

The Agentic SOC: Security Operations at Machine Speed

AI-Assisted Investigation and Human Oversight for Modern Security Teams

Exabeam helps security teams build an agentic SOC that combines AI-assisted investigation, automation, and analyst oversight across human and non-human identities. Rather than forcing analysts to manually assemble context across disconnected alerts, systems, and investigations, Exabeam helps security teams prioritize risk, accelerate investigations, and reduce operational friction while keeping analysts in control of response decisions. This approach enables security operations to scale alongside AI agents, autonomous workflows, service accounts, and machine identities.

Why Machine-Speed Security Operations Is Hard

The enterprise has become increasingly agentic.

AI agents now execute workflows, retrieve information, interact with systems, invoke tools, and make decisions alongside human employees. Enterprise identity no longer consists solely of people. It now includes AI agents, service accounts, APIs, automated workflows, and machine identities operating continuously across business processes.

Traditional approaches were designed for a slower operating model:

- Human-driven activity
- Sequential investigation
- Alert-by-alert analysis
- Discrete security incidents

Those assumptions are becoming increasingly unreliable. Machine-driven activity generates massive volumes of legitimate events, while risk often emerges through sequences of actions that appear normal when viewed individually. Security teams face growing investigation backlogs, increasing alert volumes, and shrinking response windows as machine-speed activity outpaces human-speed investigation.

The challenge is no longer collecting more alerts. The challenge is determining which activity represents meaningful risk quickly enough to prevent business impact.




Figure 1. Compression of the gap between risk execution and response.

Exabeam helps compress the gap between machine-speed risk and human-led response.

How the Agentic SOC Addresses Machine-Speed Risk

The agentic SOC is designed to help security teams investigate and respond in environments where AI agents, machine identities, service accounts, APIs, and humans all generate activity at a scale that outpaces manual investigation.

Instead of evaluating alerts in isolation, behavioral intelligence analyzes behavior across identities, systems, workflows, applications, and time. It helps security teams understand whether activity makes sense within its broader context and identify patterns that static rules often miss.

This shift is increasingly important because risk in the agentic enterprise often hides inside legitimate activity. AI agents, service accounts, APIs, and users frequently operate using approved permissions and valid credentials. What matters is not simply whether an action occurred, but whether that action fits expected behavioral patterns.

With behavioral intelligence, organizations can:

- Establish behavioral baselines
- Detect meaningful deviations
- Prioritize risk dynamically
- Maintain long-term behavioral context
- Correlate activity across human and non-human identities

The result is faster identification of meaningful risk and more effective prioritization of analyst attention.

The Model: The Agentic SOC

The agentic SOC is a security operations model that combines AI-assisted investigation, automation, and human decision making to help security teams operate at machine speed.

It is not autonomous security. Instead, it creates a collaborative operating model where AI accelerates investigation, context gathering, triage, and workflow execution while analysts retain responsibility for validation, prioritization, and response decisions. This approach allows organizations to increase operational speed without sacrificing oversight or accountability.

Security operations increasingly require machine-scale analysis, but machine-scale analysis alone is insufficient. Organizations still require business context, risk judgment, governance, and accountability. The agentic SOC creates a deliberate division of labor where AI helps manage investigative scale while humans retain authority over decisions and outcomes.

The goal is not to replace analysts. The goal is to help analysts spend less time gathering evidence and more time making decisions.

The Workflow: The Agentic SOC in Practice

Step 1: Activity Occurs

Humans, AI agents, service accounts, machine identities, APIs, and automated workflows continuously interact with systems and data using legitimate credentials and approved permissions.

Step 2: Behavioral Intelligence Evaluates Risk

Exabeam evaluates activity across time, identities, systems, and workflows to determine whether behavior aligns with established patterns or represents meaningful risk. Behavioral anomalies, drift, unusual workflow chains, and suspicious sequences can be identified even when individual events appear legitimate.

Step 3: AI Accelerates the Investigation

Exabeam Nova acts as a persistent AI investigator that helps gather background context, retrieve entity information, run additional searches, surface

related activity, and streamline investigative workflows. The Claude Skills Pack helps accelerate triage, case prioritization, and investigative tasks through prebuilt guidance and workflows.

Rather than generating additional alerts for analysts to review, AI helps assemble the evidence analysts need to evaluate incidents more efficiently.

Step 4: Human Analysts Validate and Decide

Analysts review findings, evaluate business impact, determine risk tolerance, and decide whether response is required.

The analyst remains accountable for the final decision.

Step 5: Automation Accelerates Response

Playbooks, SOAR capabilities, automation management, and risk-based prioritization help streamline response activities while maintaining appropriate human oversight.




Figure2.

The agentic SOC separates AI-assisted investigation from analyst-owned decisions and response.

Key Capabilities and Outcomes

Detect Risk Hidden in Legitimate Human and Agent Activity

The Challenge

Machine-speed environments generate enormous amounts of legitimate activity. AI agents, service accounts, and users may all operate using valid credentials and approved access. Risk often appears as a behavioral pattern rather than a single malicious event.

How Exabeam Works

New-Scale Analytics applies behavioral intelligence across human and non-human identities. Dynamic risk scoring, behavioral baselining, and Agent Behavior Analytics (ABA) help identify meaningful deviations from expected behavior. Exabeam evaluates activity as patterns and sequences rather than isolated alerts, helping uncover unusual workflows, anomalous tool use, behavioral drift, and potentially rogue agents.

Customer Outcome

Security teams can prioritize behavior-based risk instead of reacting to disconnected alerts.

Help Analysts Investigate at Machine Speed Without Losing Control

The Challenge

Analysts spend significant time collecting evidence, searching systems, pivoting between tools, and building investigative context before they can determine whether activity matters.

How Exabeam Works

Exabeam Nova works as a persistent AI investigator alongside analysts. Threat Center, investigation timelines, threat timelines, detection grouping, prioritization, and search capabilities help surface relevant evidence quickly. Exabeam Nova assists by gathering context, retrieving supporting information, and streamlining investigative steps that traditionally required manual effort.

Customer Outcome

Analysts spend less time assembling evidence and more time making informed decisions.

Operate an Agentic SOC at Machine Speed

The Challenge

Organizations need greater operational speed but cannot eliminate the need for judgment, governance, accountability, and business context.

How Exabeam Works

The agentic SOC establishes a collaborative model where AI assists with context gathering, evidence retrieval, investigation support, and prioritization while analysts evaluate findings and determine actions. This approach helps organizations scale investigations without surrendering control of critical decisions.

Customer Outcome

Security teams increase operational efficiency while preserving human oversight.

Close Visibility Gaps Across AI, Agent, and Non-Human Activity

The Challenge

Security teams frequently lack visibility into AI-driven activity, non-human identities, prompts, tool usage, approvals, and automated workflows. These blind spots create governance and investigation challenges.

How Exabeam Works

Okta AI Context Integration enriches investigations with identity metadata and helps establish behavioral baselines for non-human identities. Claude Enterprise ingestion captures prompts, tool calls, approvals, and outcomes within a unified investigation context. ABA helps identify anomalous patterns across this activity.

Customer Outcome

Organizations gain greater visibility into AI-driven activity and reduce investigation blind spots across human and non-human identities.

Not All Agentic SOC's Are Equal

Many agentic SOC solutions act primarily as advisors, generating summaries, recommendations, or additional alerts that analysts must manually interpret and operationalize.

The Exabeam agentic SOC combines behavioral analytics, AI-assisted investigation, Threat Center workflows, contextual enrichment, detection grouping, search, and automation to help operationalize investigations across human and non-human identities.

Instead of creating more work for analysts, Exabeam helps security teams:

- Gather context faster
- Prioritize risk more effectively
- Reduce investigation effort
- Maintain analyst control
- Accelerate security operations without sacrificing oversight

The result is an agentic SOC that helps security teams investigate and respond more effectively while keeping humans responsible for decisions and outcomes.

How Exabeam Enables the Agentic SOC

Exabeam brings the agentic SOC to life through a connected set of capabilities that combine behavioral analytics, persistent AI investigation, agent visibility, and operational workflows across human and non-human activity.

Exabeam Nova: Persistent AI investigator that gathers context, retrieves supporting information, runs additional searches, and supports investigations across Threat Center, search timelines, and Outcomes Navigator.

Claude Skills Pack: Prebuilt investigative guidance that helps streamline triage, prioritization, and investigation workflows.

Threat Center: Centralized investigation experience that helps analysts correlate, prioritize, and investigate incidents.

New-Scale Analytics: Behavioral analytics and dynamic risk scoring across human and non-human identities.

Okta AI Context integration: Establishes richer behavioral context and baselining for non-human identities.

Claude Enterprise ingestion: Normalizes AI platform activity including prompts, tool calls, approvals, and outcomes.

Agent Behavior Analytics: Detects anomalous agent behavior, drift, unusual tool use, and suspicious behavioral sequences.

Outcomes Navigator: Connects security operations improvements to measurable outcomes, governance priorities, and executive reporting.

Measuring Progress

An agentic SOC should be measured by whether security teams can reduce investigation friction, improve risk prioritization, expand visibility across human and non-human identities, and respond before machine-speed activity becomes business impact.

Operational metrics:

- Mean Time to Detect
- Mean Time to Investigate
- Mean Time to Respond
- Alert reduction
- Analyst productivity

Risk metrics:

- Behavioral risk trends
- Agent visibility coverage
- Human and non-human identity coverage
- Investigation backlog

Business metrics:

- Governance readiness
- Compliance visibility
- Operational resilience

How Outcomes Navigator Helps

Outcomes Navigator helps organizations connect machine-speed security improvements to measurable business outcomes, governance goals, operational efficiency, and executive reporting requirements.

Use Case Summary

Security teams use Exabeam to build an agentic SOC

where AI agents, users, service accounts, machine identities, APIs, and autonomous workflows operate side by side. Through AI-assisted investigation, automation, agent visibility, and operational workflows, Exabeam helps organizations identify risk hidden inside legitimate activity, accelerate investigations, reduce analyst effort, improve prioritization, and respond more effectively in machine-speed environments.

Typical Use Cases

- Detect anomalous behavior across human and non-human identities
- Investigate AI-driven activity and automated workflows
- Accelerate triage and investigation using AI-assisted workflows
- Reduce investigation backlog and analyst workload
- Improve visibility into sanctioned AI platforms and agent activity
- Measure security improvements through operational, risk, governance, and business outcomes

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.

Cloud-native SaaS delivery removes the operational overhead of installing, implementing, and maintaining another security application so your analysts can focus on security.

Flexible Integration

Pre-built connectors tightly integrate

with over 250 vendors and over 500 popular security and IT tools for threat detection. Simple new log parser app lets the analyst bring in new logs sources in minutes. Security analysts manage an increasing volume and complexity of threats with limited staff and manual processes

Differences in analyst domain experience and skill can lead to inconsistent and potentially incomplete incident scope and discovery.

Exabeam user and entity behavioral analytics (UEBA) uses machine learning to process over 1,700 built-in correlation rules and models to find advanced threats like credential-based attacks, insider threats, and ransomware activity missed by other tools. Automatically detect lateral movement.

Machine-built Smart Timelines automatically gather

evidence and assemble it into a cohesive story that can be used to perform an initial investigation. Each event registering as anomalous includes a risk score.

Extensive rule mapping enables analysts to do threat hunting based on abnormal TTPs, and use Outcome Navigator to map coverage against the MITRE ATT&CK framework.

Bringing in existing SIEM, XDR, and logs from other cloud data lakes or resources like CASB, SaaS event APIs, and Web gateways add depth and correlation potential to see end-to-end attack event strings.

Exabeam user and entity behavioral analytics (UEBA) uses machine learning to process over 1,700 built-in correlation rules and models to find advanced threats like credential-based attacks, insider threats, and ransomware activity missed by other tools. Automatically detect lateral movement.

Machine-built Smart Timelines automatically gather evidence and assemble it into a cohesive story that can be used to perform an initial investigation. Each event registering as anomalous includes a risk score.