McKinsey & Company

Risk & Resilience Practice

Securing the agentic enterprise: Opportunities for cybersecurity providers

The rise of agentic AI is exposing enterprises to new cybersecurity risks, creating extensive opportunities for providers that can make the world of autonomous systems safer.

This article is a collaborative effort by Charlie Lewis, Jeffrey Caso, and Marc Sorel, with Alexander Keegan and Julian Fuchs-Souchon, representing views from McKinsey's Risk & Resilience Practice.

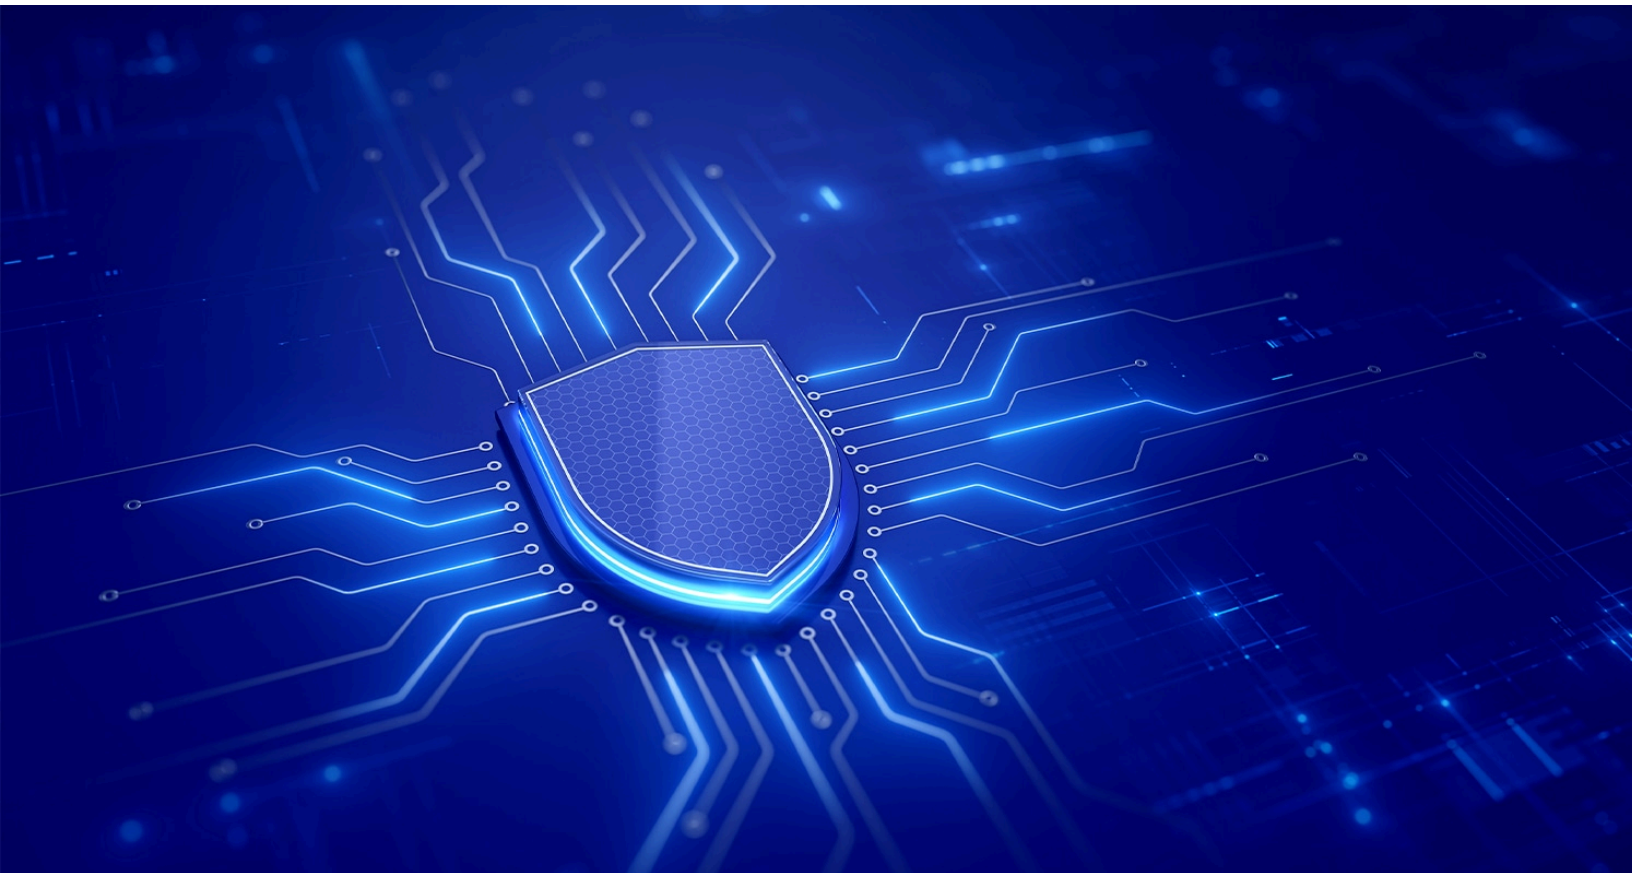

March 2026

The $220 billion cybersecurity market is projected to grow at roughly 13 percent CAGR over the next several years.[1] But unlike prior waves of growth fueled by entirely new product categories (for example, the convergence of intrusion prevention and detection into end point detection and response, and secure access service edge), the next phase will be defined by a reframing of existing ones. Identity, detection, and security operations are not being replaced—they are being rebuilt to absorb AI capabilities and govern autonomous systems.

Enterprise AI has entered the agentic phase, in which autonomous systems independently perform complex tasks at machine speed. After a wave of pilots in AI-assisted workflows, organizations are beginning to deploy agents across their infrastructure, identity, engineering, and security environments. Over the next 12 months, companies expect the share of fully implemented agentic AI solutions to more than double.[2]

As enterprises embrace AI agents, the surface for potential cyberattacks expands. For example, consider a financial operations agent authorized to reconcile accounts and generate quarterly reports. If the agent accesses manipulated inputs or misaligned policy updates, it could access restricted financial forecasts, override approval workflows, or distribute confidential information externally. This fundamentally changes organizations' risk profile because agents can autonomously decide which data and environments to access to accomplish their tasks. Such autonomy can greatly boost productivity—but it also heightens risk if an agent's actions run afoul of enterprise risk controls.

This environment creates a new mandate for cybersecurity providers to address a material shift in the control plane. The dawn of the agentic age is reallocating customer budgets, creating new unsolved opportunities (such as securing nonhuman identities at scale), and redefining competitive advantage. With AI agents likely to increasingly replace human analysts over the coming years, and large language models [LLMs] creating risks of data leakage, the shift is also expanding the total addressable market.

This article considers how agentic adoption is reshaping control planes and attack surfaces, how chief information security officer (CISO) mandates and budgets are changing as a result, the opportunities these shifts create for providers, and how providers can capture them.

How agentic AI is reshaping the enterprise control plane

The agentic transition is firmly underway. In a recent McKinsey survey of cybersecurity solutions buyers,[3] respondents say they expect the adoption of AI agents to double in the next three years. Additionally, 35 percent of respondents anticipate that AI agents will replace their tier-one security operations center (SOC) analysts, and nearly 50 percent expect AI to be embedded across the cyber stack within the same period.

Together, these responses point to an acceleration in the shift from human-scale to machine-scale security for tasks previously owned exclusively by humans, such as prioritizing and

[1] McKinsey Cyber Practice, McKinsey Cyber Market Map 2025.
[2] "Winning B2B customers in technology and telecommunications," McKinsey, February 26, 2026.
[3] McKinsey Cybersecurity Customer Survey, September 2025, n = 104. Respondents by geography: approximately 56 percent from North America and about 44 percent from the European Union. By size: all enterprises (that is, more than 5,000 employees and more than $1 billion revenue). By industry: financial services, 54 percent; software 46 percent; all confirmed decision-makers for their organization on which cybersecurity solutions to purchase (for example, 60 percent CISO or CISO direct report, 40 percent adjacent security leader with decision-making responsibility for security solutions).

remediating security alerts, provisioning and configuring cloud infrastructure, managing identity and access workflows, reviewing and deploying code, and even executing financial and operational processes. More than ever, controls will need to be automated, continuous, and embedded into run time to mitigate the risks posed.

This context will change how CISOs and cybersecurity solution providers approach system and data protection. Historically, cybersecurity programs concentrated on three questions: who has access, whether systems are hardened, and whether threats are detected and contained. The focus was on defending systems operated by humans and enforcing access controls at the perimeter, usually at the end point.

With the rise of software as a service (SaaS), cloud, and remote work over the past decade, identity became an even more critical control plane for the security stack than it was in the days of employees largely working on premises. The rise of AI agents and other nonhuman identities also means cybersecurity programs will need to consider more questions, as risk increasingly sits downstream of authentication—within multistep workflows executed autonomously. Agentic environments thus make it more important—and more challenging—to continuously monitor activity at network, data, device, and identity levels.

The CISO's mandate is no longer merely to protect the technological infrastructure. Function leaders, in collaboration with partners across the enterprise, need to ensure that humans and agents operate safely, predictably, and within policy constraints. They need to govern how agents act after access is granted, manage rapidly growing nonhuman identities, monitor decisions that agents make in real time, and ensure machine actions remain traceable and auditable despite their rapid spin-up and spin-down to execute.

This expanded mandate could change how CISOs invest their budgets. Our survey, expert interviews, and other research point to three trends. First, overall budgets are projected to grow at a CAGR of around 2.5 percent over the next three years, with the share spent on AI solutions expected to more than triple to 15 percent, from about 4 percent (Exhibit 1).

Second, the spread of AI agents will likely lead to greater spending on controls to govern autonomous systems, particularly around identity (registering agents, for example, and monitoring communications between LLMs and agents in real time). This shift will not eliminate traditional cybersecurity categories but instead reallocate spending within them, as identity, detection, and security operations evolve to govern machine activity.

Lastly, budgets are changing because risks—and the ownership of those risks—are shifting. As the CISO budget grows, so too in parallel is the expectation that other business functions will own more of the risk and embed its management into their day-to-day operations. The result is increased spending on security outside of the CISO function.

Exhibit 1

Agentic AI spend could rise to 15 percent of enterprise cybersecurity budgets in next three years, concentrated in identity, governance, and data.

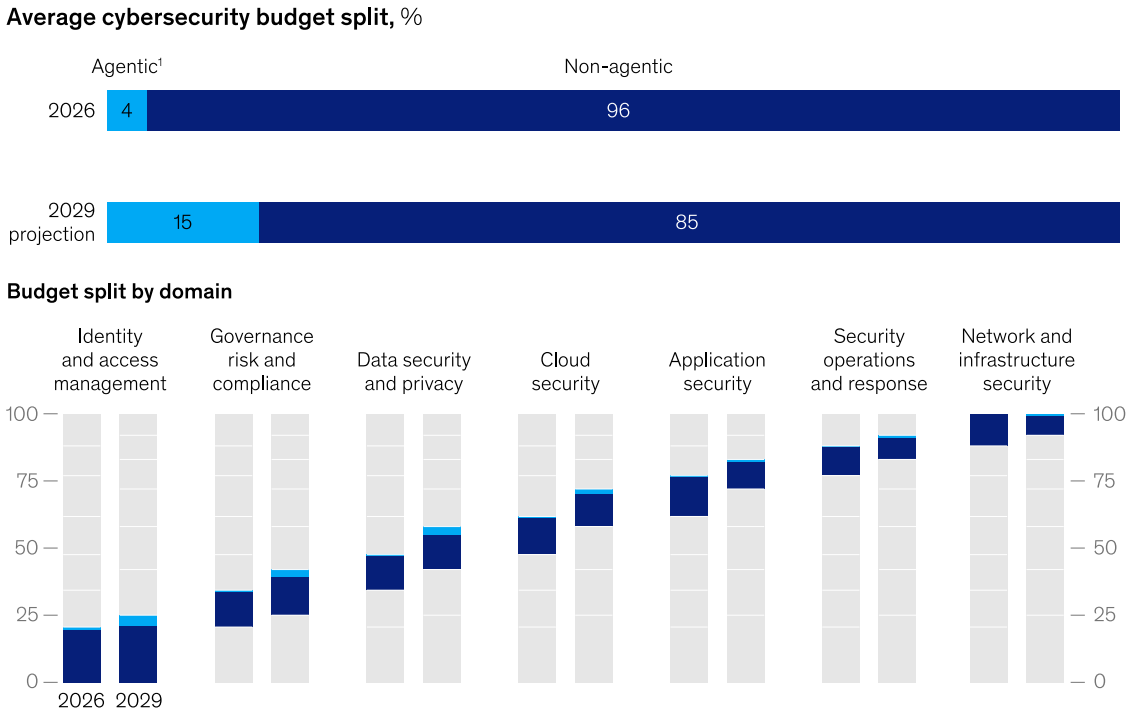

[1]Technologies purpose built to manage AI agents.

McKinsey & Company

Implications for providers and investors

With cyber risk shifting toward autonomous systems—and enterprise budgets following—two questions matter for providers and investors: Where will value concentrate within the $220 billion cybersecurity market in the coming years? And which platforms are positioned to capture it?

The emerging opportunity pools

The value for cybersecurity providers lies in incorporating more AI features and addressing enterprises' growing agentic-related security problems across identity systems and detection and security operations. In particular, three challenges stand out: identity and access management (IAM) architectures, detection, and automation of security operations.

Restructuring IAM architectures to match the pace and volume of access requirements. Traditional IAM architectures were designed to accommodate human users with predictable access patterns. Agentic systems introduce short-lived, task-specific machine identities that can be created and destroyed programmatically—often within seconds to minutes. Agents increasingly gain access dynamically as they execute workflows rather than having access assigned through static roles as part of typical IAM processes. . This creates new risks such as autonomous privilege escalation, in which agents can accumulate unintended permissions due to policy drift, tool chaining, or malicious manipulation.

In such cases, static access models—in which permissions are granted infrequently and reviewed periodically—are likely to be insufficient. Identity will evolve toward continuous authorization for nonhuman entities, granting access for a period of time or in a specific context, with the agent's behavior monitored. This will likely spur demand for solutions addressing machine identity governance, identity threat detection and response (ITDR) for nonhuman entities, and ephemeral credential life cycle management (for example, where agents are given temporary access tokens or passwords that automatically expire after a predetermined time).

Expanding detection from alerts into run time enforcement. In traditional environments, authentication and end point visibility were the primary controls. In agent-driven systems, risk shifts to what happens after agents are granted access. Not surprisingly, more than 75 percent of respondents in our survey report seeking protection against AI input manipulation (Exhibit 2).

However, current run time controls for agent behavior lack sufficient sophistication. Roughly 30 percent of survey respondents say they are not confident in their existing vendors' ability to address these AI risks. At the same time, CISOs do not anticipate closing these gaps through custom-built solutions; they want vendors to provide them. This creates an opportunity for providers to expand detection and response solutions beyond alerts to run time enforcement—similar to how a credit card provider can automatically decline a transaction if it exceeds a spending limit or is conducted in a suspicious location. Solutions could address the following five areas:

— *agent run time protection:* real-time monitoring and policy enforcement while agents execute tasks

— *behavioral anomaly detection:* recognizing when agents act outside expected patterns, permissions, or workflows

— *agent activity logging and traceability:* end-to-end visibility into actions, tool calls, and decision chains for auditability

— *decision integrity safeguards:* controls that keep AI-driven actions policy-compliant and resistant to manipulation

— *AI-native data loss prevention:* averting sensitive data exposure through prompts, outputs, and downstream integrations

Exhibit 2

Organizations are seeking to increase the types of capabilities they adopt for AI security despite additional cost.

Share of respondents highly likely to adopt AI security capability,[1] %

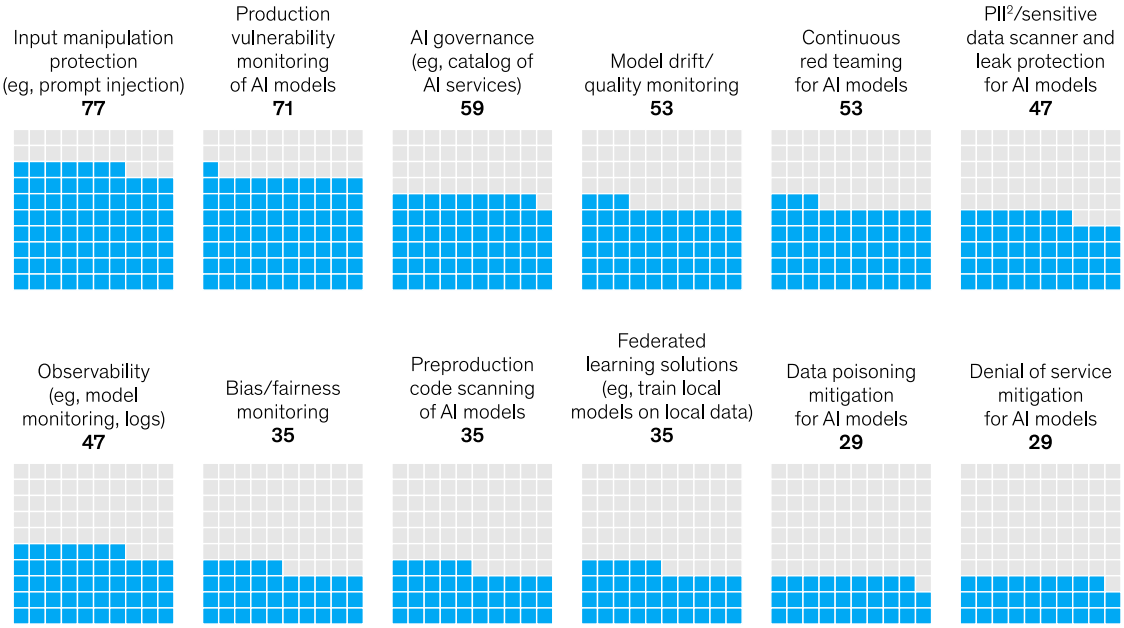

[1]Question: On a scale of 1–10, how likely would your organization be to invest in developing AI capabilities through new/existing cybersecurity vendors within the next three years (assuming the new capability gives you 20–30% more automation or better efficacy).
[2]Personally identifiable information.
Source: McKinsey Cyber Market Customer Survey, Sept 19–29, 2025 (n = 77)

McKinsey & Company

Automating security operations. The security stack of the future is likely to be human supervised but machine operated. Approximately 80 percent of respondents to our survey expect AI and machine learning capabilities to be moderately to highly integrated within the next three years (Exhibit 3). This could increase demand for automated SOCs, with AI-native detection and response platforms, secure remediation orchestration, and agent supervision

dashboards. Automated SOCs may also alter cost structures across the industry, favoring platforms that can deliver AI-supervised responses with less human involvement while maintaining control and auditability. For large enterprises, such solutions can yield tens of millions of dollars in annualized run-rate savings while improving outcomes on key operational metrics, such as mean time to remediation.

Exhibit 3

Organizations expect to embed AI and machine learning capabilities across the security stack over the next three years.

Expected level of AI/machine learning integration across cybersecurity stack in the next 3 years,[1]
% of respondents

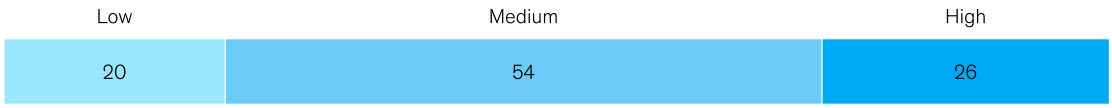

[1]Question: Over the next 3 years, what percentage of your organization's cybersecurity software stack do you expect will have AI or machine learning capabilities embedded by default (either natively in the product or via an integrated AI layer)? Respondents who answered "not sure" not shown.
Source: McKinsey Cyber Market Customer Survey, Sept 19–29, 2025 (n = 76)

McKinsey & Company

How providers can tap new market opportunities

To lead the cybersecurity market in the agentic age, solution providers may need to redesign their platforms around agent-native governance. Such an architectural reset means creating systems that continuously authorize, monitor, and constrain autonomous agents throughout their lifecycle rather than relying on periodic access reviews or static role assignments. Our projection of market segment growth reflects the central role of AI security in fueling the industry's future (Exhibit 4).

To lead this next phase of growth, cybersecurity providers—whether hardware, software, service, or hybrid—can strengthen their capabilities in five core areas:

Embed in nondiscretionary, mission-critical workflows. Compliance-focused services and software providers may have the lead in this area, given the mandatory nature of compliance work (as in solutions that help with the completion of SOC 1 or SOC 2 audits). At the same time, network firewall providers, for example, could gain an edge by adding AI-enabled firewall capabilities that inspect traffic for compliance in real time, reducing compliance assessment times upstream from the process of reporting to regulators.

Exhibit 4

Agentic adoption could drive strong growth across a number of segments in the approximately $220 billion cybersecurity market.

Cybersecurity segment market maturity, growth rate, and market size[1]

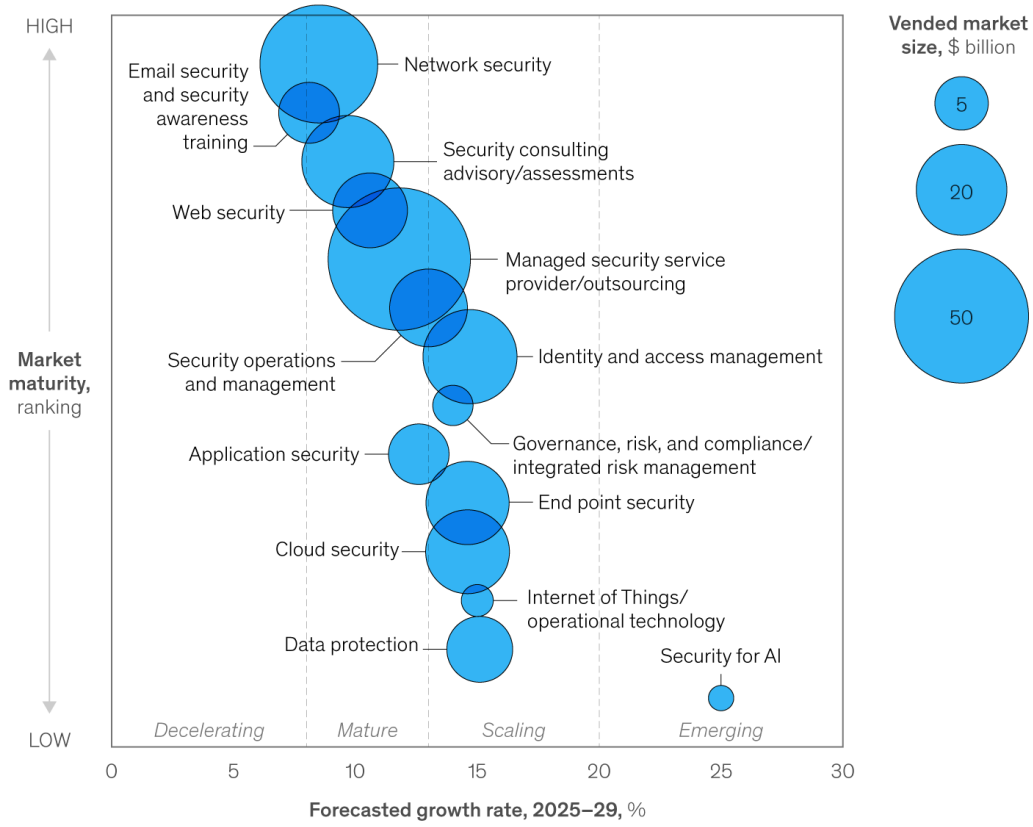

[1]Excludes hardware security modules, cyber deception/honeypots and mobile device protection.
Source: Synergy Research Group; McKinsey Cyber Market Map 2025

McKinsey & Company

Enhance offerings so they can operate as systems of record. Rather than building isolated AI security tools, providers can embed run time enforcement in their core platforms. By integrating runtime controls into IAM, security information and event management, end point detection and response, cloud security, and data loss prevention platforms, behavioral monitoring and decision

safeguards become embedded capabilities rather than bolt-on services or features. This could help create a central product offering in the security stack and the broader enterprise.

Strengthen data insights. Providers can collect or generate proprietary data to develop insights into anomalous behavior by agents and ways to remediate it. They can look across collaboration environments and internal and external sources. For example, they could learn the methods of incident responders who analyze threat actors to develop appropriate controls and solutions.

Rationalize the security stack. Simply replacing one security stack solution with a lower-cost alternative is not enough. Providers that develop functionality that can power multiple tools can gain an edge by displacing or rationalizing two or more security stack solutions, in the same way that extended detection and response merged end point detection and response, network detection and response, and some cloud security tools into one solution.

Look beyond the traditional CISO buyer. As AI deployments and cybersecurity budgets move into engineering and product organizations, providers can embed security controls directly into development workflows, AI pipelines, and infrastructure automation systems. Those focused solely on traditional CISO procurement may miss the fastest-growing budget pools.

The shift toward agentic AI will expand and reallocate cybersecurity budgets toward platforms capable of governing autonomous systems. Over the next three to four years, agentic adoption will spur new control requirements and boost security budgets within the CISO organization and beyond, as other enterprise leaders assume greater responsibility for risk management. Future cybersecurity industry leaders will be those that can address enterprises' unmet needs by redefining the core categories of cybersecurity and their control points through added functionality.

Charlie Lewis is a partner in McKinsey's Connecticut office; **Jeffrey Caso** is a partner in the Washington, DC, office, where **Alexander Keegan** is a specialist; **Marc Sorel** is a partner in the Boston office; and **Julian Fuchs-Souchon** is a knowledge expert in the Stuttgart office.

The article was edited by Arshiya Khullar, a senior editor in the Gurugram office.