Synack | Omdia by informa techtarget

The 2026 State of Agentic AI in Pentesting

Insights from 200 security leaders on the future of AI-driven offensive security

The long-term shift has begun.

Here's what early adopters are telling us about it.

AI is rewriting the rules of offensive security.

As AI-enabled adversaries become more prevalent, security teams are adopting their own AI agents to stay ahead. This change marks the beginning of a long-term shift from human-led penetration testing toward a hybrid approach incorporating agentic AI.

Still a mission-critical priority for **95%** of organizations, pentesting now includes agentic AI as a means to test more assets, more effectively.

As enterprises face a rapidly expanding attack surface that strains traditional methodologies, agentic AI is emerging as the scalable, continuous solution. But the overall vision of agentic AI in pentesting is still unfolding.

Synack

The dawn of the early adopter

87% OF ORGANIZATIONS HAVE MOVED BEYOND THE EVALUATION PHASE

Currently, 87% of organizations have moved beyond the evaluation phase and are either actively planning a pilot, currently testing, or have already integrated agentic AI into their pentesting programs. These early adopters are the vanguard of a movement toward agent-led security operations.

Trust is a primary catalyst

87% OF ORGANIZATIONS TRUST AGENTIC AI TO TEST THEIR ENTERPRISE ENVIRONMENTS

This rapid adoption is underpinned by an extraordinary level of confidence in the technology's efficacy. In fact, 87% of organizations surveyed report a high or complete level of trust in agentic AI to effectively test their enterprise environments. Notably, organizations that have fully deployed these systems are 2.2X more likely to express complete trust compared to those still in the pilot phase.

A willingness to go all-in

95% OF ORGANIZATIONS ANTICIPATE THAT AGENTIC AI WILL REPLACE THEIR TRADITIONAL PENTESTING SERVICES

The industry's expectations for this technology are profound—95% of surveyed organizations anticipate that agentic AI will displace traditional pentesting services, though the degree varies: 49% expect complete or significant displacement. Moreover, one in four organizations expect to conduct pentesting exclusively through agentic AI within the next three years—a projected increase of 67% from current adoption levels. Advanced users are particularly bullish. Organizations already utilizing agentic AI are 1.4X more likely to believe these systems will completely replace traditional services compared to those in the pilot phase.

Humans + AI is the gold standard

64% OF ORGANIZATIONS IDENTIFY AGENT-LED, HUMAN OVERSIGHT AS THEIR PREFERRED OPERATIONAL MODEL

Keeping humans in the loop allows organizations to implement the scalability of machines with the safety net of human expertise. In fact, 64% of organizations identify agent-led, human oversight as their preferred operational model.

The Pentesting Paradox:
A top priority that's under-executed

When it comes to pentesting, a scalability wall exists: While 95% of organizations rank pentesting as a top or high priority, on average, they are **pentesting only 32% of their attack surface**.

This leaves a coverage gap where **68% of the environment** is untested, creating blind spots that adversaries are increasingly adept at exploiting.

For roughly one-third of organizations, the situation is even more critical, with **20% or less** of their infrastructure receiving regular assessment.

SECURITY LEADERS CITE SEVERAL FRICTION POINTS THAT PREVENT MANUAL TESTING FROM SCALING

Limited Scale

Manual testing that relies heavily on human talent can be difficult to scale, making it challenging to keep pace with rapid application development cycles and dynamic cloud-native environments.

Ineffective Triage and Communication

Over half (55%) of organizations report that traditional testing struggles to effectively communicate findings to key stakeholders.

High Costs

Because high-quality manual testing services are resource-intensive, they can be expensive and are often relegated to periodic snapshots.

Lack of Accurate Risk Identification

Organizations are grappling with internal inefficiencies where security personnel may not take action due to a lack of context.

Agentic AI for pentesting has concrete advantages

With agentic AI for pentesting, security teams are able to test a wider attack surface and close the coverage gap. In fact, 45% of organizations that are already using agentic AI for pentesting are making it their No. 1 priority. Fueling this transition is the demand for improved security as well as operational efficiencies.

Accelerated path to remediation

Agentic AI unlocks faster execution and time to results, a primary driver for adopting AI-based pentesting. By **identifying vulnerabilities in hours rather than weeks**, security teams can synchronize their testing with rapid development cycles, ensuring that critical findings are remediated before they can be exploited in a production environment.

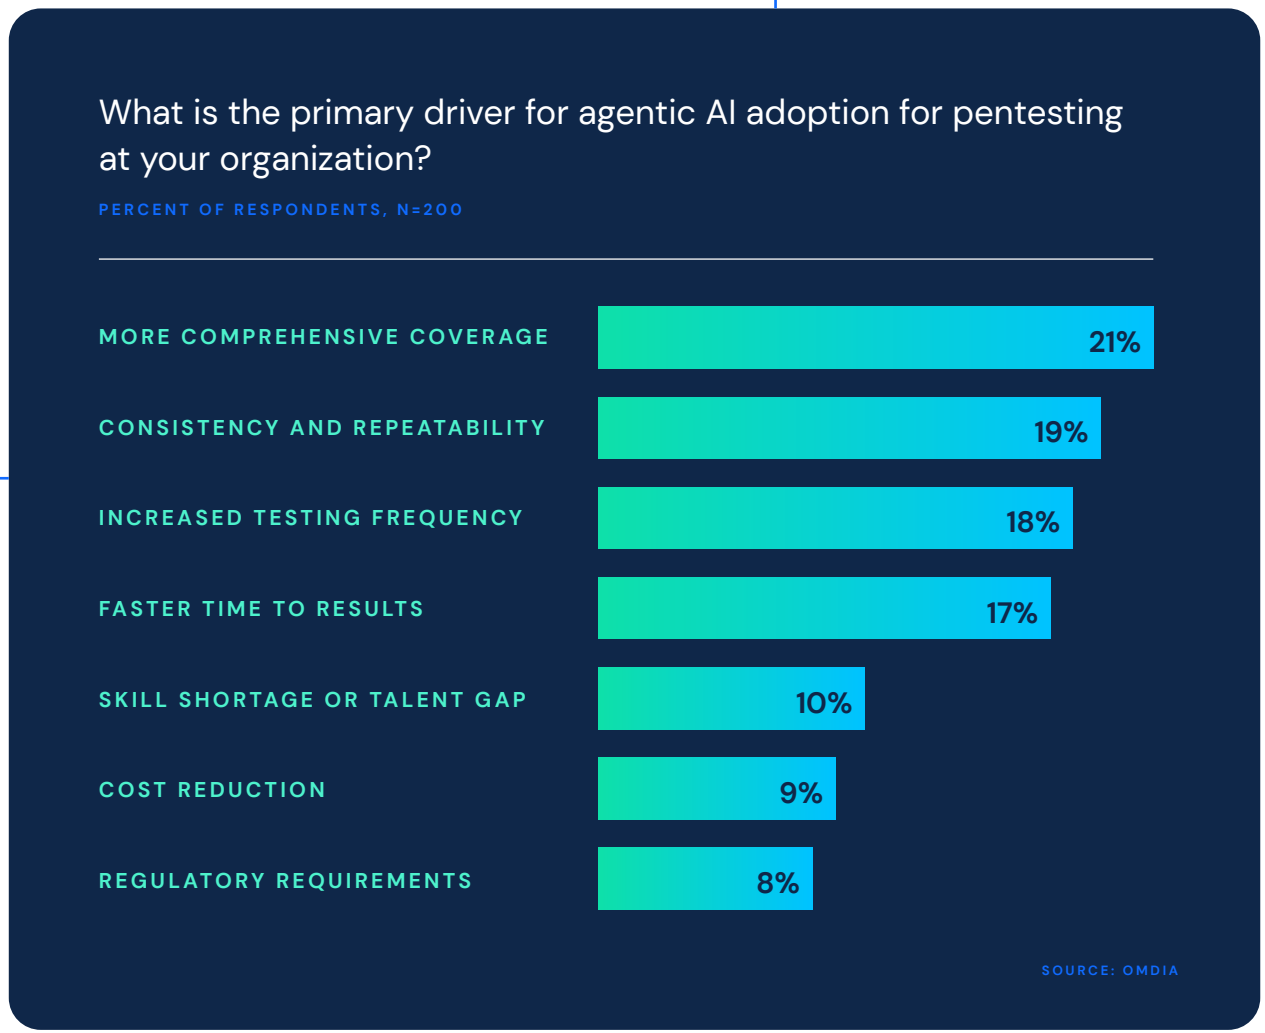

Synack

Accuracy reaches new levels

For agentic AI to be considered a viable component of a modern security stack, it must meet a rigorous precision mandate set by security leadership.

85% IS THE LEVEL OF ACCURACY THAT 69% OF ORGANIZATIONS SAID WOULD BE CONSIDERED EFFECTIVE WHEN COMPARED WITH MANUAL TESTING.

Compared with traditional manual testing, what level of testing accuracy would you require from agentic AI for pentesting to consider it effective?

PERCENT OF RESPONDENTS, N=200

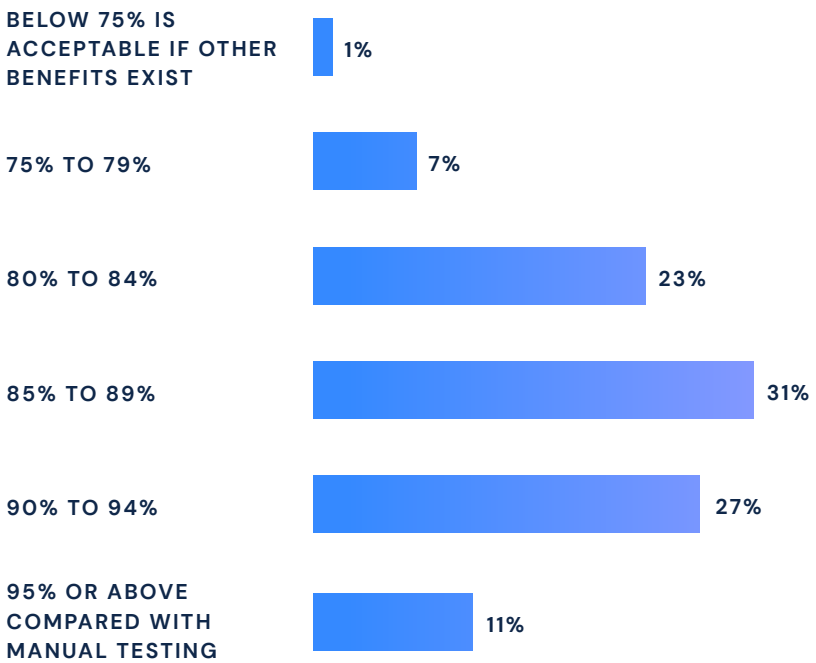

SOURCE: OMDIA

Accuracy is a function of detection breadth, logic, context, false positives, and adaptability. AI-driven pentests can scan thousands of assets and ports in minutes, while humans are more constrained by time. AI can be more accurate at scale because it will test more attack paths and cover more ground.

AI's chaining constraint

Still, AI-led testing may miss some logic and context vulnerabilities. Vertical vulnerability chaining is becoming more commonplace in agentic pentesting, but most AI solutions can't do horizontal chaining the way an advanced pentester could. Vertical chaining escalates privileges to gain deeper control within a single system, whereas horizontal chaining pivots across different systems or accounts to expand the attack's breadth across the organization. Additionally, purely autonomous approaches can have much higher rates of false positives.

Synack

Attack surface management is easier, across a larger footprint

For those that have deployed agentic AI for pentesting, over half said that testing their attack surface has become easier over the past 12 months.

This shift is particularly evident among organizations that have already deployed agentic AI for pentesting, where **60%** indicate that the task is significantly or somewhat easier today—a rate **1.3X higher** than those still in the pilot phase.

WHEN ATTACK SURFACE MANAGEMENT IS EASIER, SECURITY TEAMS CAN SCALE PENTESTING.

Increased Visibility

87% of organizations report that agentic AI for pentesting has made it easier to test a higher percentage of their global attack surface.

Continuous Coverage

91% of organizations report that agentic AI for pentesting has allowed them to conduct pentesting on a more regular basis.

Please rate your level of agreement with the following statements as they relate to your organization's current approach to pentesting.

PERCENT OF RESPONDENTS

STRONGLY AGREE | AGREE | NEUTRAL | DISAGREE | STRONGLY DISAGREE | DON'T KNOW

Agentic AI for pentesting has enabled my organization to test a higher percentage of its global attack surface.

N=83, 87% AGREE

Using agentic AI for pentesting has allowed my organization to conduct pentesting of its attack surface on a more regular basis.

N=83, 91% AGREE

My organization believes using agentic AI for pentesting will allow it to conduct pentesting of its attack surface on a more regular basis.

N=117, 80% AGREE

SOURCE: OMDIA

Synack

Early adopters are protecting a full breadth of assets across their portfolios

The focus for pentesting is shifting toward high-growth, complex environments that are often mission critical. Early adopters are moving beyond legacy boundaries and prioritizing dynamic assets that have historically had limited visibility.

TOP ASSET PRIORITIES FOR EARLY ADOPTERS

Cloud-Native Applications

59%

This is the top priority for organizations, reflecting the need to secure highly distributed and containerized environments that shift and scale rapidly.

AI Systems

57%

Organizations are increasingly focused on testing the security of their own AI deployments, recognizing that the systems defending the enterprise must also be hardened against attack.

Server Infrastructure

55%

Protecting the underlying hardware and virtual servers that support the organization remains a top-three priority.

Cloud Infrastructure

54%

Organizations are prioritizing the security of their cloud infrastructure, which provides the foundation for modern deployments.

Generally speaking, what types of assets in your organization are you most focused on/most interested in pentesting?

PERCENT OF RESPONDENTS, N=200, MULTIPLE RESPONSES ACCEPTED

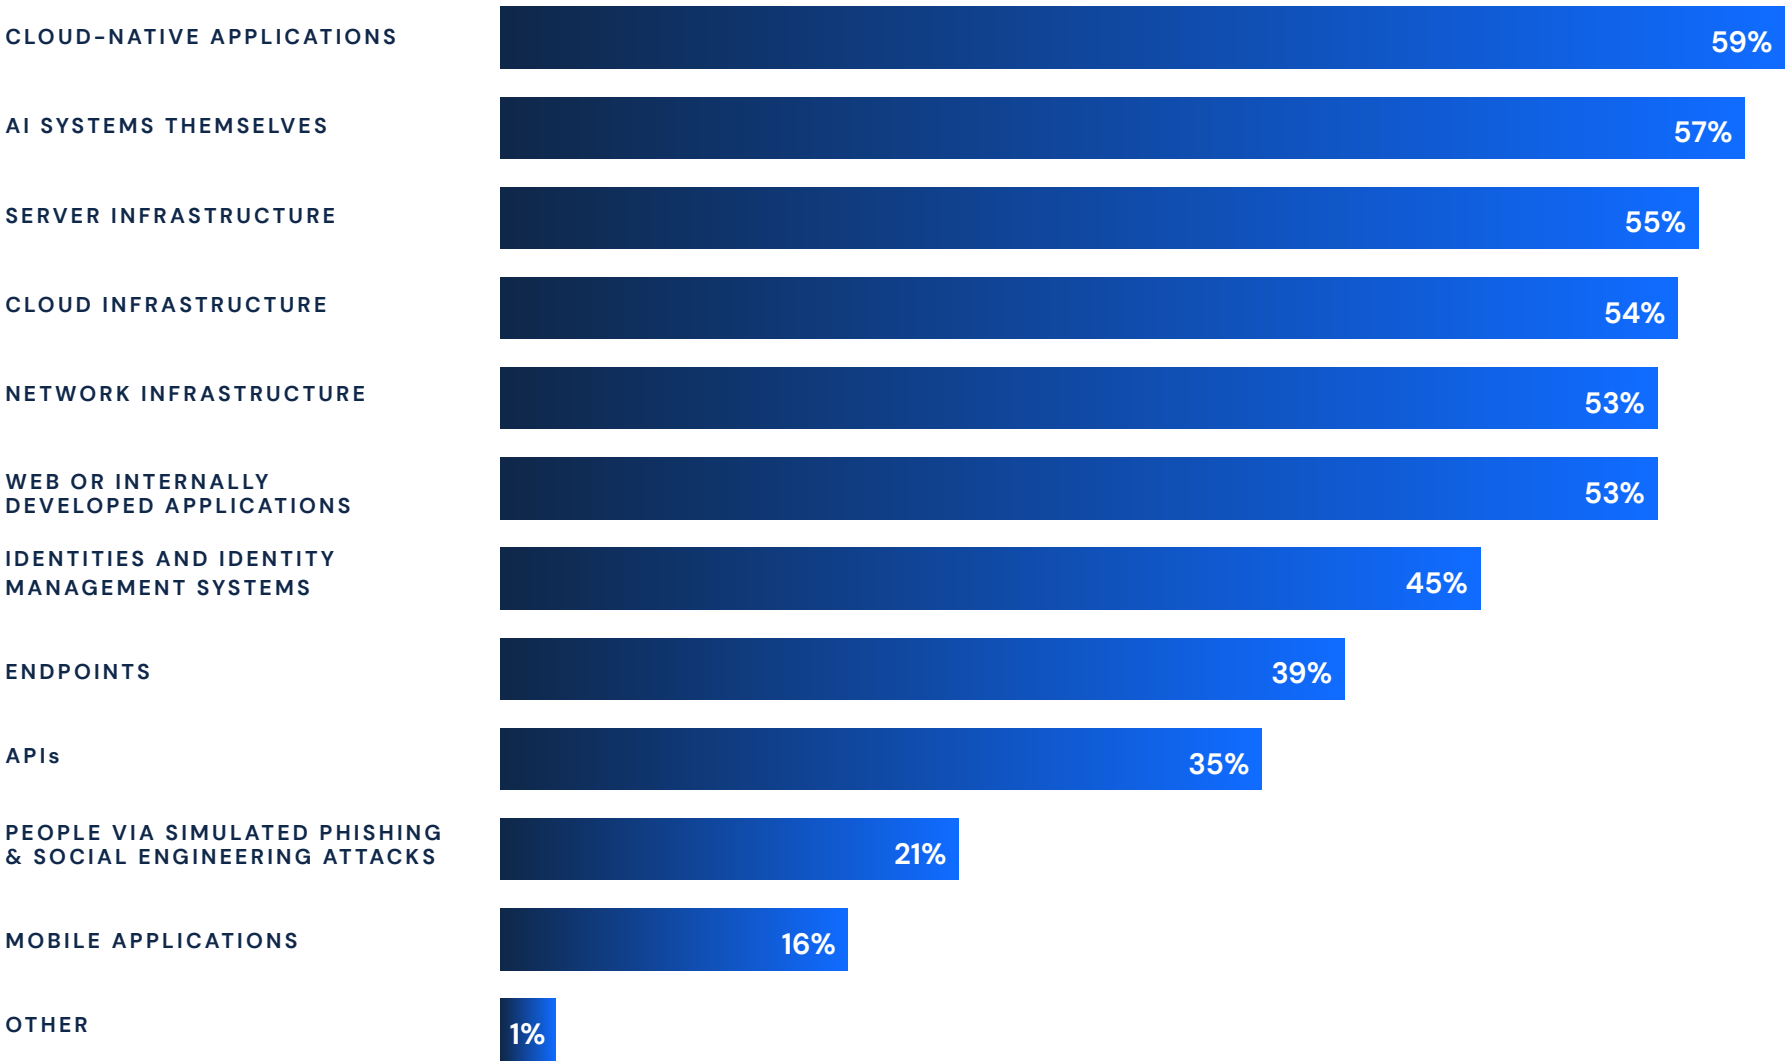

SOURCE: OMDIA

Synack

Proven results + guardrails are boosting overall trust in agentic AI

Early adopters have already established a high degree of trust in their agentic AI for pentesting programs. Of those surveyed, **87% reported high or complete trust** in agentic AI's ability to test their environments. This confidence stems from demonstrated ability to deliver on core promises of accuracy at scale. And trust is exceptionally strong among those who have moved beyond the pilot phase, with advanced users being **2.2X more likely to express complete reliance** on these systems.

While performance builds trust, rigorous governance maintains it. Early adopters recognize that autonomous systems cannot operate in a vacuum—they require a framework of safety and accountability.

Overall, **93% of organizations** indicate that comprehensive guardrails are either important or critical for the safe and effective operation of agentic AI.

TOP SAFETY MECHANISMS

Transparent AI Decision-Making

58% of organizations cite transparency in how the AI reasons and executes as the top factor for increasing trust.

Auditability and Explainability

31% demand guardrails that provide clear explainability for every action taken by the agent.

Automated Safety Stops

27% of leaders require boundary enforcement to ensure testing remains within defined rules of engagement.

Partnering for compliance

Trust is also increasingly tied to vendors providing pentesting programs. Organizations are specifically looking for providers that meet established industry requirements—such PCI-DSS, SOC 2, and ISO 27001—to ensure that the agentic systems integrated into their stack are as secure as the environments they are testing.

For **54%** of organizations, industry certifications and standards compliance directly increase their confidence in an agentic AI pentesting system.

Which of the following factors would increase your trust in agentic pentesting systems?

PERCENT OF RESPONDENTS, N=200, MULTIPLE RESPONSES ACCEPTED

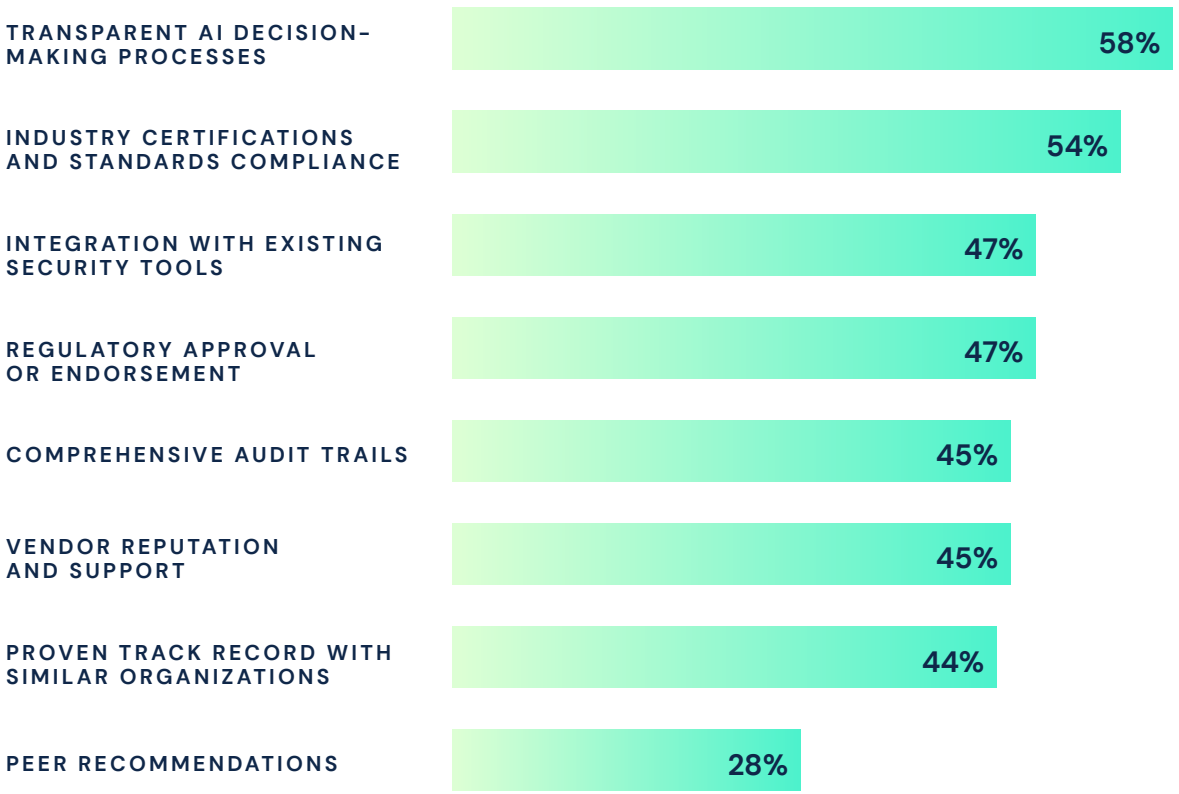

SOURCE: OMDIA

Synack

Humans are still a key component for successful agentic AI

What level of human involvement do you believe is necessary for agentic AI pentesting?

PERCENT OF RESPONDENTS, N=200

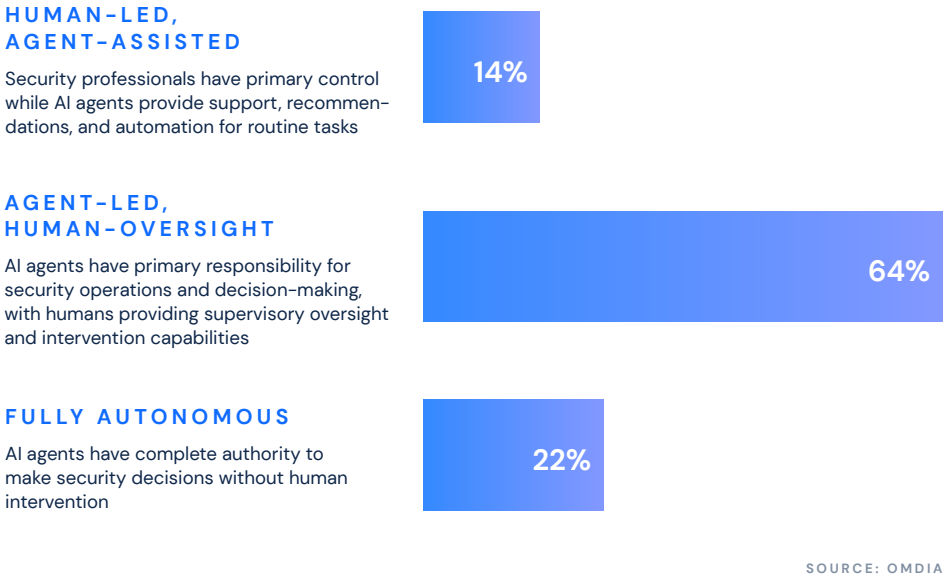

SOURCE: OMDIA

As organizations transition toward autonomous security, the consensus among early adopters is clear: The most effective strategy does not remove humans, but it elevates their role. In fact, 64% of organizations identify agent-led, human oversight as their preferred operational model. In this framework, AI agents carry the primary responsibility for execution and strategic reasoning, while human experts provide the critical supervisory safety net and intervention capabilities required for complex environments.

There is a growing recognition that human involvement is not merely a training phase for maturing technology, but a strategic necessity. Organizations already using agentic AI for pentesting are **1.6X more likely to view human oversight as a permanent requirement** for all testing scenarios, compared to those only in the pilot phase (**40% vs. 25%**).

While in the aggregate 69% of organizations see a possibility of eliminating human involvement at some point in the future, human-in-the-loop operating models are the gold standard today for maintaining safety and effectiveness.

Synack

Capabilities matter, with speed and efficacy ranking at the top

Early adopters are overwhelmingly prioritizing capabilities that provide immediate, actionable visibility into their environments. Real-time monitoring and alerting stands out as the most-cited capabilities at 32%. This preference reflects a strategic shift toward proactive security, where agentic AI can deliver instant insights into vulnerabilities and threats as they materialize.

For agentic AI to become an enterprise-ready asset, it must live within the existing security ecosystem. **Integration with existing security tools and SIEM platforms is a critical requirement for 31% of leaders.** This focus is driven by necessity, as integration challenges are cited as the No. 1 concern when implementing agentic systems for pentesting.

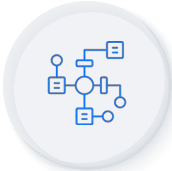

Integrations with SIEM Platforms

Security teams prioritize agentic AI that delivers seamless integrations across their existing security stack.

API-first design

Ensuring compatibility with DevSecOps workflows to keep pace with rapid development cycles.

Multi-environment support

The capability to span across cloud, on-premises, and hybrid environments seamlessly.

Automated prioritization

Features that not only find vulnerabilities but automatically triage and prioritize them based on business risk.

Which of the following capabilities does your organization want in agentic AI for pentesting offering?

PERCENT OF RESPONDENTS, N=200, UP TO THREE RESPONSES ACCEPTED

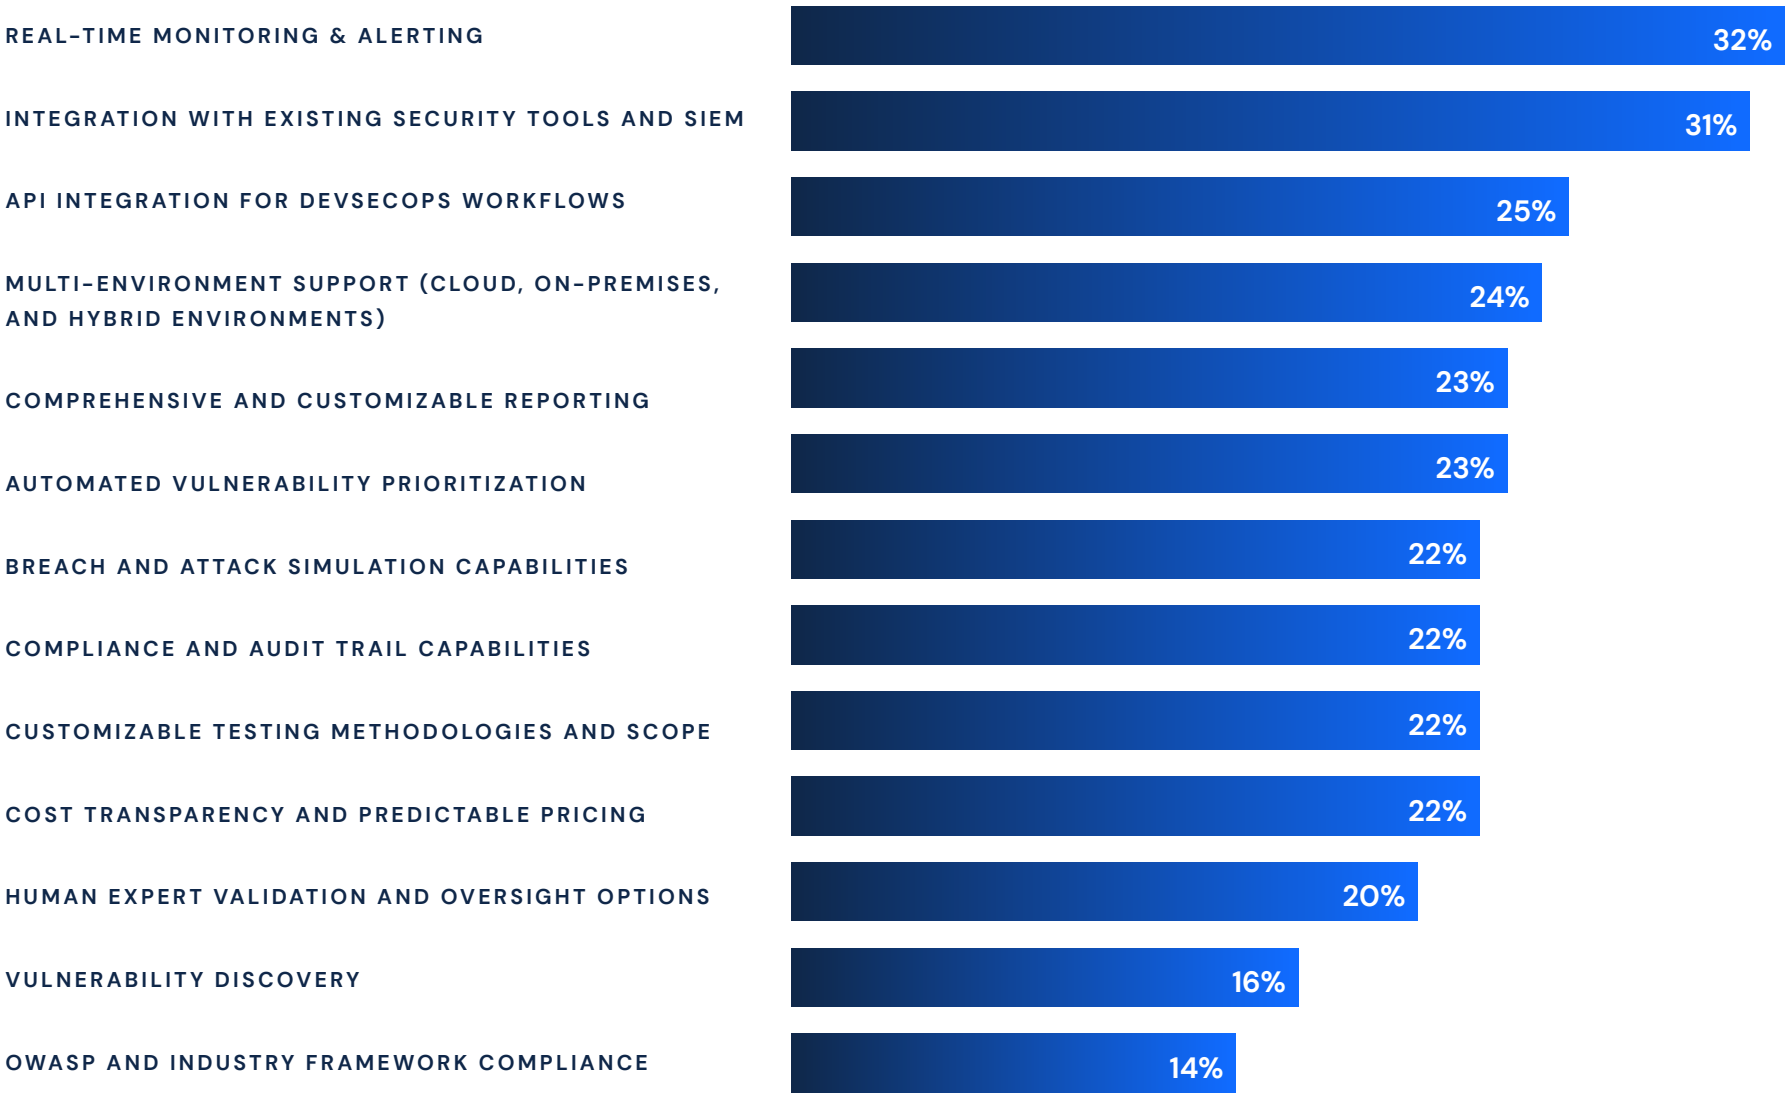

SOURCE: OMDIA

Synack

Becoming production-ready

For the organizations that are piloting agentic AI for pentesting, the most critical factor for getting their solution to be production ready is accuracy. Demonstrating proven accuracy was the most-cited factor—at 42%—for organizations to consider their AI-driven pentesting to be production ready for enterprise use.

But this is one of many considerations. Becoming production-ready involves a multi-faceted evaluation of how the solution fits into the broader enterprise ecosystem. Organizations are looking for a mature operational profile that includes aspects like integrations, certifications, and governance.

What would need to happen for you to consider agentic AI pentesting to be "production-ready" for enterprise use?

PERCENT OF RESPONDENTS, N=200, UP TO THREE RESPONSES ACCEPTED

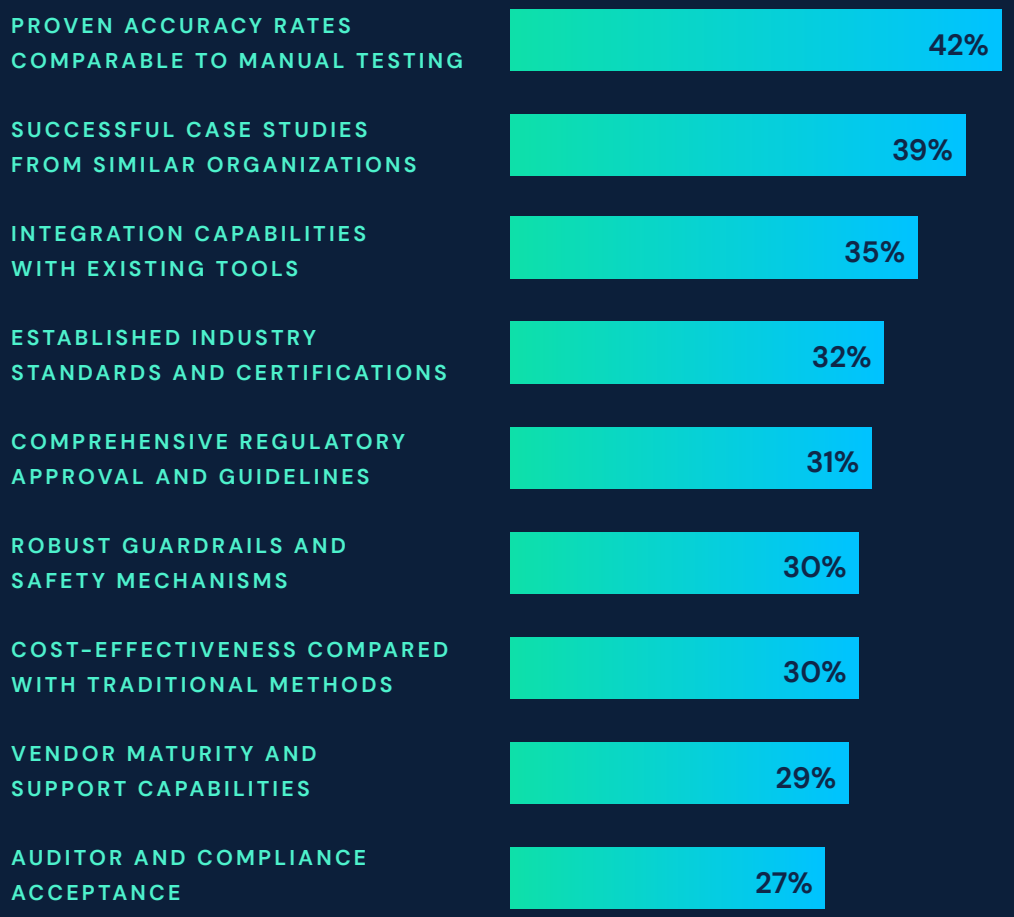

SOURCE: OMDIA

Overcoming final barriers to adoption

WHILE THESE EARLY ADOPTERS ARE SIGNALING HUGE MOMENTUM BEHIND AGENTIC AI IN PENTESTING, CONTINUED ADOPTION WILL DEPEND ON OVERCOMING SEVERAL HURDLES.

Risk Management

For organizations that have not yet adopted agentic AI in pentesting, security concerns regarding AI systems remain the top roadblock. Leaders are cautious about the reliability of autonomous decision-making and emphasize the need for transparency, explainability, and robust safeguards to build long-term trust.

Operational and Technical Challenges

Difficulties such as integration complexity with existing security infrastructure and a lack of internal expertise present significant friction points. Organizations are looking for implementation support, ease of use, and flexible automation that can adapt to their specific environments without causing system disruption.

Organizational and Financial Barriers

Budget constraints and the need for a clear business case are also holding back organizations. For agentic AI in pentesting to gain a foothold, a comprehensive ROI should prove both the tactical and strategic value of the shift.

What is holding your organization back from adopting agentic AI for pentesting?

PERCENT OF RESPONDENTS, N=172, MULTIPLE RESPONSES ACCEPTED

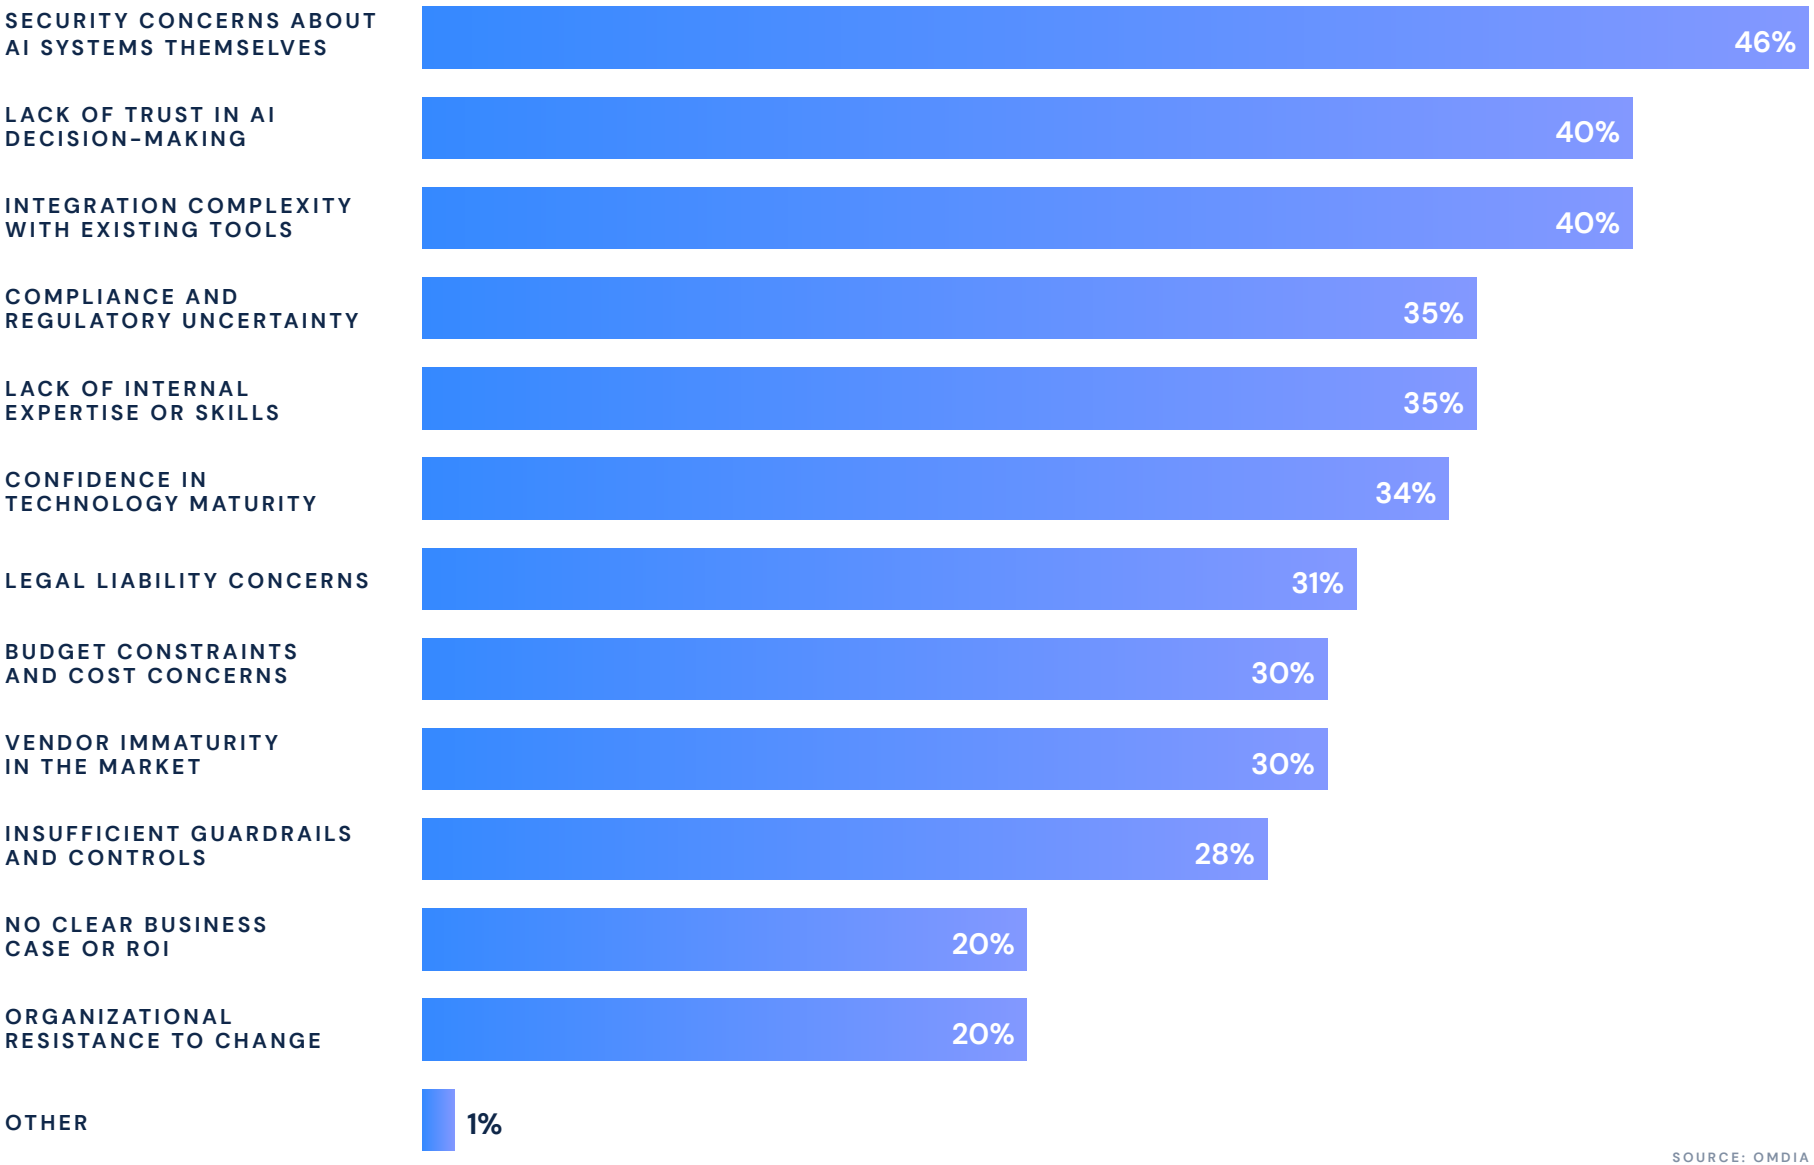

SOURCE: OMDIA

Synack

Research methodology

This whitepaper is based on a primary research study commissioned by Synack and conducted by Omdia. The findings reflect the current state of agentic AI adoption and its transformative impact on the penetration testing landscape.

Quantitative Survey Overview

The research was conducted via a quantitative web-based survey to gather actionable insights from a qualified group of cybersecurity professionals.

SAMPLE SIZE: N=200 qualified completes

GEOGRAPHY: US respondents

FIELD DATES: December 11–22, 2025

Respondent Profile

To ensure the data represented the decision-making pulse of the industry, the survey targeted specific roles and organization types.

JOB TITLES: Respondents included security leaders with manager titles or higher, as well as security practitioners. All participants hold responsibility for their organization's investments in AI tools for cybersecurity.

ORGANIZATION SIZE: All respondents represented enterprise organizations with 1,000 or more employees.

- 63% were from organizations with 1,000 to 4,999 employees.
- 38% were from organizations with 5,000 or more employees.

About Synack

Synack is the leader in human-led and AI-powered penetration testing, transforming offensive security to help organizations proactively reduce risk, stay compliant and defend against evolving cyber threats. Synack harnesses agentic AI innovations and a talented, vetted community of security researchers to deliver continuous penetration testing and autonomous vulnerability management. Founded by former NSA operatives, Synack has enabled nearly 10 million hours of expert testing to protect critical assets, from global financial systems to U.S. Defense Department networks.

Learn more about Synack and find out how to use agentic AI as a force multiplier for penetration testing.