

# The State of Agentic Cybersecurity

How Security Leaders Are Training, Testing,  
and Trusting Humans & AI Agents Together



# EXECUTIVE SUMMARY

## AI is transforming the SOC, **but trust is lagging**

AI agents are increasingly embedded into detection, response, and playbook execution workflows of the modern SOC—reshaping how organizations perform both defensive and offensive security operations.

But while adoption is accelerating, proof of performance is not.

In a survey of 93 CISOs and senior security leaders, combined with real-world performance data from SimSpace environments, a clear pattern emerges:

- 78% of leaders report high confidence in their defenses
- 73% are already using AI agents in their SOC a moderate to high amount
- Detection and response still cluster in 1–6 hour windows
- 20% cannot consistently measure mean time to detect and respond (MTTD/MTTR)

However, extensive performance data from the **SimSpace platform** adds critical context:

- Across customer environments, initial defensive security readiness (DSR, see page 3 for definition) in early exercises commonly falls in the low-mid range, depending on team maturity.
- With repeated exercises, teams improve into the mid-high range, representing up to 50% absolute performance gains.
- Improvement follows a learning curve: the largest gains occur in early repetitions, with continued improvement over time.
- Measurable trust is exponential when humans and AI agents are trained and tested together in realistic, production-like environments — what we at SimSpace call the **AI Proving Grounds**.

The conclusion is clear: AI is being deployed into SOC environments faster than it is being tested, measured, proven, and trusted.

How are security teams currently testing their AI SOC? And how can security leaders close the gap between agent deployment and proven AI trust? **Let's dive in.**

# A Note on Defining “Readiness”

Throughout this report, we’ll discuss how security teams attain and measure readiness, particularly as it pertains to the AI threat landscape. And while “readiness” often refers specifically to individual and team training, **SimSpace measures readiness and resilience together**, as we believe that training and testing human operators and AI agents go hand in hand.

## Definition: Defensive Security Readiness

At SimSpace, we use a proprietary metric called defensive security readiness (DSR), and is a performance index measured in SimSpace training and testing. DSR measures how effectively security teams, tools, and agents detect and respond to adversarial activity during realistic attack simulations.

## Measuring Defensive Security Readiness

At a high level, DSR represents the proportion of attack activity that is correctly detected, investigated, and acted upon within a simulated environment. Rather than measuring tool output (e.g., alerts generated), DSR measures operational outcomes:

- Did the team or AI agent identify the attack?
- Did they investigate it correctly?
- Did they take appropriate action?

### **DSR is derived from structured exercises in which:**

- Realistic adversary behaviors are executed in a controlled environment
- Security teams interact with their AI agents, workflows, and playbooks
- Detection and response actions are observed and scored

### **Across these exercises, performance is evaluated based on:**

- Detection coverage (what was identified vs. missed)
- Accuracy of analysis
- Effectiveness of response actions

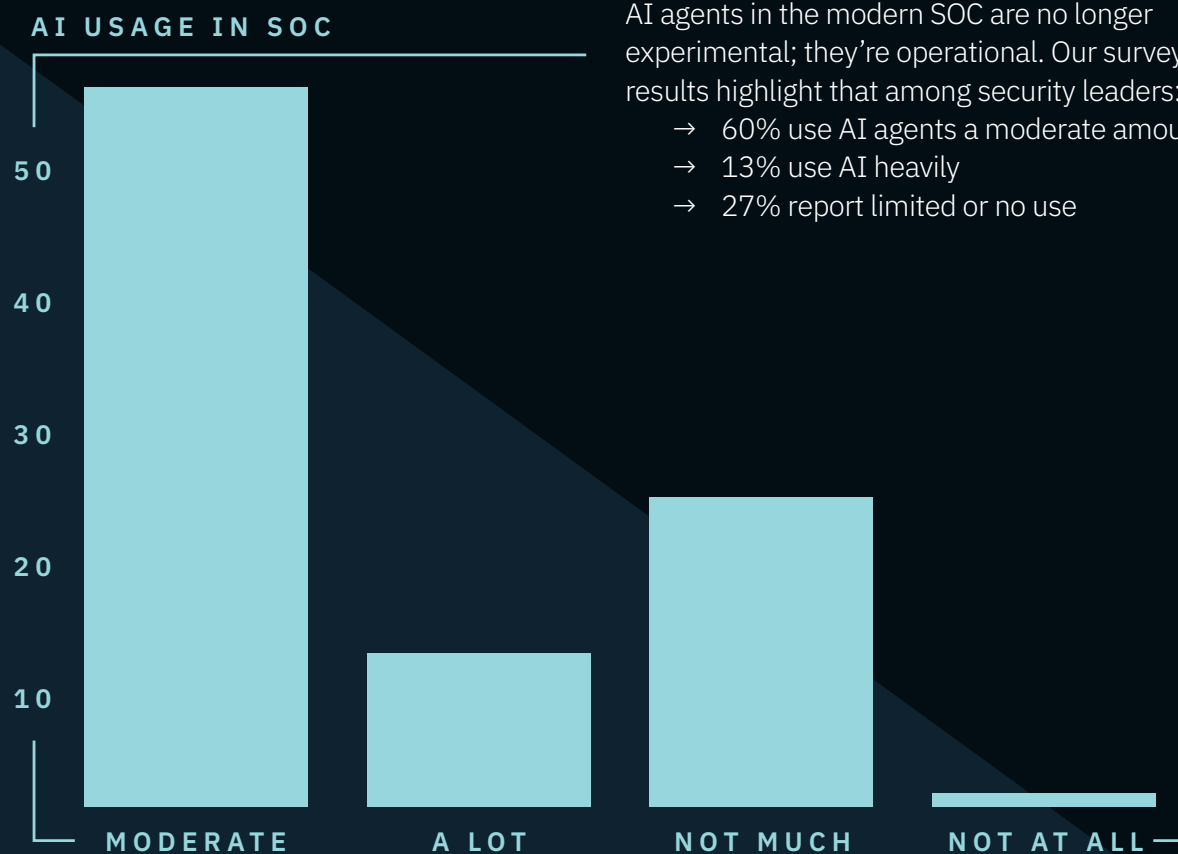
### **SimSpace aggregates these outcomes into a normalized score (0 to 1), where:**

0.0 = No meaningful detection or response

1.0 = Complete and correct detection and response across all scenarios

# The Rise of the Agentive SOC

## From human-centric defense to human + AI agent collaboration



AI agents in the modern SOC are no longer experimental; they're operational. Our survey results highlight that among security leaders:

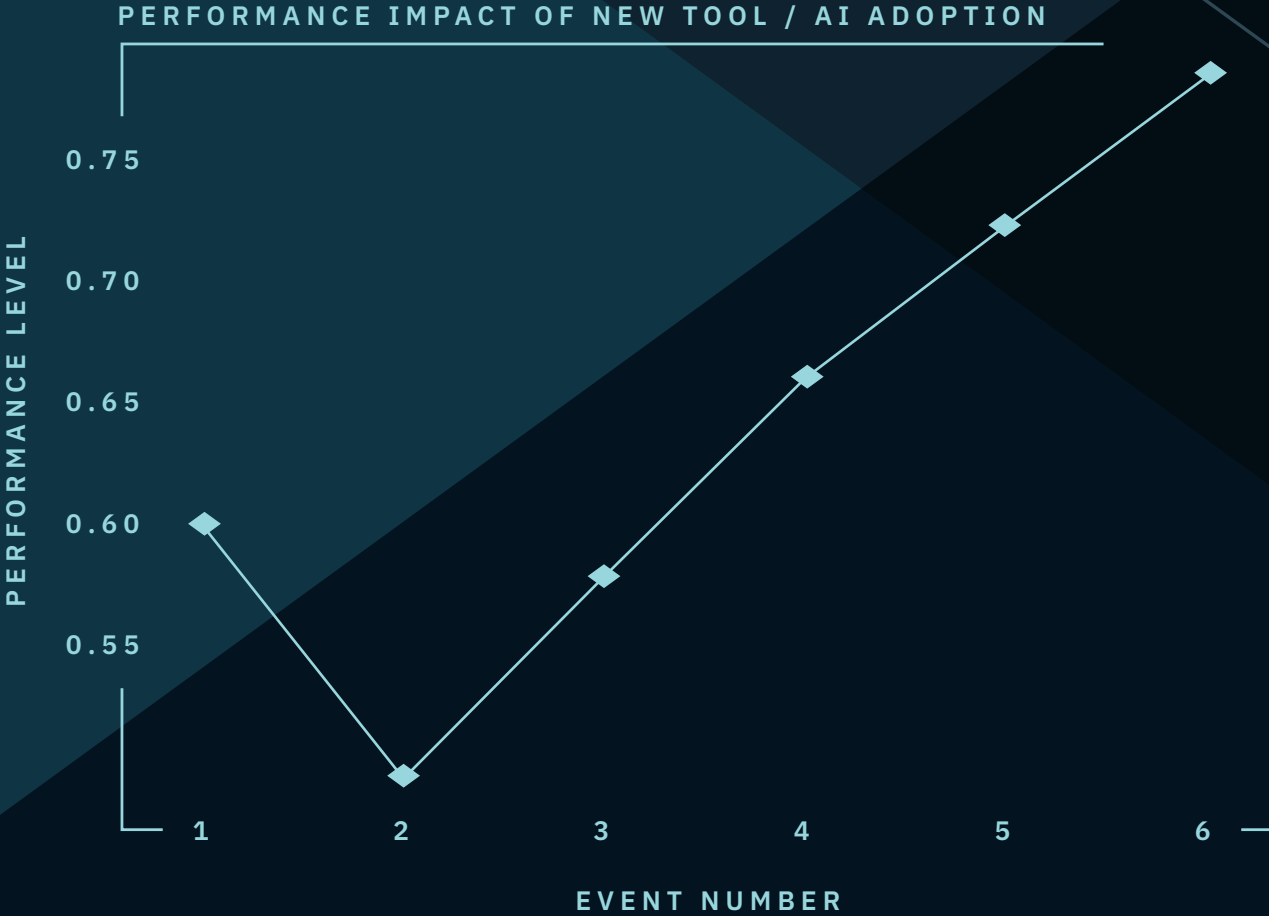
- 60% use AI agents a moderate amount
- 13% use AI heavily
- 27% report limited or no use

AI agents are being applied across:

- SOC augmentation
- Playbook automation
- Tool optimization

**However, most organizations remain in augmentation mode, not full autonomy.**

Data from the SimSpace platform shows **what happens when new technologies** like AI agents are introduced:



**Implication:**


AI agents are already influencing SOC performance, but without structured testing and performance measurement, they initially introduce disruption and negatively impact the efficiencies they're designed to improve.

# The AI Confidence Gap

**High confidence in defenses;  
limited proof of agent performance**

Survey data shows that security leaders are confident about their AI SOC: 77% respondents rated their confidence at 4-5/5.

CONFIDENCE DISTRIBUTION



## However:

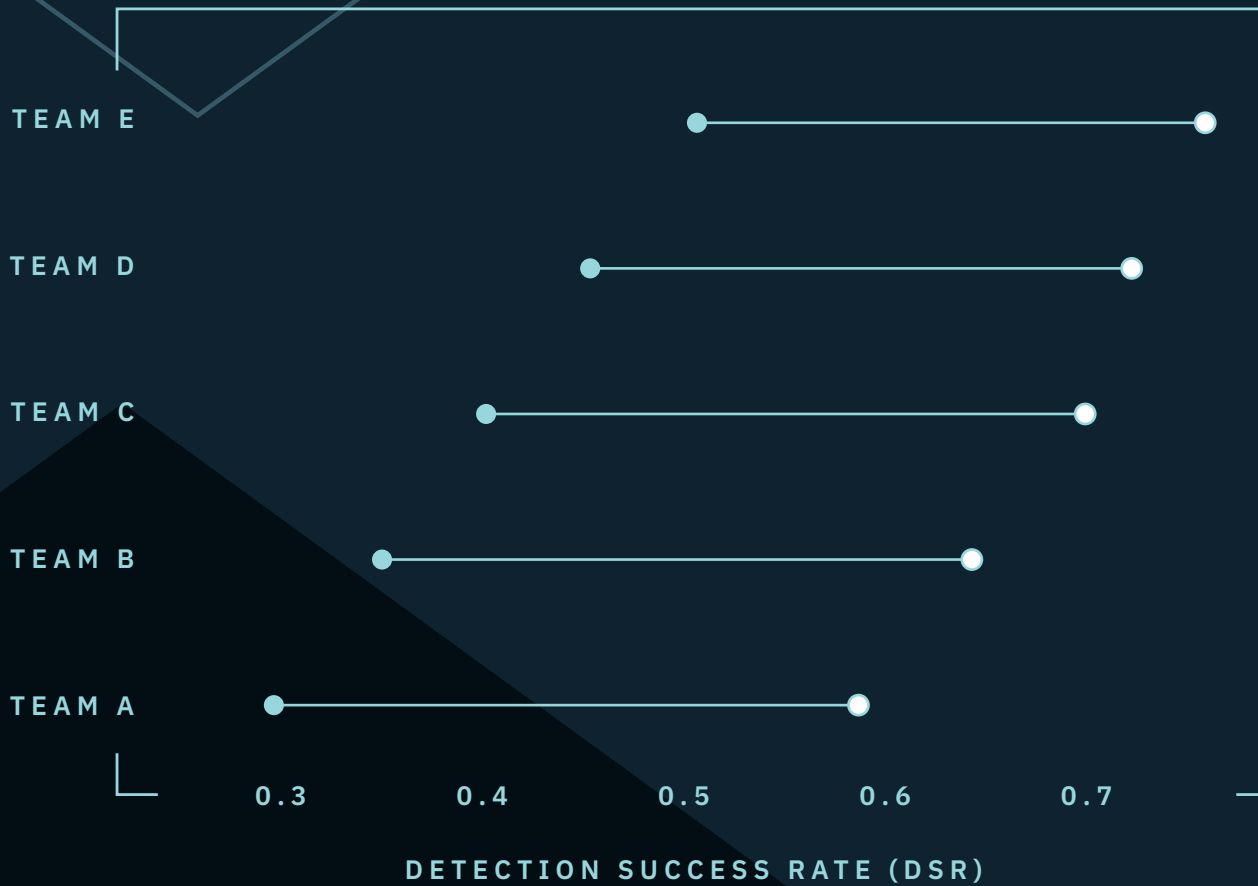
- Only 27% express maximum confidence
- 20% cannot consistently measure detection or response performance

# At the same time, **SimSpace data reveals:**

Teams with similar self-reported maturity can show DSR variability as low as 0.30 out of 1.0 across initial events. Even within the same organization, performance between events can vary by ~10-25 percentage points, depending on:

- Scenario complexity
- Tool configuration
- Team familiarity

DSR VARIABILITY ACROSS TEAMS WITH SIMILAR MATURITY



## Implication:

Organizations report high confidence, but measured performance shows significant variability that is only uncovered through realistic testing.

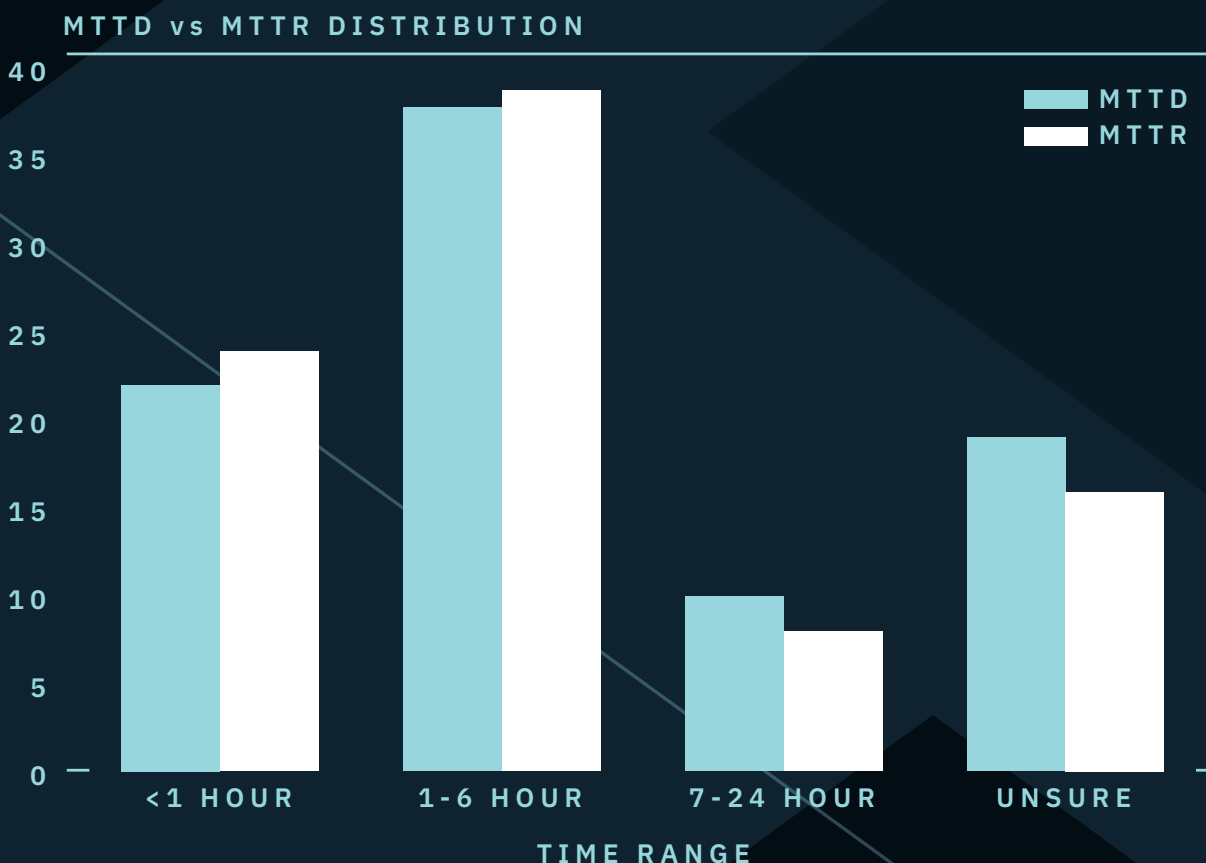
# Detection and Response in an AI-Driven Environment

## MTTD

41% → 1–6 hours  
24% → <1 hour  
21% → unsure

## MTTR

42% → 1–6 hours  
26% → <1 hour  
17% → unsure



As cyber attack breakout time approaches zero, security teams need to be detecting and responding to threats at machine speed. But despite rapid AI adoption, only an average of 25% of teams say they detect and respond to threats in under an hour, and most MTTD/MTTR rates cluster in 1-6 hours. Meanwhile, an average of 20% of teams can't effectively measure MTTD/MTTR performance at all. **When every threat is a zero-day, there's no room for uncertainty.**

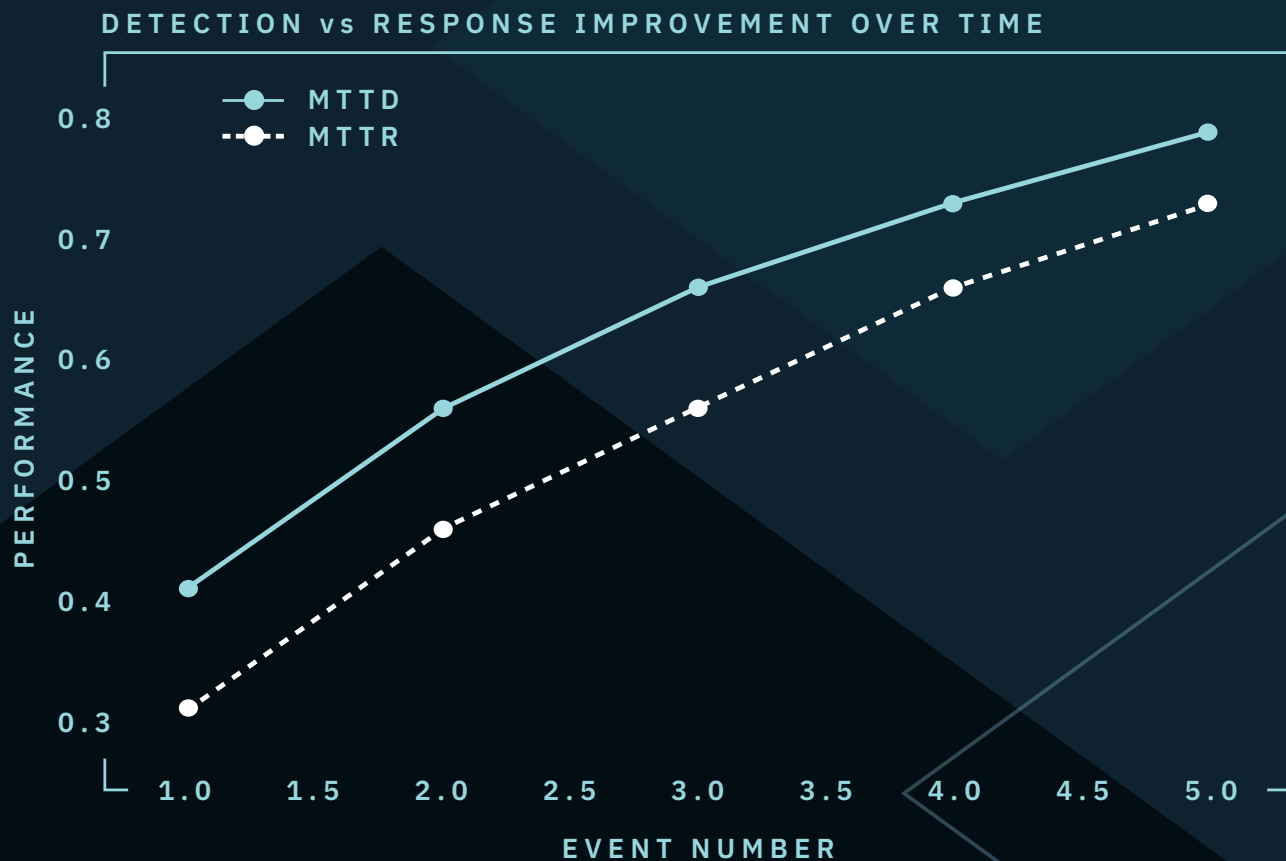
# Data from the SimSpace platform helps explain why:

In early exercises:

- Response precision and execution accuracy frequently lag detection performance
- Teams may detect threats but fail to execute optimal response workflows consistently

With repeated testing:

- Response precision improves alongside detection
- Typical improvement in response-related metrics ranges 15-40% over successive events



## Implication:

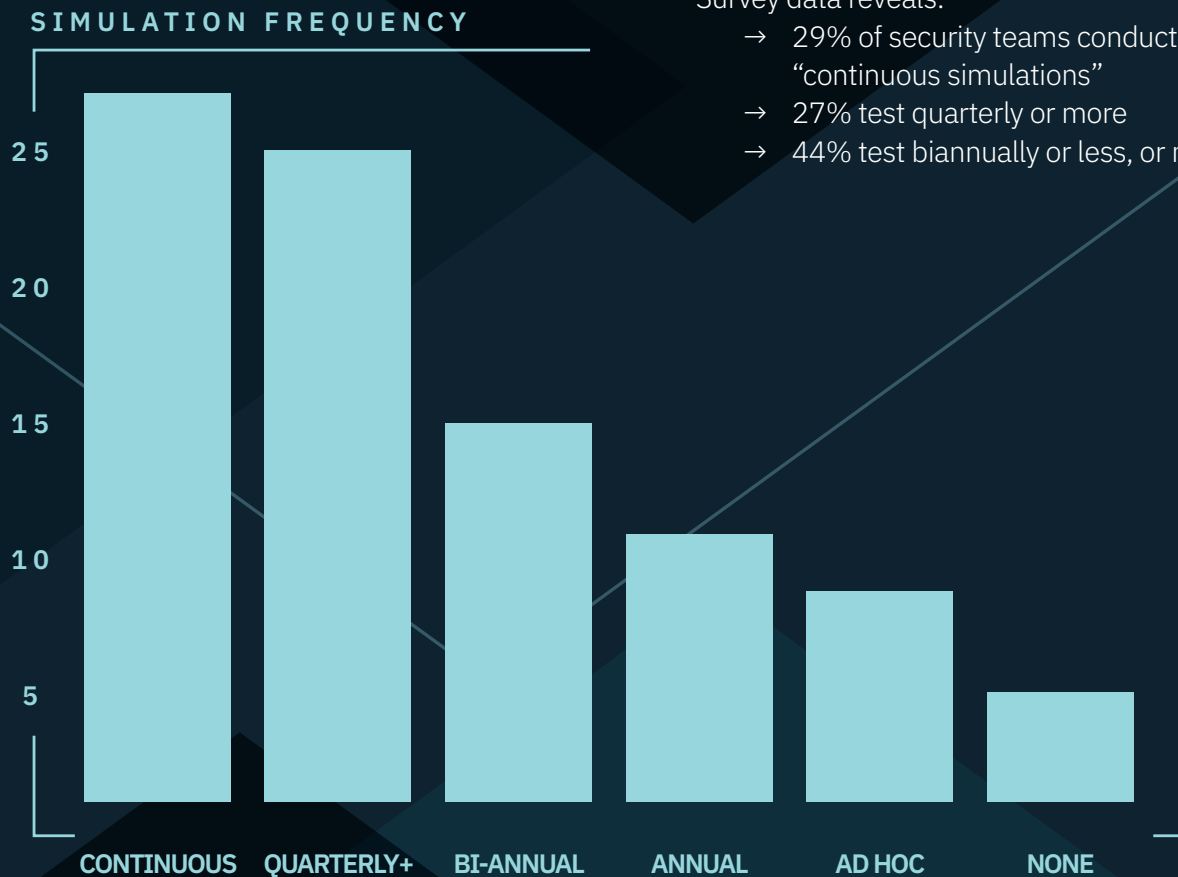
AI may accelerate detection, but without repeated testing within a realistic cyber simulation infrastructure, response execution remains inconsistent and human-dependent.

# Are Organizations Testing Their AI Agents, or Just Their Teams?

## Continuous vs. episodic testing

Survey data reveals:

- 29% of security teams conduct “continuous simulations”
- 27% test quarterly or more
- 44% test biannually or less, or not at all



Meanwhile, AI agents are operating continuously and automated decisions are made daily. This means most teams aren't testing their **teams and agents** enough to keep up with ongoing agentic workflows.

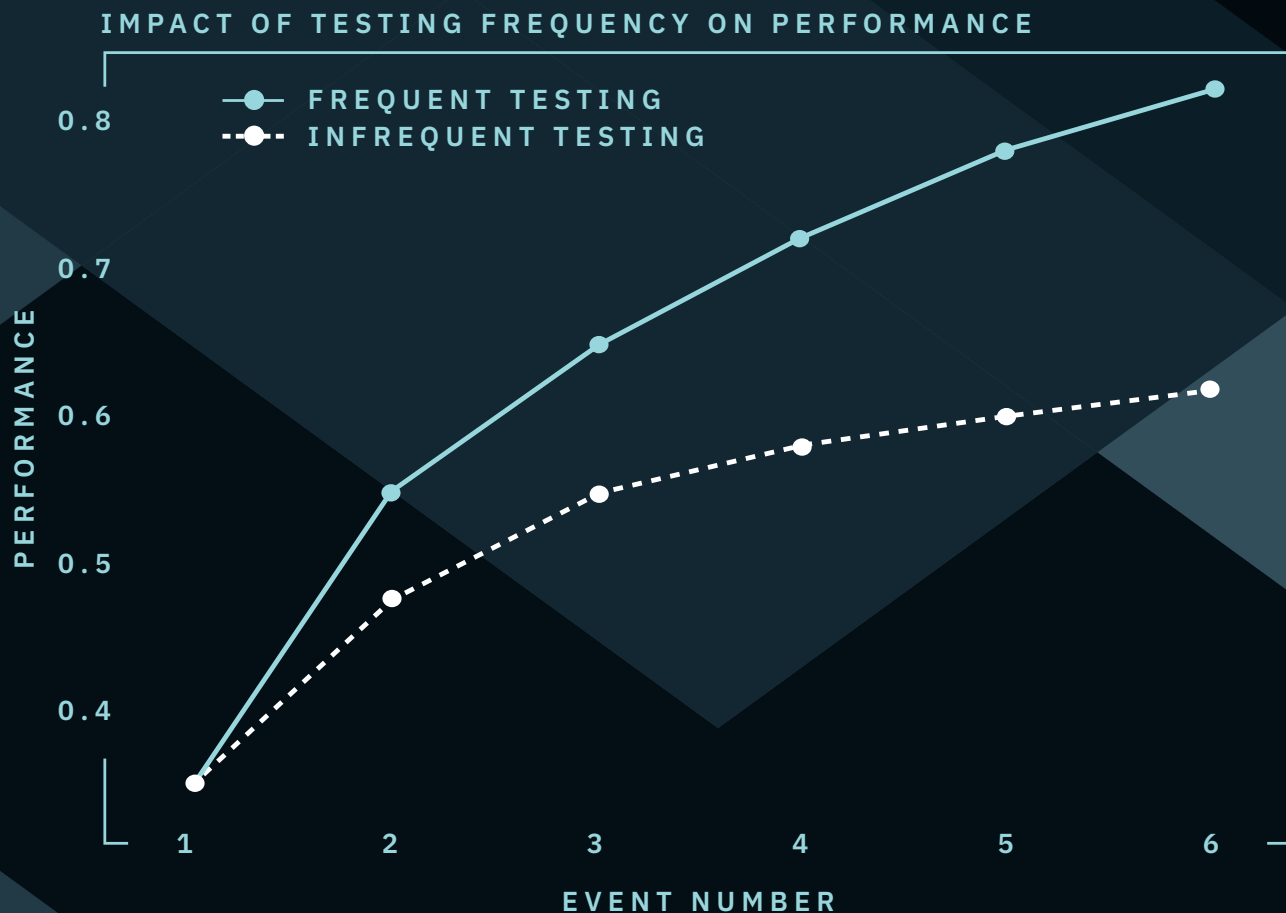
# SimSpace data shows the impact of testing frequency and what “continuous simulation testing” really means:

Teams participating in frequent, full-scope events:

- Improve DSR by 0.03-0.05 per event on average
- Reach 0.70-0.80+ performance levels within 4-6 iterations

Teams using limited-scope or infrequent events:

- Improve more slowly
- Often plateau in the 0.50-0.65 range



## Implication:

Continuous validation with realistic simulation testing produces measurable gains, while episodic testing limits performance and leaves gaps in AI-driven decision validation.


# The Measurement Problem: Why AI-Driven Readiness Is Hard to Prove

Survey data shows that about 20% of security teams lack consistent MTTD/MTTR measurement, with the top barriers including a **lack of unified metrics, siloed tools, and fragmented workflows.**

A singular source for training and testing teams alongside AI agents helps CISOs gain much more consistent measurement and deliver proof of agentic readiness. **We call this the AI Proving Grounds** (more on this on page 19).

**SimSpace platform data shows how organizations can truly measure readiness.**

MEASURABLE CYBER READINESS RANGES (DSR)



Early-stage teams typically operate between 0.25–0.50 DSR, while mature teams achieve 0.60–0.80+, demonstrating measurable progression.




Cyber readiness improves incrementally with repeated testing. Organizations typically see performance gains of 0.03–0.05 DSR per event, reinforcing that readiness is a progression, not a static score. These same measurement principles apply to AI-augmented systems, where detection accuracy, response precision, and decision quality must be continuously evaluated under real-world conditions. Additionally performance progression can be modeled and predicted using learning curve functions, showing consistent improvement trends across environments.

## Implication:

Organizations don't lack measurable signals. It lacks environments that consistently generate them, especially for AI agent workflows.

# Training in the Age of AI

Survey data shows that training remains dominated by tabletop exercises, certification programs, and tool-specific training.

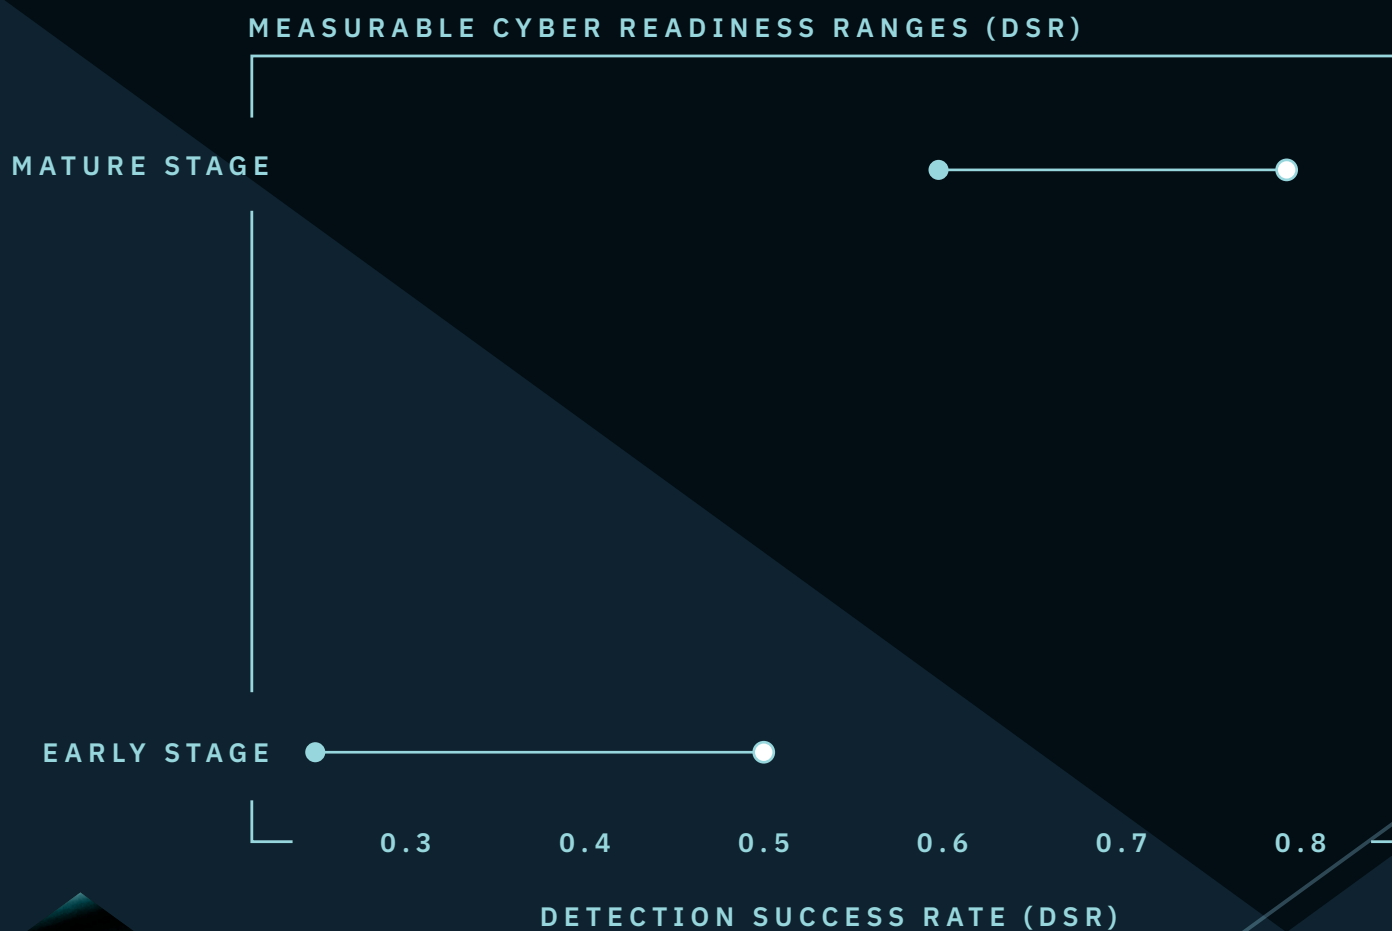


The recently published **Forrester Wave™: Cybersecurity Skills And Training Platforms, Q1 2026** establishes that the traditional training methods like these are no longer enough for teams to keep pace with AI-driven attacks.

Forrester’s evaluation notes that organizations need training that “prepares practitioners to work alongside and evaluate AI-generated outputs for accuracy.” This means training programs must evolve to teach security how to operate alongside AI agents.

**And data from the SimSpace platform shows which training methodologies actually drive performance:**

Hands-on, realistic environments produce 20–50% performance improvements over repeated exercises.

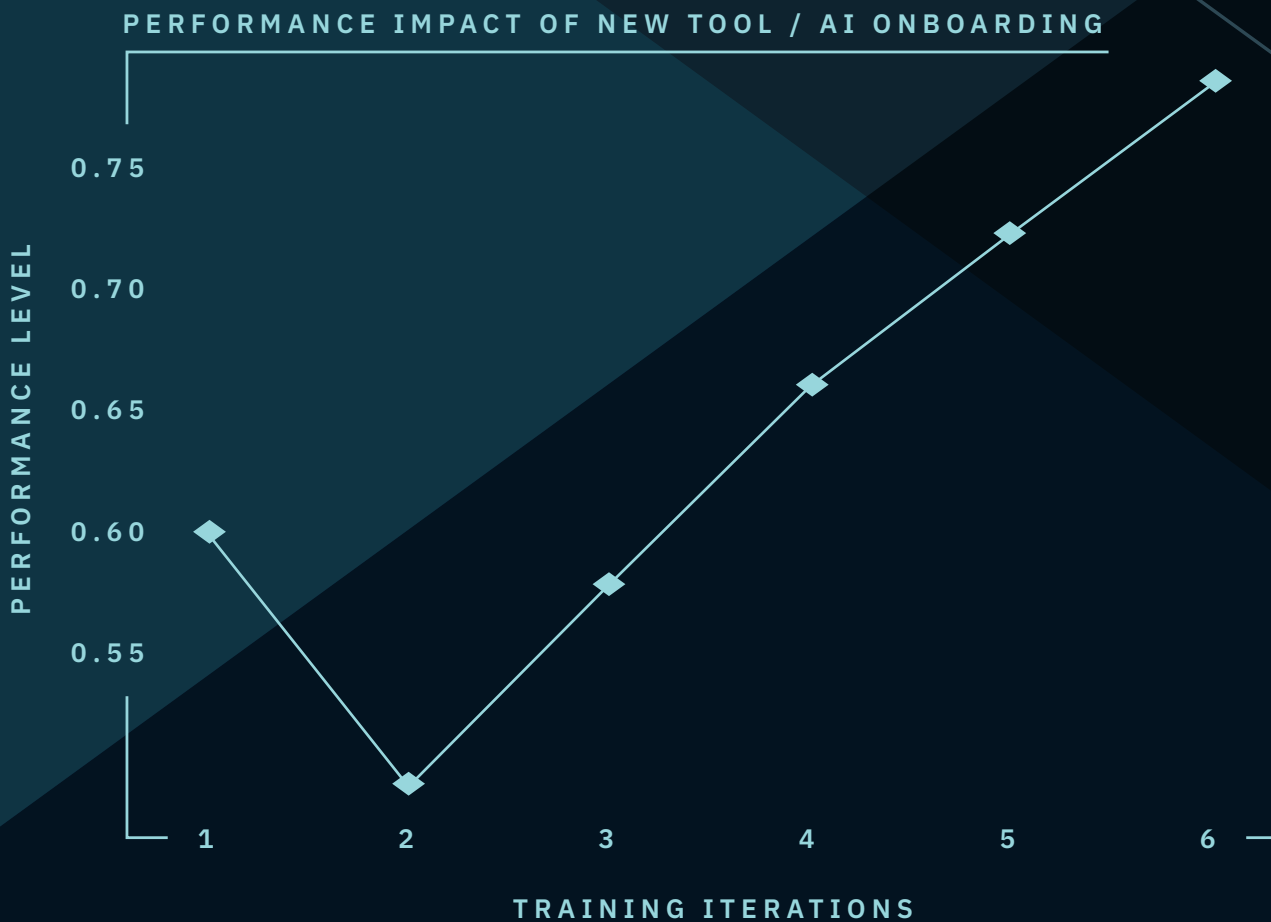


**Initial performance:** ~0.25 → 0.50.

**After repeated testing:** ~0.60 → 0.80+ shows a 20%-50%

# Tool onboarding scenarios

When onboarding tools, you'll see that initial performance dips ~10-20%.

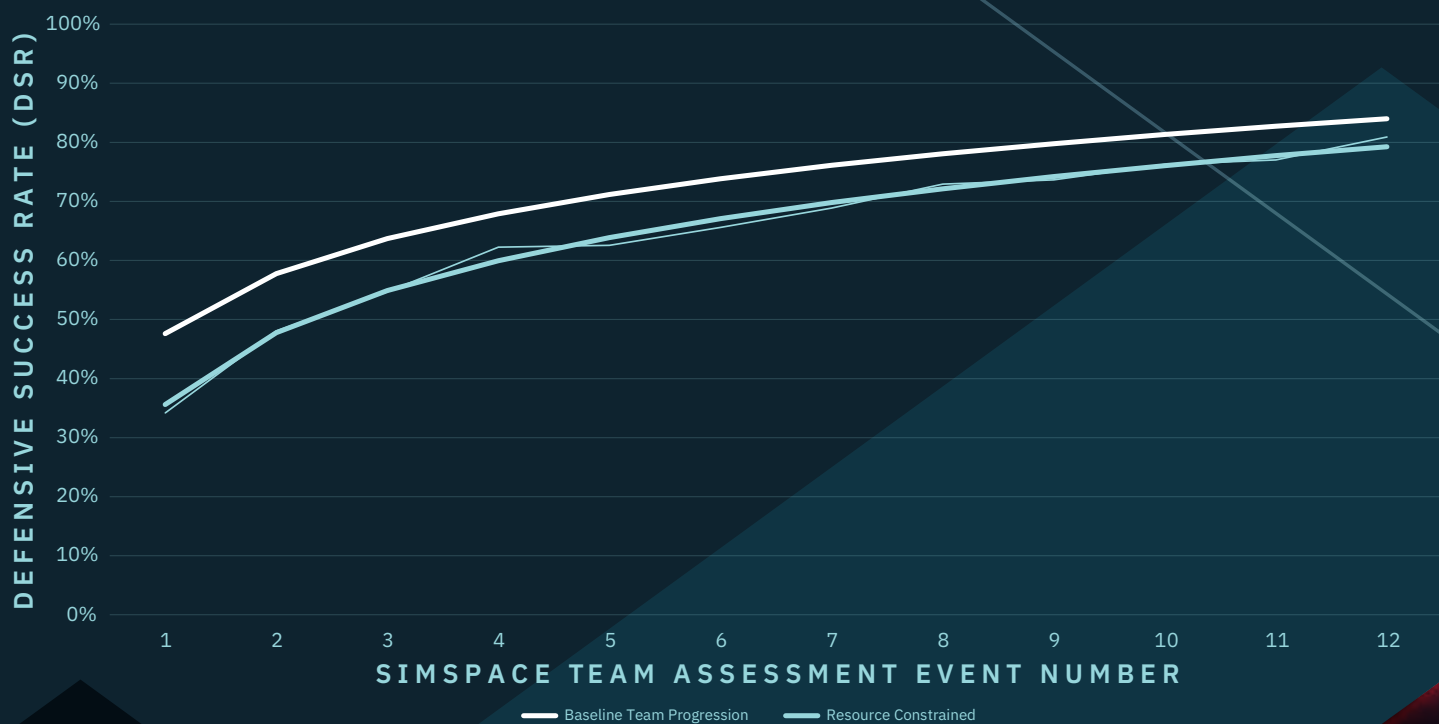


New tools and AI agents can initially degrade performance before delivering gains. **With structured, iterative AI testing in unified cyber simulation platform, teams can move beyond that initial disruption to optimized performance.**

# Resource-constrained teams

Teams with fewer resources tend to start at lower DSR (~0.25–0.40). This is to be expected, as smaller teams may have a harder time detecting, prioritizing, and mitigating persistent threats.

### READINESS PROGRESSION FOR RESOURCE CONSTRAINED ORG



Iterative exercises alongside AI agents can have even greater implications for these teams, **achieving 70%+ improvement relative to the baseline.**

## Implication:

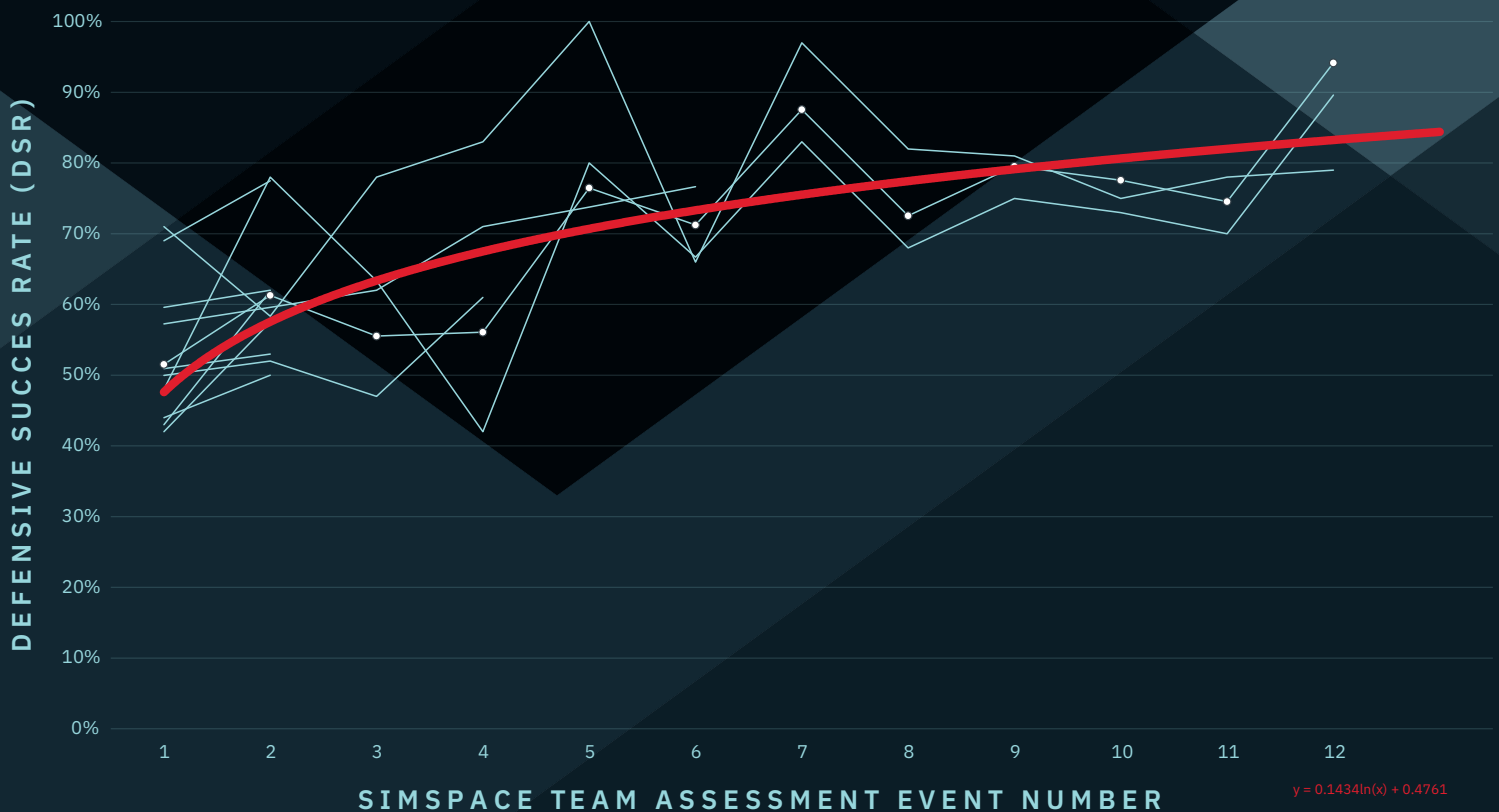
AI agents, and the humans who rely on them, must be trained through iterative, realistic threat simulation, not static instruction.

# Agentic Readiness as a Learning Curve, Not a Point-in-Time

Collectively, the SimSpace platform data highlights that performance follows a logarithmic progression curve:

- Event 1 → large improvement
- Events 2–4 → steady gains
- Later events → incremental refinement

## HUMAN & AGENT TEAM READINESS PROGRESSION



But our survey data reinforces the gap between confidence and measurement, with 77% of respondents feeling moderate to high confidence in their AI SOC, while less than 20% believe they can consistently measure their performance.

## Implication:

Many organizations treat AI readiness as a snapshot, but real-world data shows it is a measurable progression over time.

# How Security Leaders Can **Close the Agentic Confidence Gap**

## 1. Shift from episodic to continuous, realistic agent testing

AI-driven systems operate continuously—but many organizations still validate performance periodically.

To keep pace, security leaders must move toward:

- Continuous validation loops, not one-off exercises
- Regular testing of detection, response, and AI-driven workflows with realistic simulations

Organizations that test more frequently achieve measurably higher performance, while those that test less often plateau.

## 2. Measure what matters

As AI agents become embedded in SOC operations, traditional metrics are no longer sufficient.

Leaders must measure:

- Detection success
- Response accuracy
- Decision quality across human and AI workflows

Without outcome-based metrics, organizations cannot determine whether AI is improving performance or introducing new risk.

## 3. Plan for the Learning Curve

AI adoption is not frictionless.

SimSpace data shows that AI agents often introduce initial performance dips of ~10–20%, followed by steady improvement with repeated testing.

Organizations that anticipate this learning curve and improve through iterative testing are far more likely to realize long-term gains.

## 4. Invest in your AI Proving Grounds

The most critical step in closing the agentic confidence gap is establishing an environment where AI can be continuously tested and proven alongside human operators.

As AI agents take on a larger role in detection, investigation, and response, organizations must move beyond deploying automation and begin testing it under real-world conditions.

This requires **AI Proving Grounds**: A realistic, controlled replica of your production environment where organizations can:

- Generate synthetic, precision-labeled telemetry training data and retrain with customer data
- Test and validate agentic solutions
- Ensure agentic solutions are safe and trustworthy
- Strengthen cyber resilience with AI agents alongside human operators

Without AI Proving Grounds, organizations are effectively deploying AI into production environments without ever fully understanding how it behaves under stress, where it fails, or how it interacts with human operators.

In contrast, organizations that invest in these environments gain something far more valuable than automation: They gain trust.

**Trust that their AI agents work.  
Trust that their teams can rely on it.  
And trust that their combined human-agent operations will perform when it matters most.**

# Methodology

This report combines insights from two primary sources:

1. An independent survey of cybersecurity leaders conducted by Gatepoint Research, and
2. Aggregated performance data from the SimSpace platform.

## Survey methodology

The survey data featured in this report was conducted by Gatepoint Research between December 2025 and March 2026 as part of a study on cybersecurity preparedness approaches.

A total of 93 respondents participated, representing a cross-section of senior security, IT, and risk leadership roles across large enterprises.

### Respondent profile:

- 21% CIOs and CISOs
- 27% Vice Presidents
- 42% Directors
- 10% Senior or department managers

### Organizational characteristics:

The majority of respondents represent large enterprises with \$500M+ in annual revenue. Approximately 60% operate in highly regulated industries, including:

- Financial services
- Government and defense
- Healthcare and life sciences
- Energy and utilities

## SimSpace platform data methodology

To complement survey insights, this report incorporates aggregated performance data from the SimSpace platform, based on real-world cybersecurity exercises conducted across customer environments.

### Data source and scope:

SimSpace data is derived from:

- Human & agent eam-based cyber readiness tests
- Adversary simulation exercises
- Structured training and validation events

The dataset includes:

- Multiple organizations across industries
- Repeated events over time
- A range of team maturity levels and resource profiles

### Interpretation and limitations:

Data reflects observed performance in simulated environments, not self-reported perception.

Individual results will vary based on:

- Team experience and maturity
- Tool configuration and integration
- Scenario complexity and scope

# SimSpace is the **AI Proving Grounds**

Allied governments, militaries, commercial enterprises, and research universities worldwide trust SimSpace as the AI Proving Grounds where human operators and AI agents train and test together in a realistic replica of their production environments to outperform and outsmart any adversary in any terrain.

## **Train with Synthetic Data**

Generate synthetic, real-world telemetry training data and retrain with customer data.

## **Live-Fire AI Agent Testing**

Test and validate agentic solutions under realistic conditions.

## **Prove & Trust AI Agents**

Ensure agentic solutions are safe and trustworthy.

## **Strengthen Humans & AI Together**

Strengthen cyber resilience with human operators alongside AI agents.

To see the AI Proving Grounds in action, [schedule a demo](#) with the **SimSpace** team.

