risk more effectively. This approach shifts organizations from reactive alert handling to proactive risk management. By focusing on behavior, context, and identity, teams can detect issues earlier, prioritize high-impact activity, and reduce overall exposure.

With New-Scale Fusion, New-Scale SIEM, and New-Scale Analytics, along with built-in Exabeam Nova capabilities, organizations gain visibility into both human and non-human activity. Behavioral analytics, dynamic risk scoring, and automated response workflows help identify and contain insider threats before they escalate.

Over time, this model leads to:

- Faster detection and response
- Reduced impact from insider incidents
- Stronger governance across identities and systems
- Improved visibility into AI-driven and automated processes

A structured, behavior-driven program enables organizations to manage insider risk as an ongoing capability rather than a reactive function.

About Exabeam

Exabeam is the leader in behavior intelligence for the agentic enterprise. As organizations deploy digital workers and confront machine-speed adversaries, Exabeam delivers flexible, industry-proven solutions for insider threat coverage of humans and agents and faster, more accurate threat detection, investigation, and response (TDIR). Learn more at www.exabeam.com.

exabeam™

Learn more at
www.exabeam.com →