

Securing the Agentic Enterprise




Table of Contents

03	Abstract
04	Introduction
05	Identity Has Expanded Beyond Humans
06	Risk Hides Inside Legitimate Activity
07	Manual Security Operations Cannot Keep Pace
08	Agentic Enterprise Security
09	A Lifecycle Approach: Verify, Observe, Analyze, Improve
12	Conclusion

Abstract

The enterprise has become agentic. AI agents have moved out of experimentation and into production, where they reason, call tools, access sensitive data, run workflows, and act on valid credentials of their own. Enterprise identity no longer means people alone. It now includes service accounts, machines, and AI agents operating inside business processes right alongside employees. That shift changes the security model, because risk can hide inside activity that looks completely legitimate to traditional controls.

Security leaders must balance AI-driven productivity with the need to manage operational, financial, regulatory, and reputational risk. That starts with knowing whether an agent's behavior is expected or a warning sign. Static rules, prompt guardrails, human-only insider threat programs, and event-based monitoring weren't built for workflows that adapt, chain actions, and move across systems on their own.

This shift requires a new security approach. Throughout this paper, we refer to that approach as Agentic Enterprise Security: one behavioral risk model covering human and non-human identities side by side. This framework helps security teams understand agent behavior in real workflows, uncover risk hidden within legitimate activity, and adapt security operations to keep pace.

Exabeam secures the agentic enterprise with behavioral intelligence across human and non-human identities. Behavioral analytics, risk scoring, investigation timelines, automation, and AI-assisted workflows work together to make agent activity visible, understandable, and actionable, before risk turns into business impact.

Introduction

AI agents are becoming digital insiders. They operate on legitimate credentials, reach into sensitive systems, call tools, and take action inside enterprise workflows, much like any employee with a badge. As agents join the workforce, security teams need a model that goes beyond traditional identity and access controls.

None of this makes AI agents inherently malicious. The risk is that autonomous systems are hard to govern, observe, and defend once they're embedded in a business process. Their activity blends into normal workflows, so risky or malicious behavior can hide in plain sight. Security teams need a fast, reliable way to tell expected behavior from a warning sign.

Traditional security was built around isolated events, static rules, and human-centric identity models. It struggles to tell whether a sequence of actions, spread across users, agents, tools, APIs, and systems, adds up to something normal or something dangerous. As agentic systems scale, organizations need a model that tracks behavior over time, not just the moment an event fires.

AI agents have stopped being tools people use and started being participants in the workflow itself. They act autonomously, run workflows, and hold legitimate credentials of their own. Enterprise identity now spans humans, service identities, and AI agents alike. Traditional tools tend to focus on events rather than behavior, which means they miss context and let risk hide inside activity that looks perfectly normal.

Securing the agentic enterprise requires a shift from event-based detection to behavioral intelligence, an approach that delivers visibility, context, and risk-based response across identities, systems, workflows, and autonomous activity. Exabeam helps organizations surface hidden risk in legitimate activity, uncover shadow and rogue agent behavior, and speed up security operations for a workforce that now includes both humans and AI agents.

This paper examines how the rise of AI agents is changing enterprise identity, why risk increasingly hides within legitimate activity, and why traditional security operations struggle to keep up. It also introduces Agentic Enterprise Security and outlines a lifecycle-based approach for securing AI agents through behavioral visibility, analysis, and continuous improvement.

AI adoption is accelerating through commercial, custom-built, and open-source technologies. At the same time, agent behavior is becoming harder to predict and govern using traditional security approaches. Organizations need more than guardrails and runtime monitoring. They need a repeatable operating model that helps them verify expected behavior, observe activity, analyze risk, and continuously improve controls throughout the agent lifecycle.

Identity Has Expanded Beyond Humans

Enterprise identity no longer means human identity alone. It now spans employees, contractors, service accounts, APIs, machines, autonomous systems, and AI agents, all operating under the same roof.

The growth of commercial, custom-built, and open-source AI technologies is also increasing the number of non-human identities organizations must govern. Business teams can create workflows and deploy agents outside traditional security oversight, increasing the risk of shadow agents, excessive permissions, and unmanaged behavior.

Traditional identity programs still play a role, but permissions and access controls can't answer the real question in an agentic enterprise: is this identity behaving the way it should? Legitimate credentials and authorized access don't guarantee safe behavior. Every identity needs to be evaluated in behavioral context, not just checked at the door.

That's especially true for non-human identities. Many organizations lose visibility into how service accounts and AI agents operate once access is granted. Behavior fills that gap, giving teams a way to tell expected activity apart from misuse, compromise, excessive privilege use, or risky automation.





Figure 1. **Enterprise identity has expanded.** Organizations must understand behavior across humans, machines, service accounts, and AI agents.

Risk Hides Inside Legitimate Activity


Modern threats increasingly show up as sequences of actions, not single isolated events. Agents, employees, contractors, service accounts, and machines all operate on authorized access. As activity volumes grow, the hard question stops being whether an action happened. It becomes whether that action made sense, given the identity, workflow, role, system, and business process behind it.

Rules catch events. Behavior reveals risk. Static detections and prompt guardrails can't predict every downstream action, adaptive workflow, tool invocation, or business scenario a system might run into. Behavioral analytics, long-term baselining, risk scoring, and contextual correlation give organizations a way to catch emerging risk hidden inside activity that otherwise looks legitimate.

WHAT RULES SEE



WHAT BEHAVIOR REVEALS



Cumulative risk 88 — prioritized for investigation

Figure 2. Risk increasingly hides inside legitimate activity. Behavioral context reveals what traditional controls miss.

Manual Security Operations Cannot Keep Pace

AI agents act at a speed and scale traditional security operations were never built for. Manual triage, investigation, and response can't keep up with machine-generated activity. Organizations need a security operations model that moves at the same velocity as the threats and the AI systems it's defending against.

Machine-scale risk demands machine-speed detection, investigation, prioritization, and response. Organizations need AI-assisted workflows, automation, behavioral analytics, and investigation tools that point analysts straight at the highest-risk activity, before it turns into business impact.




Figure 3. **Behavioral intelligence** connects activity over time to uncover hidden risk within legitimate behavior.

Agentic Enterprise Security


One Behavioral Risk Model for Humans and Agents

Agentic Enterprise Security protects environments where human and non-human identities work side by side, under one behavioral risk model. That unified view is necessary because humans, service identities, machines, and AI agents are now woven into the same workflows, access patterns, and business processes.

Exabeam defines Agentic Enterprise Security as the ability to detect and prioritize risk in behavior across human and non-human identities at scale. This approach helps organizations catch threats hiding in legitimate activity before they can negatively impact the business.

For security leaders, that means securing AI adoption without slowing innovation or piling up unmanaged operational risk. It means audit-ready visibility into AI-driven activity across human and non-human identities, and a clear line from agentic enterprise security investment to governance, compliance, board reporting, and operational resilience.

For technical teams, it means visibility into AI agent behavior, faster detection of anomalies across human, machine, and agent activity, risk prioritized with real behavioral context, and investigations accelerated by unified timelines, detection grouping, automation, and AI-assisted workflows.



A Lifecycle Approach: Verify, Observe, Analyze, Improve

Securing AI agents takes more than runtime monitoring. Agentic risk can emerge long before an agent reaches production. Security teams need to understand:

- **How an agent is configured**
- **Which tools it can call**
- **What data it can reach**
- **Whether its intended behavior already exceeds its authorized scope before it ever reaches production**

A strong operating model for Agentic Enterprise Security runs on four continuous stages: verify expected behavior, observe agent activity, analyze behavioral deviations, and improve security outcomes.




Figure 5. **Agentic Enterprise Security** requires continuous verification, observation, analysis, and improvement.

1. Verify Expected Behavior Before Production

Every AI agent needs a defined behavioral remit before it takes a single action. Organizations need to know what it's authorized to do, which tools it can use, what data it can touch, which systems it can reach, and just as important, which actions fall outside that scope entirely.

Pre-production verification cuts risk before an agent ever touches a live workflow. By checking whether an agent's configuration, prompts, memory, and integrations actually match its authorized role, security teams can catch scope creep before deployment, not after.

2. Observe Activity Across Agentic Environments

Once agents are live, security teams need a reliable record of what they're doing. Observability should capture prompts, tool calls, approvals, actions, outputs, outcomes, and the workflow context analysts need to reconstruct what happened, and why.

Standardized agent telemetry creates one consistent behavioral record across a mix of AI environments. Observra normalizes agent activity into a common information model, so prompts, tool calls, actions, approvals, and outcomes are captured the same way every time, ready for behavioral analysis.

Exabeam Agent Sensor extends that visibility into AI CLI activity, capturing session activity, prompts, tool calls, token usage, and cost data. It then normalizes all of it to the Exabeam Common Information Model for consistent detection and investigation.

3. Analyze Behavior for Risk, Drift, and Misuse

Telemetry alone isn't enough. Security teams still have to determine whether activity is normal, risky, drifting, or abusive, in the context of the user, agent, role, tool, system, workflow, and time it happened.

Exabeam applies behavioral analytics across human and non-human identities, including UEBA and Agent Behavior Analytics (ABA), to baseline activity and catch deviations as they happen. Behavioral risk scoring, dynamic prioritization, Threat Center, Threat Timelines, Investigation Timelines, detection grouping, fast search, and long-term visibility all work together to correlate activity across identities, systems, tools, and workflows.

That turns agent activity into investigation-ready context. Analysts move beyond isolated alerts to evaluate whether behavior points to misuse of legitimate access, excessive privilege use, suspicious deviation, agentic drift, or a longer chain of risk.

4. Improve Outcomes Continuously

Agentic Enterprise Security works best as an operating model for continuous improvement, not a one-time monitoring project. Security teams need to close coverage gaps, tune detections, align to security and compliance frameworks, demonstrate control, and sharpen readiness over time.

Outcomes Navigator connects security operations to measurable outcomes and governance priorities. Automation, AI-assisted investigation, Threat Center, and Exabeam Nova reduce manual work, sharpen prioritization, and accelerate detection, investigation, and response.

What This Means for Security Teams

Security leaders and security operations teams both need to reduce AI-driven risk without slowing innovation. Leaders need defensible visibility, reporting that supports executive and regulatory scrutiny, and a clear line from security investment to measurable outcomes. Operations teams need the context to catch misuse of legitimate access across humans and AI agents, cut alert noise, speed investigations, and correlate behavior across identities, systems, and workflows.

Behavioral intelligence gives both groups a shared operating model. It helps leaders strengthen compliance readiness, operational resilience, and board-level visibility into an expanding identity surface, while giving analysts unified timelines, risk scoring, detection grouping, automation, and AI-assisted workflows to investigate prioritized behavioral risk at machine speed.




Figure 6. Behavioral intelligence helps organizations govern AI-driven activity without slowing investigation.

Conclusion

The enterprise has changed. AI agents now operate inside business workflows with legitimate access, sensitive data, and the ability to act at machine speed. Traditional tools built around isolated events, static rules, or human-only identity models can't reliably tell whether that activity is expected or a threat in the making.

Securing this environment requires Agentic Enterprise Security: one behavioral risk model covering human and non-human identities alike. Organizations need to verify expected behavior before agents run, observe activity across environments, analyze behavioral deviations over time, and continuously improve security outcomes.

Exabeam secures the agentic enterprise with behavioral intelligence that detects and prioritizes risk across human and non-human identities at scale, giving leaders defensible visibility into AI-driven exposure and giving security teams the context to investigate and respond before risk becomes business impact.

Slowing AI adoption was never the goal. Making AI-driven activity visible, understandable, and actionable is. Organizations that build security operations around behavior will be better positioned to govern AI agents, reduce hidden risk, and defend an enterprise where the workforce now includes both humans and autonomous systems.

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.