

The Speed of Risk

Why Machine-Speed Threats Require Machine-Speed Security Operations

Executive Summary

The enterprise has entered a new operating era. AI agents now execute workflows, make decisions, access systems, and interact with sensitive data alongside employees. Automation connects applications throughout the business, while service accounts and APIs perform work continuously. Enterprise activity now comes from people, services, machines, and AI agents operating together inside business processes. The numbers back it up: [69% of enterprises already run AI agents in production](#).

That gap between how fast businesses develop and operate and how fast security teams can respond is widening. The challenge is no longer limited to preventing external and traditional insider threats. [Only 1 in 5 organizations has mature governance for autonomous agents](#), and most security teams are in danger of losing pace with the enterprise.

Traditional security was designed for environments dominated by human activity. Analysts often investigated alerts sequentially, and rules, signatures, and isolated events helped identify known threats. Those methods remain important, but they are incomplete when risk develops through valid access, authorized actions, and connected sequences of behavior.

Machine-speed risk requires a new operating model. Applying behavioral intelligence to their security data helps teams baseline behavior over time, prioritize meaningful risk, investigate activity in context, and improve outcomes. The agentic SOC applies AI-assisted investigation and automation while analysts retain judgment, accountability, and response authority.

The future of security operations is an agentic SOC designed to operate at machine speed while keeping people in control.

The Enterprise Has Accelerated

The Rise of the Agentic Enterprise

For decades, enterprise security focused primarily on people. Users authenticated, accessed systems, performed actions, and generated the activity that security teams monitored. Service accounts and automation existed, but people initiated much of the meaningful business activity.

That model is changing. Organizations increasingly rely on AI agents, autonomous workflows, machine-driven decisions, and connected automation platforms to see productivity increases and improved business intelligence. These systems can access applications, retrieve information, invoke tools, move data, and execute actions with legitimate credentials and approved permissions.

The enterprise now runs on human employees, AI agents, service accounts, APIs, and automated workflows. Together, these actors form the operational fabric of the agentic enterprise.

A New Operating Model

Automation can reduce repetitive work, accelerate decisions, increase productivity, and help organizations scale. It can also increase the speed at which risk develops and spreads.

When workflows execute automatically and decisions occur at machine speed, security teams have less time to determine what happened, why it happened, and whether they need to intervene. Business operations are accelerating. Security operations must evolve with them.

Business Speed and Security Speed

Activity	Operational Reality
Machine actions	Can execute in seconds and trigger connected workflows
Investigation	Often requires analysts to gather and correlate context
Remediation	May depend on validation, coordination, and human decision authority

Table 1. **Machine actions** can execute in seconds, but investigation and remediation still depend on manual analyst work, widening the gap between business speed and security speed.

Understanding Machine-Speed Risk

AI agents are not inherently malicious. The challenge is that they can operate continuously, at scale, and across multiple systems. Like people, they authenticate, access data, use approved credentials, execute workflows, and trigger downstream actions. Unlike people, they can perform these activities at machine speed.

This creates a new risk environment. Individual actions may appear legitimate, yet connected actions can create significant exposure before a security team has enough context to respond.

When Risk Moves Faster Than Defenders

Four Sources of Machine-Speed Risk

1. Shadow AI

Employees may adopt AI tools, copilots, browser extensions, and assistants outside approved governance processes. Security teams may lack visibility into what these systems access or how people use them, which can be a source of private data leaving the organization.

2. Vendor-Embedded AI

Enterprise platforms may introduce AI assistants and copilots directly into applications. These capabilities can inherit existing permissions and access business-critical data, and, if exploited by an attacker, could potentially result in a catastrophic breach.

3. Citizen-Built Automation

Low-code and no-code platforms let business users connect sensitive systems without traditional security review. A productivity workflow can become an unmanaged source of risk, especially if vulnerable to attack.

4. Custom Enterprise Agents

Purpose-built agents may access databases, APIs, SaaS platforms, and sensitive workflows. Broad authority and limited behavioral oversight can increase risks to business continuity and exposure to attack.

Although these technologies differ, they share common risks: unauthorized tool use, excessive privilege consumption, unusual workflow chains, sensitive data exposure, and abnormal behavioral sequences. That risk is already top of mind in the C-suite. [37% of CISOs now call securing AI agents their number one priority.](#)

Why Traditional Security Falls Short

As machine-driven activity grows, security operations teams face more alerts, more complex investigations, shorter response windows, and larger backlogs. Manual processes can turn the security operations center into a bottleneck.

The Human-Speed SOC Problem

Many security operations centers were built around four assumptions:

- Activity is primarily human driven.
- Incidents are discrete events.
- Investigations happen sequentially.
- Alerts represent meaningful indicators of risk.

These assumptions become less reliable as autonomous activity grows. Machine-driven systems generate large volumes of legitimate activity. Finding events is the easy part. Determining which activity represents meaningful risk is the real challenge.

Rules, signatures, and policy controls remain necessary for known threats and explicit violations. They are less effective on their own when misuse, compromise, misconfiguration, or behavioral drift occurs through legitimate access. Security teams need both static and dynamic detection, supported by context that connects activity over time.

The Cost of Delay

Risk can develop before an investigation concludes. By the time analysts complete triage, data may be exposed, privileges may be escalated, workflows may finish, and business processes may already be affected. That delay carries a real cost: [Organizations using AI and automation in security operations see \\$1.93 million lower average breach costs and resolve incidents 65 days faster than those without.](#)

The objective is to identify and prioritize risk before it becomes business impact.

Risk Is a Pattern, Not an Event

Traditional detection often evaluates individual activities such as a login, an API call, or an access request.

Agentic environments operate through sequences. Risk can emerge through gradual escalation, behavioral drift, excessive tool use, unusual workflow chains, and connected activity between people and agents.

One action may appear legitimate. A sequence of actions can tell a different story. The application of behavioral intelligence to security data helps teams:

- Establish behavioral baselines
- Detect meaningful deviations
- Apply dynamic risk scoring
- Maintain long-term behavioral context
- Connect from people, AI agents, and service accounts

Instead of asking only whether an action occurred, security teams can ask whether the action fits the broader context. That distinction helps teams surface risk hidden inside legitimate activity, and it pays off: [Behavioral analytics can reduce average incident cost by 68%.](#)

\$1.93M

lower average breach cost

65 days

faster breach resolution

Figure 1. **Reduction in breach cost** and resolution time when security operations use AI and automation, per Ponemon Institute research.

With AI and automation in security operations, vs. without (Ponemon, 2026)




Figure 2. **The agentic SOC** operating model divides detection, investigation, and response between agentic capabilities and human analysts at each stage.

The Agentic SOC: A New Operating Model

The agentic SOC helps security teams investigate, prioritize, and respond at machine speed while analysts retain judgment, accountability, and response authority. The model spans detection, investigation, and response, with agentic capabilities and human analysts each playing a distinct role.

Detection: Understand Behavior

- User and entity behavior analytics (UEBA)
- Agent Behavior Analytics (ABA)
- Behavioral baselining
- Dynamic risk scoring
- Pattern identification

The objective is to understand behavior, not simply collect more events.

Investigation: Build the Narrative

- Threat Center
- Machine-built timelines
- Detection grouping
- Context correlation
- Long-term behavioral visibility

Analysts need a connected narrative, not a collection of isolated alerts.

Response: Accelerate Meaningful Action

- Automation
- Security orchestration, automation, and response (SOAR)
- AI-assisted workflows
- Risk-based prioritization
- Human oversight

The goal is to reduce investigation friction and accelerate action without removing human decision authority.

The future SOC is neither fully manual nor fully autonomous. In the agentic SOC, agentic capabilities provide speed, scale, and persistent investigation assistance. Analysts provide judgment, business context, accountability, and response authority.

Agentic capabilities can help with:

- Data collection
- Correlation
- Enrichment
- Investigation support

Analysts remain responsible for:

- Prioritization
- Business context
- Risk evaluation
- Response decisions

This division of labor addresses the real bottleneck in most SOC: not expertise, but time. AI-assisted investigation and automation can reduce repetitive work so analysts can focus on decisions that require expertise and business understanding.

Measure Security at Machine Speed

Organizations also need to evolve how they measure success. A balanced measurement model can connect operational performance, risk visibility, and business outcomes.

Operational Metrics

- Mean time to detect (MTTD)
- Mean time to investigate (MTTI)
- Mean time to respond (MTTR)
- Alert reduction
- Analyst productivity

Risk Metrics

- Behavioral risk trends
- Agent visibility coverage
- Human user and service account coverage
- Investigation backlog

Business Metrics

- Governance readiness
- Compliance visibility
- Operational resilience

Executives increasingly need defensible, audit-ready visibility into AI-driven risk and measurable evidence that security is keeping pace with enterprise transformation.

A Roadmap to Machine-Speed Security Operations

Organizations don't need to rebuild security operations overnight. Five steps can establish a foundation for machine-speed security.

1. **Gain visibility into every actor:** Understand activity from people, machines, services, and AI agents.
2. **Establish behavioral baselines:** Define normal behavior before attempting to detect abnormal behavior.
3. **Prioritize risk, not alert volume:** Focus analyst attention on behavioral risk with the strongest context.
4. **Automate investigation and response:** Use AI-assisted workflows and automation to reduce repetitive work while retaining human oversight.
5. **Adopt the agentic SOC model:** Use agentic capabilities to accelerate operations while analysts retain accountability and response authority.

Together, these steps create an operational foundation for machine-speed security.




Figure 3. **Five steps** to build machine-speed security operations. Most teams already have pieces of this in place.

The Speed of Risk Is Redefining Security

The enterprise is becoming agentic. People, AI agents, service accounts, and automated workflows now operate inside the same business processes. This shift can increase speed, scale, and productivity. It also creates a new risk environment.

Risk increasingly hides inside legitimate activity. Security models built around isolated events and human-centric investigations can struggle to keep pace. Rules and preventive controls remain essential, but organizations also need behavioral context, AI-assisted investigation, automation, and measurable improvement.

Machine-speed risk requires machine-speed security operations. Behavioral intelligence helps reveal risk across sequences of activity. The agentic SOC helps teams investigate and respond faster while keeping analysts responsible for security decisions.

The future SOC succeeds by empowering people with systems designed to detect, investigate, and respond at the speed of the modern enterprise.

Assess Your Organization's Machine-Speed Readiness

- Evaluate visibility into human users, AI agents, and service accounts.
- Measure investigation bottlenecks and response delays.
- Identify opportunities for behavior-driven detection.
- Assess readiness for AI-assisted investigations and automation.
- Determine whether security operations can keep pace with the business.

In the agentic enterprise, the question isn't whether risk exists. The question is whether security operations can move fast enough to see and address it.

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.