

Detection Strategies for the Agentic Enterprise

When to Use Rules, Behavioral
Analytics, and Hybrid Models




Table of Contents

03	Executive Summary
04	Why the Agentic Enterprise Changes Security Operations
06	Understanding Detection Approaches
08	Detection Evaluation Scenarios for the Agentic Enterprise
10	Detection Strategy Framework
12	Building a Detection Strategy for the Agentic Enterprise
13	What an Effective Detection Strategy Requires

Executive Summary

Agentic systems are changing the operating environment for security teams. AI-enabled applications, automation platforms, copilots, and AI agents now participate in more of the workflows that keep a business running, and that shift is pushing more activity through legitimate credentials and approved systems than most security programs were built to evaluate. The job itself hasn't changed: Distinguish known conditions, expected activity, and meaningful deviations while maintaining operational efficiency. What has changed is the volume and speed at which that job must happen, and how much of it now runs through authorized access.

Rules remain valid. What you need to decide is which detection approach best fits each problem you're solving. Organizations need a detection strategy that combines deterministic rules, behavioral analytics, risk prioritization, and investigation context to address both known and emerging risks in an increasingly automated environment.

Key message: The agentic enterprise doesn't make rules obsolete. It increases the need for complementary detection methods that can evaluate activity, context, and risk at scale.

Why the Agentic Enterprise Changes Security Operations

More Activity Happens Through Legitimate Access

Historically, many detection programs focused on malicious indicators, unauthorized access, or clearly suspicious events. Rules built to catch a known-bad IP address, a blocked file hash, or a login from an impossible location did most of the work, because most of the risk a security team needed to catch showed up as something visibly wrong.

Today, that's not where most activity lives. Users work in more systems, and AI agents, service accounts, automated workflows, and AI-enabled applications now perform more of the actions that used to require a person at a keyboard. Increasingly, that activity runs on valid credentials, through approved systems, doing exactly what those systems were built to do.

That creates a different evaluation problem. For any given event, a security team now has to determine whether the activity is authorized, whether it's expected, whether it's risky despite being authorized and expected, and whether it warrants additional investigation.

None of those four questions gets answered by confirming the credential is valid, because a valid credential is no longer evidence of anything on its own.

Key message: Legitimate access does not automatically mean low risk.

The Operational Problem Is Scale

As organizations adopt AI-enabled workflows, the volume problem gets harder before it gets easier. Activity volume rises as agents, automations, and copilots generate log and event data around the clock, at a pace no human-driven workflow ever produced. Investigation demand climbs with it, because more activity worth a second look means more open cases at once. And alert fatigue, already a persistent problem on most security operations teams, doesn't ease up just because the new alerts come from agentic activity instead of endpoints. It gets worse because analysts now have to evaluate a wider range of account types and activity patterns in the same number of hours per shift.

Prioritization determines whether that added volume turns into useful detections or just more noise. Security teams need to know which activity is most critical, which detections warrant a human's time, and which events represent conditions the team already understands well enough to deprioritize.

Key message: Agentic environments increase the volume of activity analysts must evaluate, making prioritization and context increasingly important.

Human Users and Non-Human Accounts Are Critical

Security teams now monitor a wider population than the word “user” traditionally covers: human users, yes, but also service accounts, machine accounts, automated processes, and a growing set of AI agents acting with their own credentials and their own scope of access.

Evaluating that population well means answering the same four questions for every account type a program touches. Which users, agents, and accounts are represented in the environment? What telemetry exists for each account type? How is activity attributed back to the account or agent that generated it? And what behavior can realistically be baselined, given the data available?

A detection strategy that only answers these questions for human users leaves the fastest-growing part of that population effectively unmonitored.

Key message: Account context is increasingly important as organizations rely on both human and non-human activity.

Understanding Detection Approaches

Before building a strategy, it helps to be precise about what each detection approach does well. Treating rules and behavioral analytics as competing philosophies misses the point. They answer different questions, and a mature program needs both.




Figure 1.

The hybrid detection model. Deterministic rules catch known threats and policy or control violations. Behavioral analytics surfaces meaningful deviations from established baselines. Risk and investigation context determine what an analyst works first. Together, they form a hybrid detection strategy that fits each method to the problem it handles best.

Deterministic Rules

Correlation rules, threshold rules, signature detections, detections based on indicators of compromise (IOCs), and alerts on specific policy or control violations share one design: each detects a known condition with explainable logic and a predictable outcome.

Deterministic rules are the backbone of most detection programs, and for good reason: if a defined condition occurs, the rule fires, and an analyst or an auditor can trace exactly why. Exabeam, like most mature security platforms, uses rule-based detection extensively, for the same reason every serious program does. Rules are still the most direct way to catch a known threat or flag a specific policy or control violation. What rules don't do well is catch the condition nobody has written a rule for yet, which is the gap agentic activity tends to open.

Key message: Deterministic rules remain foundational for known threat detection and for alerting on specific policy or control violations.

Behavioral Analytics

Behavioral analytics starts with a different question than rules do. Instead of asking whether an event matches a known bad condition, it asks whether current activity from a user, agent, or account is consistent with how it has behaved before, and it does that through baselining, anomaly detection, rare-event identification, deviation analysis, and risk scoring that feeds prioritization.

That approach answers four questions rules can't: Is activity unusual? Is behavior changing? Is activity consistent with prior patterns? Does investigation priority increase?

Key message: Behavioral analytics helps evaluate whether activity differs meaningfully from expected patterns.

Hybrid Detection

A hybrid approach combines deterministic rules, behavioral analytics, risk scoring, and investigation context, and each piece has a specific job. Rules preserve coverage for the known conditions a program has already invested in detecting. Behavioral analytics adds context around activity rules can't classify on their own. Risk scoring improves how that combined output gets prioritized, and investigation context carries the analyst work that follows. This is the model behind New-Scale Fusion: rule-based detection for known conditions and policy or control violations, Agent Behavior Analytics (ABA) and user and entity behavior analytics (UEBA) for baselining and deviation spanning human users, agents, and service accounts, and risk-based prioritization that brings both into a single queue an analyst can work.

Key message: Most mature security programs rely on multiple detection approaches rather than a single method.

Detection Evaluation Scenarios for the Agentic Enterprise

The distinctions above become concrete once you apply them to specific situations. The four scenarios below cover the conditions a detection strategy for the agentic enterprise has to handle, and each one points to a different best fit.

Scenario 1: Known Threat Detection

Example: A known malicious indicator appears in endpoint or network telemetry.

Best fit: Deterministic rule-based detection.

Evaluation question: Can the platform reliably identify known conditions and present investigation context alongside the alert?

This is the scenario deterministic rules were built for, and it's still common. A malicious hash, a known-bad IP address, a signature match: these conditions are defined well enough in advance that a rule can catch them with high confidence and low ambiguity. Most platforms can fire the alert. The evaluation question is whether the platform gives the analyst enough context alongside that alert to act on it without pulling data from three other tools first.

Scenario 2: Unusual Use of Legitimate Access

Example: A user, agent, or service account begins operating outside established patterns.

Evaluation question: Can the solution identify meaningful deviations when sufficient baseline data exists?

This scenario is where rules alone run out of runway. Nobody writes a rule for "this service account normally touches four systems and now touches eleven," because that condition doesn't exist until it happens. Detecting it takes a baseline of what normal looks like for that specific account, and the ability to recognize when current activity has moved meaningfully away from it.

Key message: Authorized activity may still warrant investigation when behavior changes significantly.

Scenario 3: AI-Related Activity Monitoring

Example: An organization introduces AI-enabled applications, copilots, or automated workflows.

Evaluation questions: What telemetry is available? How is activity represented? Which detections exist today? How are human users and non-human accounts modeled?

This is the scenario security teams are least prepared to evaluate accurately, mostly because it's new enough that assumptions outrun facts. Not every platform ingests telemetry from every AI-enabled application, and not every platform models a copilot or an automated workflow as its own actor with its own baseline. Before assuming a detection program has agent visibility, confirm what telemetry exists, how that activity is represented once ingested, and how the platform distinguishes a human user from a non-human account.

Key message: Buyers should evaluate actual telemetry and supported detections rather than assume universal agent visibility.

Scenario 4: Detection Prioritization at Scale

Example: Analysts face thousands of alerts generated from normal business activity.

Evaluation questions: How is risk prioritized? What context is available? How are detections grouped? Which investigations should happen first?

At scale, the harder problem is whether an analyst can tell which of a thousand alerts to work first. Grouping related detections into a single case, attaching risk scores that reflect actual severity, and surfacing supporting context up front all determine whether a detection program is usable in practice or just generates more noise for a team that's already stretched thin.

Key message: Detection quality increasingly depends on prioritization, not simply alert volume.

Detection Strategy Framework

Evaluation Dimensions

The table below maps each detection approach against seven evaluation dimensions specific to agentic environments. Use it as a working framework when evaluating your own detection stack, or a vendor’s: for each row, know which approach carries the weight, and why, for a program built to handle machine-speed, non-human, and AI-related activity alongside everything else a SOC already monitors.

Dimension	Deterministic Rules	Behavioral Analytics	Relevance
Known conditions	Strong fit	Adds context	Known threats remain important.
Deviations from expected activity	Limited without predefined logic	Stronger fit	Emerging risk may not match predefined conditions.
Explainability	High	Requires clear supporting evidence	Analysts must understand why detections occurred.
Investigation prioritization	Limited alone	Strong contributor	Security teams must focus attention efficiently.
Compliance and control evidence	Can alert on specific policy or control violations	Can add context around activity relevant to controls	Compliance reporting typically requires evidence that controls are in place, not detection logic alone.
Non-human account activity	Depends on telemetry	Depends on telemetry and baselines	Account context grows more important in agentic environments.
AI-related activity	Depends on available telemetry	Depends on available telemetry	Visibility varies by platform and implementation.

Table 1.

Evaluation dimensions for detection strategy in the agentic enterprise. Use this table to map deterministic rules and behavioral analytics against seven dimensions that determine whether a detection approach fits a given problem, and to see where compliance reporting sits alongside, rather than inside, detection logic.

Key message: Match each detection method to the security problem it handles best and treat compliance reporting and control evidence as related but distinct requirements.

Detection and compliance are easy to conflate. Compliance programs typically need reporting and evidence that specific controls are in place and operating, not just a detection rule that happens to touch the same control. A rule that alerts on a policy violation is useful, but it isn't a substitute for the audit trail a compliance team needs. Build both into the strategy deliberately, rather than assuming one produces the other.

Building a Detection Strategy for the Agentic Enterprise

Turning this framework into a working program takes a sequence, not a single decision. Work through the seven steps below in order.




Figure 2.

A seven-step roadmap for building a detection strategy for the agentic enterprise. The first six steps establish what you already detect, where risk sits inside legitimate access, and how analysts work today. Step 7 turns those findings into a hybrid strategy your SOC can run.

- 1. Inventory current deterministic and behavioral detections.**
Know what you already have before deciding what to add.
- 2. Map detections to human users and non-human account types.**
If your inventory was built human-first, make the non-human gap visible.
- 3. Identify areas where activity increasingly occurs through legitimate access.**
This is where risk hides in plain sight, inside access nobody would flag as wrong.
- 4. Validate available AI-related telemetry and monitoring coverage.**
Confirm what's being ingested before assuming visibility exists.
- 5. Assess baselining, anomaly detection, and prioritization capabilities.**
This is where behavioral analytics either earns its place in the stack or doesn't.
- 6. Evaluate analyst workflows and investigation experience.**
A detection is only as good as the investigation it enables.
- 7. Develop a hybrid strategy combining known-condition detection with contextual behavioral analysis.**
This is the synthesis step, where the first six steps come together into a program a SOC can run.

What an Effective Detection Strategy Requires

The agentic enterprise increases the amount of activity security teams must understand, prioritize, and investigate. Rules remain essential for detecting known threats, explicit attack patterns, and specific policy or control violations, and nothing in this framework argues otherwise. Behavioral analytics helps organizations evaluate deviations, prioritize activity, and add context that rules alone can't provide, especially for a growing population of agents and service accounts. Compliance programs still need their own reporting and evidence that controls are in place, addressed alongside detection strategy rather than assumed to follow from it automatically.

None of these three pieces replaces the other two. A detection strategy built for the agentic enterprise needs all three, applied deliberately, not assumed by default.

How Exabeam Puts This Framework Into Practice

Exabeam applies this same hybrid model in New-Scale Fusion: rule-based detection for known threats and policy or control violations, Agent Behavior Analytics (ABA) and user and entity behavior analytics (UEBA) for baselining human users, agents, and service accounts, and risk-based prioritization that brings rule-based and behavioral detections into one queue an analyst can work efficiently. Exabeam Nova AI adds investigation support on top of that detection layer, drafting timelines and surfacing context so analysts spend their time on the judgment calls only they can make.

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.



Learn more at
www.exabeam.com →

Without limitation, the Exabeam and LogRhythm names and logos, related product, service, and feature names, and related slogans are service marks, trademarks, or registered marks of Exabeam (or its affiliates) in the United States and/or other countries. All other brand names, product names, or trademarks belong to their respective owners.
© 2026 Exabeam, LLC. All rights reserved.