Abnormal

NEW RESEARCH

2026 State of Behavioral Security

The Challenge of Defining Normal Across People, Machine Identities, and Autonomous AI Agents

Broad adoption of behavioral security has exposed a harder challenge: defining normal precisely enough to separate meaningful risk from noise, turn anomalies into confident action, and extend that understanding to non-human identities and autonomous AI agents.

Asked where their greatest behavioral-monitoring investment will go over the next 24 months, **43%** of security leaders chose AI agents, more than three times any other category. Almost as many, **46%**, say they are not at all effective at baselining agent behavior today.

METHODOLOGY

525
cybersecurity professionals

1,000+
employees per org

6
countries

85%
CISO-level or SOC leadership

Key Findings

01. Behavioral Security Has Gone Mainstream

Adoption is settled; precision is not.

- **81%** rely on behavioral anomaly detection to a major or extensive degree
- **87%** say behavioral baselining is mostly or comprehensively operationalized
- **91%** expect it to be very or critically important two years from now

02. The Knowing-Normal Gap

The leading obstacles to behavioral baselining aren't budget or staffing.

- **48%** cite too many false positives
- **44%** cite difficulty defining normal behavior
- **39%** cite difficulty explaining or trusting detections

03. The Understand-to-Act Gap

Detecting an anomaly is just the beginning.

- **90%** say understanding why an anomaly triggered is very or critically important to how they respond
- **51%** typically need four hours or more to investigate and remediate one
- **10%** need at least a full day

04. The Machine Identity Confidence Paradox

Confidence runs highest where coverage runs thinnest.

- **72%** say they are better prepared to detect anomalous behavior from non-human identities than from human identities
- **53%** do not establish normal for machine identities at all
- **4 of 5** non-human categories: more organizations rate themselves not at all effective than very or extremely effective

05. The AI Agent Investment Paradox

The spending is decided ahead of the baselines.

- **43%** name AI agents their greatest behavioral-monitoring investment for the next 24 months, three times API integrations at 14%
- **54%** already maintain behavioral baselines for AI agents
- **46%** say they are not at all effective at establishing and maintaining them

Evaluating Agent Behavior

- **Is this agent authorized?**
- **Is this action permitted?**
- **Is this behavior expected for this agent, pursuing this goal, on behalf of this user?**

Conclusion: The future of behavioral security will be defined not by how much unusual activity organizations can observe, but by how well they can baseline normal activity—and act when it changes.

Read the 2026 State of Behavioral Security

The complete findings, the data behind every figure, and what CISOs say about securing AI agents that make their own decisions.

Get the Full Report ›

FROM THE CISO INTERVIEWS

> **There is no plan...we don't know what the solutions are, but we're all gonna put resources to go after it."**

— CISO, B2B software organization

Figures are rounded to the nearest whole number. Multi-select questions do not total 100%. Full methodology and question wording appear in the report.

© 2026 ABNORMAL AI, INC. ALL RIGHTS RESERVED Abnormal