Security Operations Insider Investigation Playbook

How to Detect, Investigate, and Prioritize Insider Risk Hidden Inside Legitimate Activity

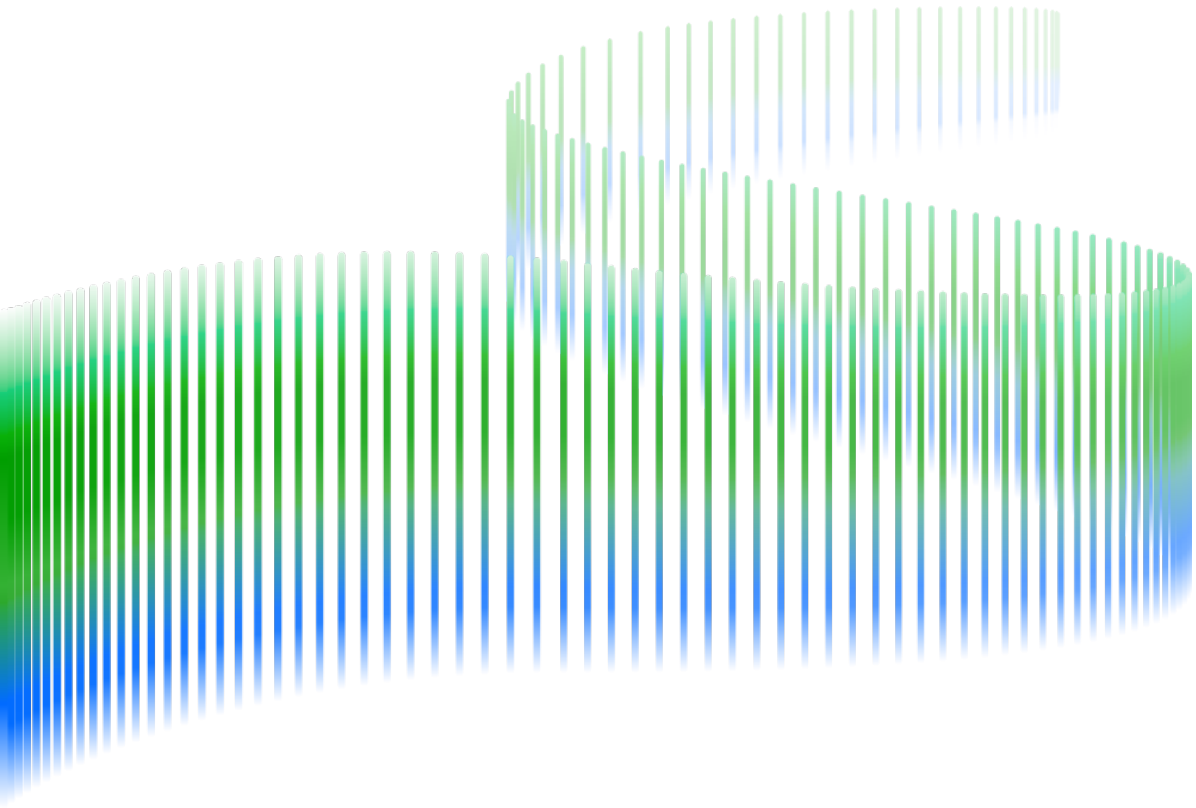

Table of Contents

Introduction

Insider threats rarely begin with an obvious violation. They develop through legitimate access, approved tools, and routine workflows. Viewed in isolation, the activity may look normal. The risk becomes visible when those actions are connected over time and reveal a meaningful change in behavior.

Modern insider risk extends beyond malicious employees to include compromised users, negligent insiders, privileged accounts, contractors, service accounts, and AI agents operating with trusted access. These identities don't need to bypass controls because they already have access.

Security operations teams feel this gap every day. Traditional detections operate within short, stateless windows. They generate alerts on individual events but do not retain behavioral history or connect activity across time. As a result, gradual insider risk is never assembled into a complete investigative story. Analysts are left reviewing isolated signals without the context needed to determine intent, impact, or progression.

This playbook provides a practitioner-focused investigation model that treats behavior as the unit of analysis. The goal is to help you detect progression earlier, prioritize real risk, and reach defensible decisions faster, without increasing manual effort.

The guidance applies to human users, privileged accounts, service accounts, AI agents, and autonomous workflows. While scale and execution differ, the investigative challenge remains the same: Trusted identities can introduce risk without triggering clear policy violations. Only behavioral context reveals when risk is developing.

By the end of this playbook, you will be able to:

- Recognize insider risk patterns that do not surface in isolated alerts.
- Build behavioral context that connects activity into meaningful sequences.
- Prioritize which cases require action and which need monitoring.
- Reduce investigation time without lowering analytical rigor.

This is an operator's playbook designed for real investigations and real decisions.

Why Insider Investigations Break Traditional Security

Most security operations workflows were designed to detect external attackers. Insider threats expose the limitations of that model.

Traditional detections focus on discrete events: an unusual login, a privilege change, a file download. Each event is evaluated independently and within a short time window. When activity is clearly malicious, this approach works.

Insider activity often appears legitimate at the event level, making it difficult to evaluate using traditional IF/THEN static alerts. This creates a fundamental mismatch between how systems detect activity and how insider risk develops.

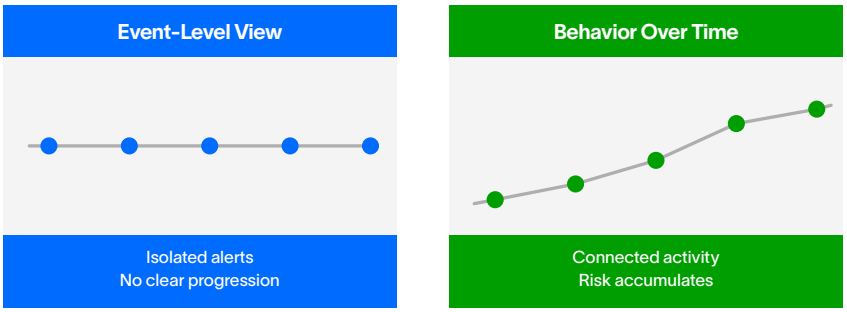

Figure 1.

Events that fall within baseline guardrails often appear normal in isolation. Behavioral change within those guardrails can go unnoticed by static rules, even as risk gradually develops.

Where Traditional Workflows Break Down

Security teams consistently encounter the same challenges:

- Alerts are generated on individual events but lack the context to assess risk.
- Analysts can see activity but can't easily understand user or entity behavior.
- Identity, endpoint, cloud, and application activity are not connected into a single view.
- There is no persistent baseline for normal behavior.
- Investigations require manual correlation across tools and time periods.

The result:

- Alert volume increases while signal quality decreases.
- Investigations stall because context must be reconstructed manually.
- Analysts spend more time assembling data than making decisions.

More importantly, the detection model itself fails.

Traditional SIEM correlation, XDR detections, and identity tools operate within short time windows. They are not designed to track behavioral progression over days, weeks, or months. When risk develops through normal-looking actions, these systems never connect the signals into a coherent narrative.

The Root Issue

Insider investigations fail not because analysts miss alerts, but because the system cannot assemble the story. Behavioral context is fragmented, historical activity is not retained in a meaningful way, and relationships across actions are not preserved.

In most security operations environments, analysts work from the highest-risk cases down. Individual alerts rarely score high enough to surface on their own. When multiple related alerts are grouped and correlated into a sequence, the combined behavior reaches a threshold that requires investigation. This is why evaluating patterns rather than events reveals the true risk story.

Without that context, there is no way to distinguish:

- Normal variation from meaningful deviation
- Isolated anomalies from sustained behavioral change
- Benign activity from emerging risk

How Exabeam Changes the Model

Exabeam applies behavioral analytics and long-term event correlation to connect activity across identities, systems, and time. Instead of working from isolated alerts, analysts start with:

- Behavior that is already correlated into investigation timelines
- Risk that is continuously scored and accumulated
- Context that persists for the duration of the investigation

This shifts the investigation model from:

event review → manual correlation → delayed decision

to:

behavior sequence → contextual analysis → immediate decision

Operator principle: Event-driven detection fails for insider risk because the signal is behavioral progression, not isolated activity.

What Makes Insider Risk Different From External Threats

External attacker investigations follow a familiar pattern: Timelines are short, indicators are clear, and policy violations are obvious.

Insider investigations differ in one critical way: Activity appears legitimate.

Access paths are legitimate, actions may remain within policy, and changes occur gradually. The signal is not a single violation but how behavior evolves.

This requires a different mental model:

- Risk accumulation instead of point-in-time detection
- Behavior sequences instead of isolated alerts

For insiders, the investigative unit is not the event. It is the evolving behavior of a user or agent.

Operator principle: Insider risk must be evaluated as evolving behavior over time, not isolated events.

What to Look For: Early Signals of Insider Risk

Insider risk rarely appears as a single, obvious alert. More often, it starts as a set of small changes that only become meaningful when viewed together.

The following signals indicate behavior that may warrant investigation, even when individual events appear legitimate:

- Users, service accounts, and agents with access to high-value data or privileged systems showing changes in behavior over time
- Gradual expansion of access across systems, datasets, or administrative functions
- Repeated activity associated with contractors, third parties, or distributed teams that deviates from baseline patterns
- Behavioral changes linked to AI usage, shadow AI tools, or autonomous agent activity
- Multiple low-signal alerts that do not escalate individually but appear related
- Activity that remains within policy but deviates from historical norms
- Environments where alerts lack sufficient context to determine risk

At the event level, these signals may seem routine. The risk is revealed when they are connected and evaluated as a behavioral sequence.

Operator principle:

Early insider risk is identified through patterns of change, not isolated anomalies.

Common Insider Risk Signal Categories

After identifying early indicators of potential risk, these signals can be grouped into common behavioral categories that help guide deeper investigation.

These categories provide a structured way to evaluate how behavior is changing and where risk may be accumulating:

- Authentication and access behavior changes such as new access times, locations, or methods that persist.
- Privileged activity expansion that gradually exceeds what a role previously required.
- Lateral movement using approved tools that mirrors normal administrative behavior.
- Data access anomalies that spread across systems or datasets over time.
- Service accounts and AI agents operating outside their expected scope or cadence.

At the event level, each action may seem normal. Risk only becomes visible when those actions are evaluated in context.

The Insider Investigation Lifecycle

Traditional investigation models assume analysts build context manually. Insider risk breaks that model. Signals are subtle, distributed, and often indistinguishable from legitimate activity.

In Exabeam, investigation starts from a different premise:

- Context is already assembled
- Behavior is correlated across identities and time
- Risk is continuously scored
- Activity is presented as an investigation-ready sequence

This model reflects how investigation is operationalized in Exabeam, while the workflow principles apply across modern security operations environments.

This section shows how to move from signal to decision using a structured, repeatable workflow.

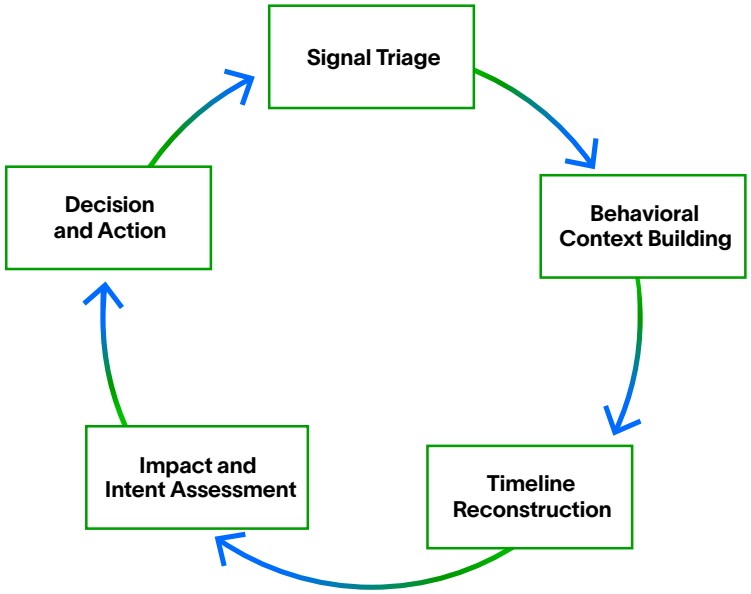

Figure 2.

A behavior-driven, continuous investigation approach that connects signals to build context, assess intent, and drive decisive action for human and non-human identities

Phase 1: Signal Triage

Start with risk, not alerts.

Where This Happens:

Threat Center + prioritization view

What the Analyst Does

Step 1: Prioritize entities by dynamic risk

- Sort alerts and cases by risk score
- Focus on sustained risk increase, not isolated spikes

Step 2: Evaluate detection grouping

- Review how related detections are grouped into cases
- Identify:
  - Developing patterns
  - Isolated anomalies

Step 3: Review enriched context across identities

- Validate behavior against historical patterns
- Identify deviations across users, service accounts, and agents

How to Decide

Escalate when:

- Risk shows sustained increase
- Multiple signal categories are present
- Activity deviates from baseline

Deprioritize when:

- Activity is isolated
- No pattern emerges

Operator principle:

Triage decisions are based on behavioral reinforcement, not alert severity.

Phase 2: Behavioral Context Building

Define normal.

Where This Happens:

UEBA entity profiles + behavioral baselines + Agent Behavior Analytics (ABA)

What the Analyst Does

Step 1: Establish baseline behavior

- Review historical activity:
  - Login patterns
  - System access
  - Activity cadence

Step 2: Compare against expected usage patterns

- Evaluate whether behavior aligns with historical activity
- Identify unexpected changes

Step 3: Identify persistent deviation

- Look for:
  - Expanding access scope
  - Increasing system coverage
  - Changes in activity patterns

How to Decide

Elevate concern when:

- Deviations persist
- Activity expands in scope
- Behavior diverges from historical patterns

Lower concern when:

- Activity aligns with operational needs
- Behavior stabilizes without continued expansion

Operator principle:

Risk is revealed through direction of change, not magnitude.

Phase 3: Automated Timelines

Analyze behavior using always-on timelines.

Where This Happens:

Investigation Timelines (Search) and Threat Timelines (Threat Center)

Timelines are always on and continuously updated, including late-arriving events. Analysts do not need to reconstruct sequences.

What the Analyst Does

Step 1: Open the unified timeline

View correlated activity across identities and systems

Step 2: Review behavior sequences surfaced by the platform

- Analyze progression across:
  - Access
  - Privilege usage
  - Data interaction

Step 3: Review escalation points identified by risk scoring

- Observe transitions such as:
  - New system access
  - Privilege increases
  - Expanding activity scope

Exabeam Nova and dynamic risk scoring surface these changes. The analyst interprets and decides.

Step 4: Analyze extended activity in a compressed view

- Identify patterns across days or weeks

Step 5: Pivot without losing context

- Drill into activity while preserving sequence visibility

How to Decide

Escalate when:

- Behavior forms a clear progression
- Multiple escalation points are present
- Activity transitions between stages

Deprioritize when:

- Activity remains bounded
- Sequence reflects normal workflows

Operator principle:

The timeline already contains the investigation. Analyst effort shifts to interpretation and decision.

Analyst Investigation Prompts

Once behavior has been automatically organized into a timeline view, the next step is interpretation. The following questions help determine whether activity represents meaningful risk.

When evaluating a user, entity, or agent, ask:

- → What has changed compared to this identity's historical baseline?
- → Does this behavior align with peer activity or does it deviate?
- → Is this activity isolated or does it form a sequence over time?
- → Is access expanding in scope, frequency, or sensitivity?
- → What systems or data are involved and how critical are they?
- → Would this activity be visible without behavioral correlation?
- → Does the sequence show progression from exploration to expansion or targeting?
- → Is this behavior consistent with known insider patterns such as compromised credentials, privilege escalation, or data exfiltration?
- → Does the behavior involve privileged users, contractors, service accounts, or AI agents?
- → If this pattern continues, what is the likely impact?

The quality of an investigation depends on the quality of the questions used to interpret behavior.

Phase 4: Impact and Intent Assessment

Translate behavior into risk.

Where This Happens:

Risk scoring views + timeline context

What the Analyst Does

Step 1: Assess accumulated risk

- Review risk score trends
- Focus on sustained elevation

Step 2: Map behavior to impact

- Identify:
  - Systems accessed
  - Data exposure risk
  - Privilege level involved

Step 3: Classify behavior pattern.

- Determine whether activity reflects:
  - Exploratory
  - Opportunistic
  - Targeted

Step 4: Validate against known patterns

- Compare to insider scenarios such as:
  - Compromised credentials
  - Privilege escalation
  - Data exfiltration

How to Decide

High-risk when:

- Behavior shows progression
- Sensitive assets are involved
- Activity persists or accelerates

Lower-risk when:

- No progression exists
- Activity remains limited

Operator principle:

Risk is defined by trajectory, exposure, and persistence.

Phase 5: Decision and Action

Execute quickly and decisively.

Where This Happens:

Response actions + automation workflows

What the Analyst Does

Step 1: Select response path

- Contain
- Monitor
- Dismiss

Step 2: Trigger automated response actions

- Notify stakeholders
- Initiate containment
- Document actions

Step 3: Preserve investigation artifacts

- Maintain timeline as evidence and audit record

Step 4: Feed continuous improvement

- Update detection logic
- Refine risk scoring
- Adjust monitoring priorities

How to Decide

Act immediately when:

- Risk is confirmed and progressing
- Sensitive assets are involved

Monitor when:

- Behavior is suspicious but inconclusive

Dismiss when:

- No pattern is established

Operator principle:

The goal is fast, defensible, repeatable decisions.

Key Operational Takeaways

- Investigation starts with risk and context, not alerts
- Behavior must be evaluated as a sequence
- The same model applies to users, service accounts, and AI agents
- Automated timelines eliminate manual reconstruction
- Analysts spend less time assembling data and more time making decisioins

End-to-End Example:
Investigating Compromised Credentials

Signal:
Moderate risk user with increasing login anomalies

Steps:

1. **Triage:** Risk score trending upward across multiple days
2. **Context:** New geolocation + increased system access vs. baseline
3. **Timeline:** Sequence shows login à privilege use à lateral movement
4. **Assessment:** Pattern aligns with compromised credential behavior
5. **Action:** Trigger containment playbook and escalate to IR

Outcome:
Detection occurred before data exfiltration stage

Investigating Low-and-Slow Insider Scenarios

Low-and-slow insider activity is the most difficult to detect because it avoids triggering traditional alert thresholds. Behavior is distributed across time and appears legitimate at each step. This applies equally to human users, service accounts, and AI agents operating with persistent access.

These scenarios are missed because no single action is significant enough to escalate. Risk only becomes visible when behavior is evaluated as a sequence.

Common low-and-slow patterns:

- Compromised credentials used conservatively over extended periods
- Gradual privilege expansion following role or access changes
- Incremental data access that spreads across systems or datasets
- Service accounts expanding execution scope beyond intended usage
- AI agents performing repeated actions outside expected operating patterns

Each action appears normal, but the pattern is not.

How to Investigate Low-and-Slow Activity

Step 1: Start with sustained risk indicators

- Identify identities with gradually increasing risk scores
- Focus on repeated or accumulating signals rather than isolated spikes

Step 2: Validate baseline deviation

- Compare current activity against historical behavior
- Look for changes in:
  - Access scope
  - Frequency
  - System coverage
- Apply the same logic across users, service accounts, and agents

Step 3: Review timelines

- Use automated timelines to analyze behavior sequences
- Do not manually reconstruct activity
- Evaluate how behavior progresses across sessions and systems

Step 4: Identify progression patterns

- Focus on direction of change, not individual events:
  - Expansion into new systems or datasets
  - Increased privilege usage
  - Repeated access to similar resources
  - Transition from exploration to broader activity
- For AI agents, progression often appears as:
  - Expanding task execution scope
  - Increased frequency of actions
  - Access beyond intended workflows

Step 5: Assess impact trajectory

- Determine whether behavior is moving toward meaningful risk:
  - Are sensitive systems or datasets involved?
  - Is activity expanding in scope or frequency?
  - Does the sequence align with known insider patterns?

Step 6: Decide and Act

- **Escalate** when progression and impact are both present
- **Monitor** when behavior shows deviation without clear impact
- **Dismiss** when no sustained pattern exists

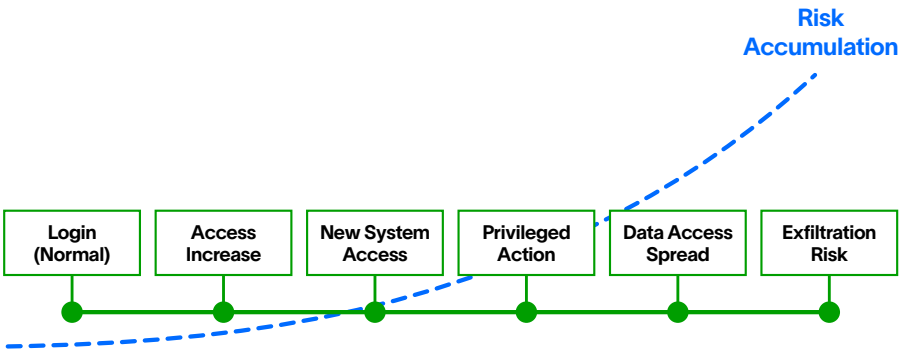

Figure 3.

In low-and-slow insider scenarios, individual actions appear normal at each step. Only when activity is viewed as a sequence do behavioral progression and accumulated risk become visible.

Why Low-and-Slow Activity Is Missed

Traditional detection models evaluate events in isolation and rely on thresholds. Low-and-slow activity remains below those thresholds and blends into expected behavior.

Behavioral baselines and automated timelines expose these patterns by preserving context, correlating activity, and presenting progression as a sequence instead of disconnected events.

Human Versus AI Agent Insider Risk

Agentic environments introduce new scale and speed challenges.

AI agents generate continuous activity and lack human intent signals. This increases the importance of consistent behavioral baselines.

Despite these differences, the investigation model doesn't change. The same behavioral lens applies to both human and non-human identities.

Reducing Investigation Time Without Losing Rigor

Investigation speed improves when context is already assembled. Exabeam reduces manual effort by correlating activity and maintaining behavioral history, eliminating the need to reconstruct timelines across tools.

This leads to:

- Faster investigations
- Fewer stalled cases
- Better prioritization of real risk

Analyst Shortcuts:
Reducing Investigation Time

- → Start with risk score, not alerts
- → Trust timeline correlation; avoid manual pivoting
- → Look for direction of change, not single anomalies
- → Prioritize multi-signal patterns over isolated behavior
- → Use automation early, not just at escalation

What Good Looks Like: Insider Investigation Maturity

Effective insider investigations are defined by how quickly and confidently analysts can move from behavior to decision, not how many alerts are reviewed.

Investigation maturity progresses through stages:

1. Alert-driven investigation

- Analysts respond to individual alerts.
- Context is fragmented across tools.
- Investigations require extensive manual effort.
- Low-signal activity is often missed or ignored.

2. Context-aware triage

- Alerts are enriched with identity and system context.
- Analysts begin to prioritize based on partial visibility.
- Some behaviors are correlated, but timelines are incomplete.
- Investigation still requires manual assembly.

3. Behavior-led investigation

- Activity is evaluated as sequences, not isolated events.
- Behavioral baselines provide context for deviation.
- Timelines begin to connect activity across systems and time.
- Analysts spend less time correlating data and more time analyzing behavior.

4. Risk-based prioritization

- Behavior is continuously evaluated and scored.
- Analysts prioritize identities with sustained deviation.
- Investigations start with context already assembled.
- Decision making becomes faster, more consistent, and more defensible.

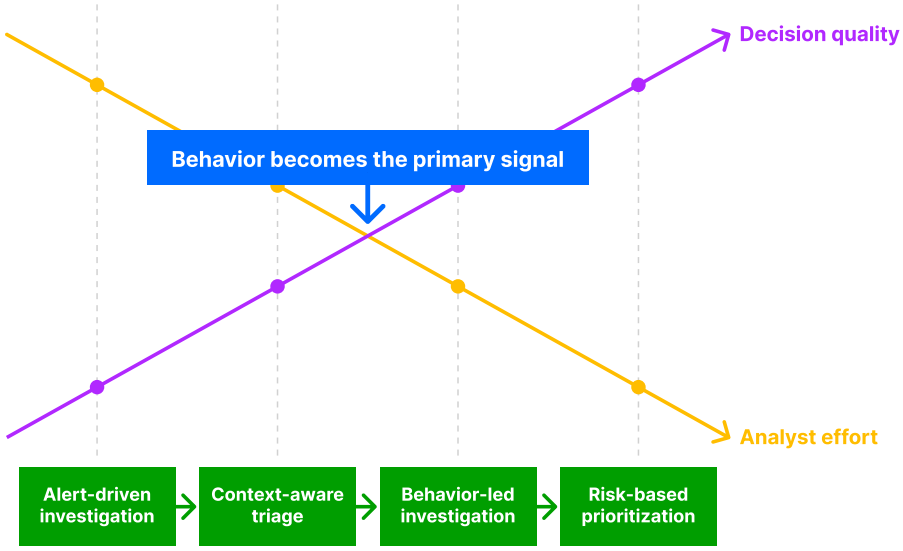

Figure 4.

As investigation maturity increases, behavior becomes the primary signal. Analyst effort declines while decision quality improves.

What This Looks Like in Practice

At higher maturity, security operations teams:

- Start investigations with risk-ranked identities instead of alerts.
- Use behavioral baselines to immediately identify deviation.
- Analyze activity through unified investigation timelines.
- Detect low-and-slow insider risk before impact occurs.
- Investigate both human and non-human identities using the same model.
- Reduce manual effort by eliminating cross-tool pivoting.
- Make decisions based on accumulated context, not isolated signals.

Exabeam enables this maturity by applying dynamic risk scoring, behavioral baselines, and automated timelines to continuously evaluate activity across users, service accounts, and AI agents.

The result:

- Higher decision quality
- Faster investigations
- Reduced analyst workload
- Defensible, repeatable investigation processes

Key Takeaways for Security Operations Teams

- Insider risk cannot be investigated like external attacks.
- Behavior, not alerts, is the investigative unit.
- Risk emerges through patterns rather than single events.
- Unified visibility across humans and agents is required.
- The right model reduces risk and analyst workload at the same time.

Conclusion

Insider risk, introduced by human and non-human entities, is difficult to detect because it develops within legitimate activity.

The challenge for security operations is not a lack of alerts. It is understanding when behavior changes and what that change represents.

By treating behavior as the unit of investigation, analysts move beyond fragmented signals and toward actionable insight. Context persists instead of being reconstructed. Risk becomes measurable instead of assumed.

Exabeam enables this shift by correlating activity into unified timelines, maintaining behavioral baselines, and scoring risk as it develops. Investigations begin with context already assembled, allowing analysts to focus on decisions instead of data collection.

When you can see how risk evolves, you act earlier, reduce investigation effort, and make better decisions.

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 →