[real]is
Solutions Group

Machine Speed, Human Judgment

How AI Is Reshaping the Security Services Market

An Executive Strategy Paper

July 2026

Mike Reed, M.S.
Chief Operating Officer, Realis Solutions Group

Executive Summary

For more than a decade, the managed security market evolved by adding capabilities. Managed firewalls became Managed Security Service Providers (MSSPs). MSSPs evolved into Managed Detection and Response (MDR). MDR providers expanded into exposure management, identity protection, cloud security, and incident response.

Artificial intelligence changes that trajectory.

The next evolution is not simply another category of security product or managed service. AI is changing what organizations expect from their security partners and, more fundamentally, changing what they are buying. In the World Economic Forum's Global Cybersecurity Outlook 2026, produced with Accenture and now in its fifth edition, 94% of survey respondents identified AI as the most significant driver of change in cybersecurity for the year ahead, and 87% identified AI-related vulnerabilities as the fastest-growing cyber risk over the course of 2025.

Historically, organizations evaluated providers on operational capability: who can monitor my environment, who can investigate alerts, who can respond to incidents, who has the best analysts. Those questions remain important, but they are no longer sufficient. As organizations adopt AI across engineering, operations, customer service, finance, and executive decision-making, security becomes less about operating tools and increasingly about enabling the business to use AI with confidence.

The result is a market organizing itself around four distinct customer buying models. These models are not a ladder that every organization climbs in sequence. They coexist, often within the same enterprise, and each reflects a different problem the customer is trying to solve and a different definition of value. What is changing is where the spending gravity sits. In its July 2025 forecast, Gartner projected worldwide end-user spending on information security would reach roughly $240 billion in 2026, an increase of 12.5% over 2025. But the composition of that spending is shifting: budget, board attention, and margin are moving from the first two models, where AI is compressing the cost of labor-intensive operations, toward the second two, where AI is creating new demand for judgment, governance, and transformation. One caution travels with

that claim throughout this paper: the compression depends on AI's own input economics, which are volatile. The cost relationship between AI-delivered and human-delivered security work is best treated as a floating exchange rate to be managed through vendor portability and contractual cost discipline rather than as a permanent discount.

This paper makes three arguments and adds one discipline. First, the four buying models are a more useful map of the security services market than traditional categories such as MSSP, MDR, or XDR, which describe delivery mechanics rather than customer value. Second, the buyer changes at each model, which means incumbency does not transfer: the provider that wins the SOC contract is not automatically positioned to win the governance engagement. Third, the providers that lead the next decade will not be those with the most analysts or the largest SOCs. They will be those that can help executive leadership navigate AI as both a strategic opportunity and an enterprise risk, turning security from a cost center into a catalyst for trusted innovation. The discipline is this: the industry has been through an automation hype cycle before, with SOAR, and that history counsels that the direction of today's predictions is more reliable than their speed. This paper treats analyst forecasts accordingly, as directional signals rather than schedules.

The Four Buying Models

Today's security market is increasingly organized around four customer buying models rather than four technology categories. Each model answers a different question, is purchased by a different buyer, and measures success differently. The largest firms increasingly span more than one model, so the examples below should be read as representative of where each firm's center of gravity sits, not as exclusive classifications. A note on terms: throughout this paper, vendor means a firm selling software or a platform, provider means a firm selling an operated or delivered service, and firm or partner is used where a claim applies to both. Exhibit 1 summarizes the framework.

Exhibit 1. The four buying models. Spending gravity shifts toward judgment, governance, and transformation.

Model 1: Platformers

"I need one platform." These organizations are primarily purchasing technology. Their objective is simplification: consolidating multiple security products into an integrated platform that reduces operational complexity while improving visibility and automation. Representative vendors include CrowdStrike, Microsoft, Palo Alto Networks, Sophos, and Zscaler, though each of these firms also sells services and several are pushing aggressively into the operational models described below. The consolidation preference is real and durable, but it is a gradual buying behavior, not an imminent event; enterprises rationalize platforms over multiple budget cycles, and the history of such predictions argues against building urgency around any particular consolidation deadline.

In Model 1, the AI conversation is operational. Customers want to know whether AI can reduce alert volume, investigate incidents, automate response, and improve analyst productivity, and how quickly those gains will show up. The evidence that it can is now substantial. IBM's 2025 Cost of a Data Breach report, conducted by the Ponemon Institute across 600 breached organizations worldwide, found that organizations using AI and automation extensively throughout their security operations saved an average of $1.9 million in breach costs and reduced the breach lifecycle by an average of 80 days (Exhibit 2). For context, the global average breach lifecycle fell to 241 days that year, a record low. Success in this model is measured through exactly that kind of operational efficiency: lower security operating costs, faster incident response, better utilization of existing technology investments, and reduced analyst workload. At this stage, customers are purchasing AI embedded within security products.

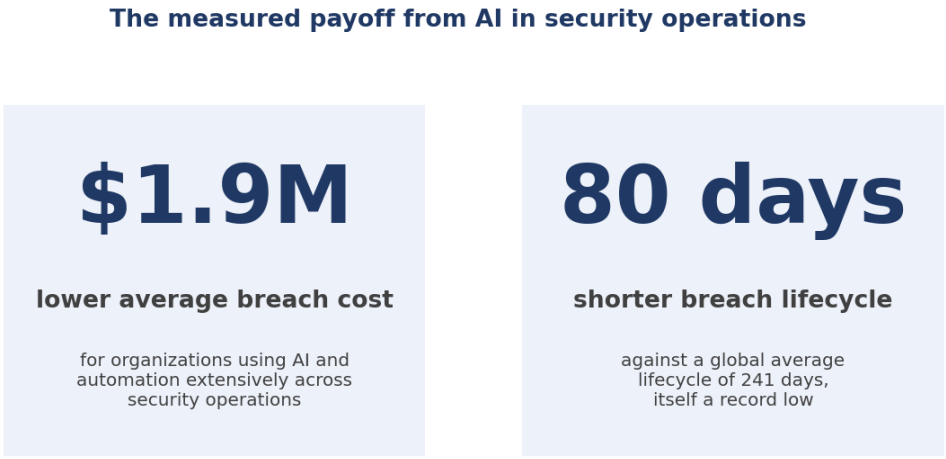

Exhibit 2. Measured effect of extensive security AI and automation. Source: IBM Cost of a Data Breach Report 2025 (Ponemon Institute, 600 organizations).

Model 2: Operators

"I need someone to run my SOC." These organizations recognize that technology alone does not create effective security operations. Instead of purchasing software, they purchase expertise delivered through a managed Security Operations Center. Representative providers include Arctic Wolf, Critical Start, Expel, eSentire, and Proficio. The demand base remains strong: Mordor Intelligence values the global MDR market at approximately $4.2 billion in 2025 and projects it to reach $11.3 billion by 2030, a compound growth rate near 22% (Exhibit 3), and other analyst houses publish forecasts in a similar range.

It is worth being precise about who these customers are, because two very different populations buy in this model. The first has a security team and wants to extend it. The second, common in the midmarket and in regulated smaller enterprises, has no SOC at all and never will: for them, the outsourced provider is not augmenting an internal capability; it is the security operations capability. For this second population, the industry's debate about whether AI can replace SOC analysts is largely beside the point. They are buying an outcome and an accountable partner; how the provider blends analysts, automation, and AI behind the scenes is the provider's problem. That structural fact insulates a large share of Model 2 demand from the "AI displaces the SOC" narrative.

Here, AI becomes an operational multiplier rather than simply a product feature. Customers evaluate how providers use AI to investigate alerts autonomously, correlate evidence across environments, accelerate threat hunting, improve detection engineering, execute response playbooks, and raise analyst productivity. The objective is no longer reducing alerts; it is improving operational outcomes: lower mean time to detect and respond, higher-quality investigations, and broader security coverage without proportional headcount growth. Customers are purchasing AI-enabled security operations.

This is also the model under the most direct pressure from AI, because MDR economics were built on labor arbitrage: a shared bench of analysts spread across many customers. As AI absorbs Tier-1 triage and portions of Tier-2 investigation, the labor component of that value proposition shrinks. But the ground truth from practitioners suggests the displacement is slower and messier than vendor narratives imply. The SANS 2025 SOC Survey, authored by Christopher Crowley, puts numbers on the gap: AI and machine learning are a defined part of operations in only 14% of SOCs, while a further 40% report staff using the tools with no defined workflow around them, and AI and machine learning tools sit at the bottom of the technology satisfaction rankings. Meanwhile 69% of SOCs still report their metrics through manual or mostly manual processes. Providers in this model must differentiate on the quality of their AI-supervised operations or migrate toward the judgment-based models below, but they have more time to do it than the hype suggests.

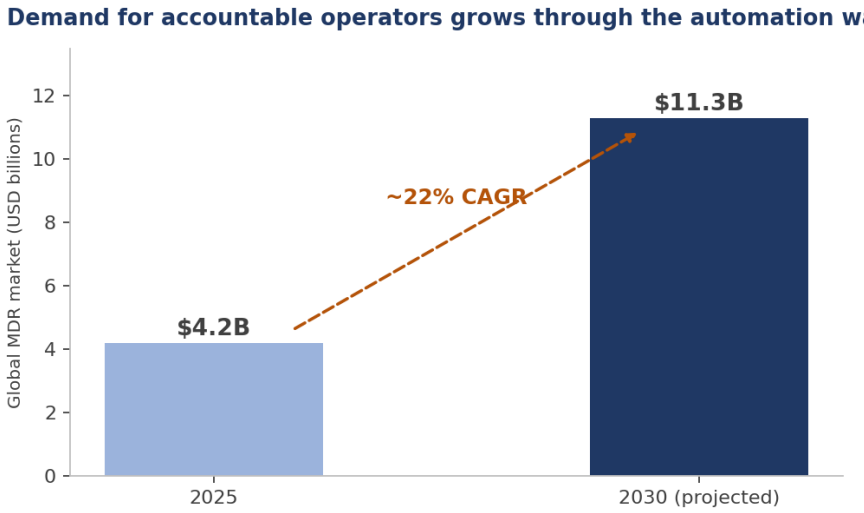

Exhibit 3. Global MDR market growth. Source: Mordor Intelligence, 2025.

Model 3: Delegators

"I need someone to run my security program." As organizations mature, operational excellence becomes an assumption rather than a differentiator, and leadership begins asking fundamentally different questions. Instead of asking whether AI can investigate an alert, executives ask which AI initiatives to approve, what governance model to implement, which regulations apply, how to evaluate AI vendors, how autonomous agents should be governed, and how AI risk should be communicated to the board. This model also includes organizations that outsource the security program itself: companies without a CISO, or with a part-time one, that purchase program leadership, virtual CISO services, and governance as a package alongside operations.

The demand signals here are unambiguous. The WEF found that the share of organizations assessing the security of AI tools before deployment nearly doubled in a single year, from 37% to 64%, while roughly a third still have no such process at all: a governance market being built in real time. Deloitte's Global Boardroom Program survey of nearly 700 directors and executives found that 66% say their boards do not know enough about AI, and 40% are rethinking board composition because of it. Gartner expects spending on AI governance platforms to reach $492 million in 2026 and to exceed $1 billion by 2030, driven by regulatory proliferation.

Representative providers include advisory-led organizations such as Optiv and GuidePoint Security, along with a growing set of firms combining managed operations with program leadership. In this model, the value of AI shifts from automation toward judgment. Customers are purchasing partners capable of helping leadership make better decisions about AI adoption, governance, security investment, regulatory readiness, and organizational risk. Success is measured through reduced enterprise risk, stronger governance, better

investment decisions, executive confidence, and sustainable AI adoption. Customers are purchasing AI-enabled executive decision support.

This model contains two distinct products that have historically been sold together. The first is program operation: the fractional CISO engagement, the program leadership, the staffing of roles the customer cannot fill internally. The second is program judgment: deciding which AI initiatives to approve, which governance model fits, what the regulatory exposure actually is, and whether the money being spent is producing the risk reduction it claims. Advisory-led firms typically sell both, which has made the distinction easy to overlook. It is becoming harder to overlook, because a growing number of customers are buying the second without the first.

The reason is structural and trust-motivated rather than skill-based. Any provider assessing whether its own service is worth the customer's money holds a position in the answer. That has always been true, and it has always been managed through procurement discipline and competitive tension. AI weakens both. As more of the delivery moves inside the provider's models and agents, the customer's ability to independently evaluate what it is receiving degrades, and it degrades at precisely the moment spending is rising and boards are asking harder questions. The same dynamic applies to a platform vendor assessing whether its embedded agents reduce enterprise risk, and to any provider scoping its own renewal. The answer is not better operators. Operators are not the problem, and the operational market is growing. The answer is that judgment about the operation increasingly needs to sit somewhere with no position in it.

This separation also tracks the buying center. Program operation is purchased by the security organization and judged on security outcomes. Program judgment is purchased by the CISO alongside the CFO, the general counsel, and the board committee, and it is judged against enterprise outcomes: whether AI initiatives cleared governance faster, whether regulatory exposure is defensible, whether the spend produced measurable risk reduction, and whether the business moved at the speed it needed to move. Those are different questions, answered for different people, on different timelines.

Model 4: Transformers

"I need business transformation." The most mature buying model treats AI not as a security initiative but as an enterprise transformation initiative. Security becomes one capability supporting a broader objective: enabling trustworthy innovation across the business. Representative providers include Accenture, Deloitte, IBM Consulting, Kyndryl, and other global transformation partners.

The size of this opportunity is defined by a gap. McKinsey's 2025 State of AI survey found that 88% of organizations now use AI in at least one business function, up from 78% the prior year, and nearly 80% use generative AI specifically. Yet only 39% attribute any EBIT impact to AI at all, just 7% report it fully scaled across the enterprise, and roughly 6% qualify as high performers attributing more than 5% of EBIT to AI. Accenture's State of Cybersecurity Resilience 2025, drawing on 2,286 CISOs and CIOs at organizations above $1 billion in revenue across 17 countries, found that only 10% of organizations sit in what it calls the Reinvention Ready Zone, while 63% remain in an Exposed Zone lacking both cohesive strategy and technical capability, and 77% lack the essential data and AI security practices needed to protect critical business models, data pipelines, and cloud infrastructure. Nearly every enterprise is adopting AI; almost none has yet converted adoption into governed, scaled value (Exhibit 4). Closing that gap is precisely what Model 4 customers are buying.

Executive discussions in this model shift away from operational metrics and toward business outcomes. Which AI investments will create competitive advantage? How do we accelerate adoption without increasing enterprise risk? How should operating models evolve? How should governance scale with autonomous systems? How do we measure AI return on investment? At this stage, AI is not primarily reducing cost; it is creating enterprise value. Customers are purchasing AI-enabled business transformation.

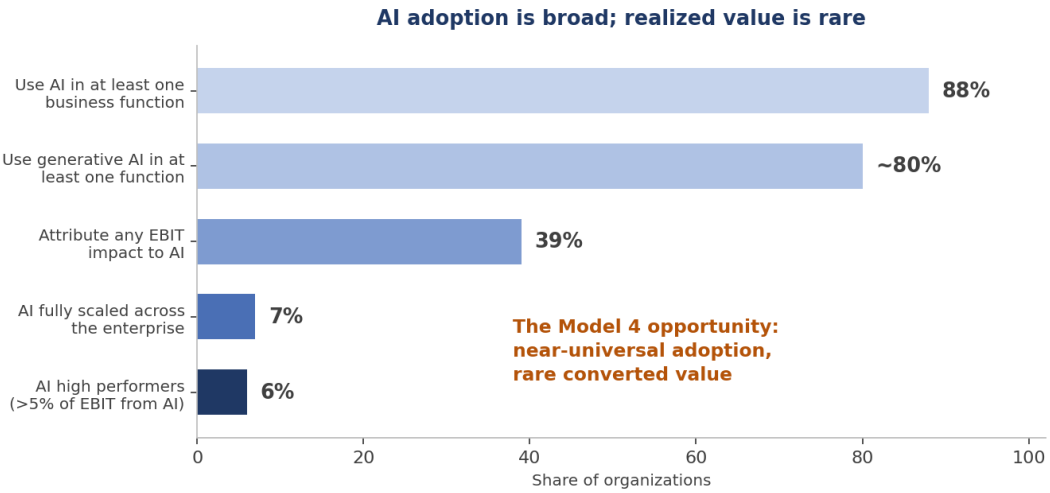

Exhibit 4. AI adoption versus realized value. Source: McKinsey, The State of AI, 2025 global survey.

Reading the Models Correctly

A note of caution on reading Exhibit 1 as a maturity curve. Large enterprises routinely buy in all four models simultaneously: a Fortune 500 company may standardize on a security platform, outsource a regional SOC, retain an advisory firm for its security program, and engage a global consultancy for AI transformation, all in the same fiscal year. Equally, many organizations enter the market at Model 2 or Model 3 and stay there permanently, because building an internal SOC or security organization will never make economic sense for them. The models describe distinct problems and buying motions, not stages every organization passes through in order. The strategic significance is directional: as AI compresses the cost of operations, the share of security-related spending and executive attention devoted to Models 3 and 4 grows.

The Buyer Changes at Each Model

The least appreciated feature of this market structure, and the one with the largest strategic consequences for providers, is that the buying center changes from model to model (Exhibit 5).

In Model 1, the buyer is typically a security director or CISO working within an IT budget, evaluating technology on features, integration, and total cost of ownership. In Model 2, the buyer is the CISO or VP of security operations. In organizations without a security team it is the CIO or even the CFO, purchasing an operating capability outright. In Model 3, the buying center expands: the CISO is joined by the CFO, the general counsel, and increasingly the audit or risk committee of the board, because the questions being answered are questions of enterprise risk and regulatory exposure. In Model 4, the buyer is the CEO, COO, or board itself, and the budget frequently comes from transformation or digital initiatives rather than from security at all.

The upward migration of the conversation is visible in governance disclosures. Analysis published through the Harvard Law School Forum on Corporate Governance found that the share of public companies disclosing board-level AI oversight rose more than 84% year over year, and more than 150% since 2022. Harvard Business Review made the boardroom framing explicit in 2026: AI is reshaping cyber risk, and boards must manage it as a strategic governance issue rather than a technical one. The buyers for Models 3 and 4 are, quite literally, being told by their own governance literature to enter this market.

This has a hard implication for provider strategy: incumbency does not transfer. A provider that runs a customer's SOC brilliantly has earned credibility with the security operations leader, but the governance engagement in Model 3 will be decided by executives who may never have seen a SOC metric, and the

transformation engagement in Model 4 will be decided by leaders who think of security as one workstream among many. Providers hoping to ride a customer up the value chain must build relationships, references, and language for a different buyer at each step. Land-and-expand works within a model; it does not work across models without deliberate investment.

The same logic explains why the traffic also fails to flow downhill. Global consultancies with board relationships have historically struggled to win SOC operations against specialized MDR firms, because the Model 2 buyer values operational depth and price discipline over boardroom credibility. Each model is defensible terrain. The question for every provider is which terrain it actually holds, and which adjacent terrain it can credibly contest.

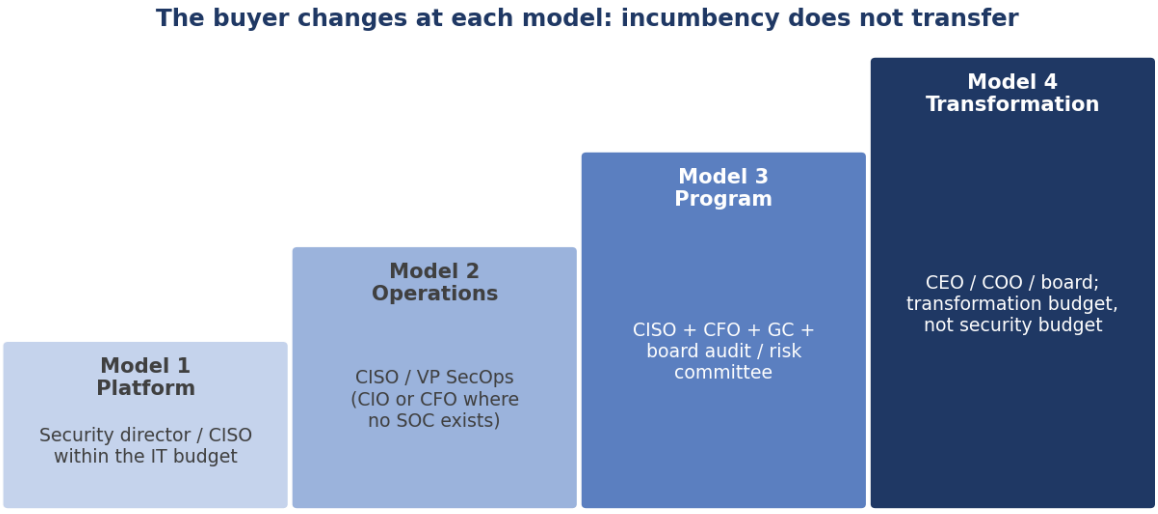

Exhibit 5. The buying center escalates across the four models; relationships must be rebuilt at each step.

Why Traditional Security Categories Are Breaking Down

The inherited categories of MSSP, MDR, XDR and platform describe how services are delivered rather than what customers are buying. That distinction mattered less when delivery models and value propositions lined up neatly. It matters enormously now for three reasons.

Skills, Knowledge and Governance

First, the lanes themselves have blurred. The line that once separated the independent software vendor from the managed service provider was structural rather than cosmetic: one built the product, the other operated it on the customer's behalf, and each needed the other to reach the market. That line has largely dissolved. Platform vendors now sell managed detection and response directly, staffed by their own analysts, competing for the same engagements as the providers who built practices on their software. MDR providers sell advisory. Consultancies operate SOCs. The same capability reaches the buyer through three different business models, and analyst category boundaries are redrawn every cycle to keep pace. A label increasingly describes where a firm started rather than what a customer receives. A buyer comparing an "MDR provider" against a "platform vendor with managed services" is not comparing categories; it is comparing two firms that happen to enter the same deal from different directions.

Second, the categories say nothing about the questions customers now consider most important. There is no analyst quadrant for "who should help our board govern autonomous agents" or "who should decide which

AI initiatives are safe to approve." The buying models framework captures those questions; the traditional categories cannot.

Economics and Tokenomics

Third, AI erodes the economic foundations of the labor-based categories. MSSP and MDR pricing was ultimately anchored to the cost of analyst time. As AI performs triage, correlation, and initial investigation at declining marginal cost, the pricing logic of the entire middle of the market comes under pressure. The industry's own workforce research reflects the shift: ISC2's 2025 Cybersecurity Workforce Study dropped its long-running headline estimate of the workforce gap entirely, because respondents now rank the shortage of specific skills, led by AI at 41%, as more pressing than the shortage of people. Nearly 90% of respondents reported at least one significant security event attributable to skills shortages. Categories built on labor arbitrage do not survive a market that no longer prices labor as the scarce input.

There is, however, an assumption buried in that erosion argument that deserves executive attention: it treats AI's marginal cost as permanently low and falling. That is an input-price assumption, not a law of nature. AI-delivered security work is ultimately priced in tokens, and tokenomics have so far moved in the buyer's favor; Stanford's 2025 AI Index documents a more than 280-fold drop in the cost of GPT-3.5-class inference in roughly 18 months. But the same economics carry real risk in the other direction (Exhibit 6). Per-task consumption is rising even as per-token prices fall, because reasoning and agentic workloads consume orders of magnitude more tokens per task than simple queries, and premium reasoning models have shown little of the price decline seen in commodity models. Much of today's inference pricing is subsidized by venture and hyperscaler capital and may not survive vendor consolidation. Any security capability, whether run internally or purchased from a provider, that is foundationally built on token-based consumption has therefore repriced its cost structure from salaries, which are stable and renegotiated annually, to a commodity input that can move quarter to quarter. The risk is not hypothetical. Uber exhausted its entire 2026 AI budget by April, four months into the year, after agentic coding tools spread to roughly 5,000 engineers faster than its financial models had anticipated, and it responded by capping individual spending at $1,500 per month per tool. Its chief operating officer publicly questioned whether the token spend was yet connected to useful customer-facing features. Meta imposed comparable internal caps in the same period. These are engineering organizations with sophisticated financial controls, and consumption still outran the forecast.

The business's strategic response is the one traditional manufacturing industries learned long ago: treat the input as a commodity to be hedged. This is a sea change for the CFO that trickles down through the rest of the business. For providers, that means model portability by design, multiple AI vendors qualified and swappable as tokenomics dictate, and customer pricing that does not silently pass through input volatility. For customers, it means contractual spend and cost constraints on AI-delivered services, and honest recognition that the cost-per-outcome comparison between AI and human delivery is a floating exchange rate, not a fixed one; a sustained move in token prices could shift work back toward human delivery as quickly as it moved away. This is one more reason the accountable-operator model endures. Absorbing and managing tokenomics risk on the customer's behalf, including the vendor diversification and switching capability few customers will build internally, becomes part of what a Model 2 or Model 3 provider is paid to do.

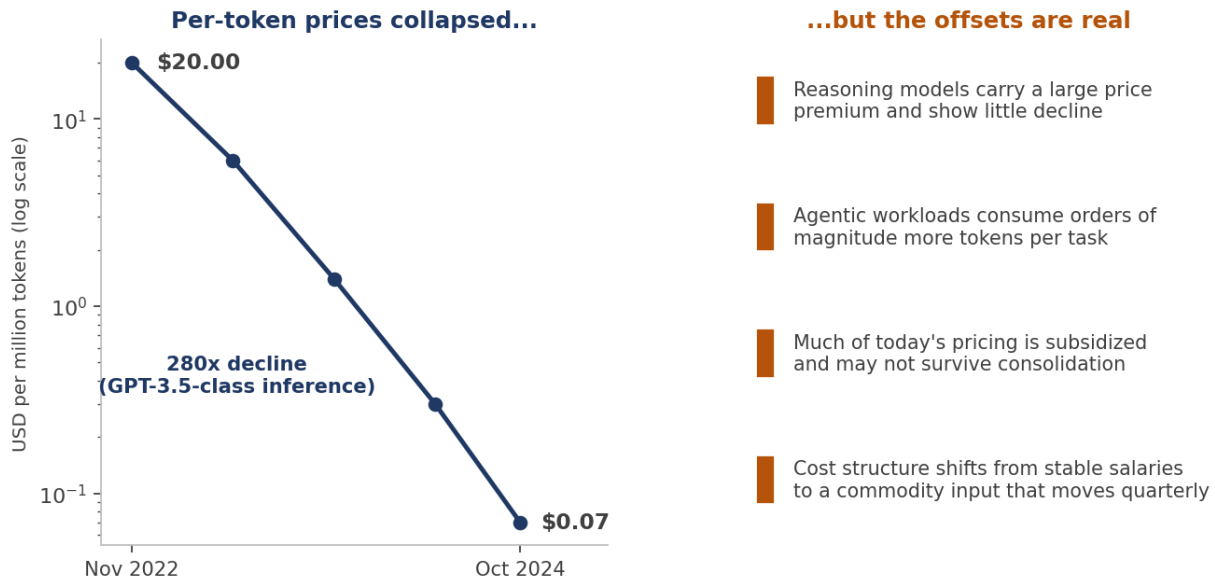

Exhibit 6. The tokenomics of AI-delivered security. Sources: Stanford HAI AI Index 2025; author's analysis.

The AI Security Operating Model: From Copilots to Autonomous Agents

Understanding how AI actually enters security operations clarifies why the market is reorganizing. Adoption is progressing through three broad phases (Exhibit 7), and each phase changes what customers need from providers. The direction of this progression is well established; the pace is routinely overstated, a point the next section examines in detail.

The first phase is assistive. AI copilots summarize alerts, draft investigation notes, translate queries into detection logic, and answer analyst questions. Humans make every decision; AI compresses the time each decision takes. Most of the market is here today, and the benefits accrue primarily to Models 1 and 2 in the form of productivity.

The second phase is supervised autonomy. AI agents conduct end-to-end investigations, assemble evidence, and recommend or execute containment actions within guardrails defined by humans. The analyst's job shifts from performing investigations to supervising them: reviewing AI conclusions, handling escalations, and tuning the boundaries of what agents may do without approval. The SOC becomes, in part, an AI supervision function. Two things are true about this phase simultaneously. It is coming, and the market is already pulling back from its most breathless version. Forrester's 2026 predictions describe a "year of reckoning" in which enterprises defer roughly a quarter of planned AI spending into 2027 as buyers demand proof over promises; the same research found that 80% of security teams have observed risky behavior from AI agents deployed in their own environments, and Forrester expects an agentic AI deployment to cause a major public breach in 2026. Gartner similarly projects that more than 40% of agentic AI projects will be canceled by the end of 2027, largely over inadequate risk controls and unclear business value. The reason for the pullback is not that the technology fails to work; it is that investigation and response ultimately turn on judgment, and judgment is precisely what practitioners report the tools still lack. SOC analysts weigh business context, organizational history, attacker intent, and the cost of being wrong. Current AI does none of this reliably, which is why the durable architecture is supervised autonomy with humans owning the judgment layer, rather than the fully autonomous SOC of vendor keynotes.

The third phase is governed autonomy at enterprise scale. Autonomous agents are not confined to the SOC; they are embedded across the business, in engineering, finance, customer operations, and decision support. At that point the security question is no longer "can AI defend the enterprise" but "who governs the

enterprise's AI." The early evidence suggests governance is trailing deployment: IBM found that 13% of organizations have already experienced a breach of an AI model or application, and 97% of those reported not having AI access controls in place. Among breached organizations, 63% either have no AI governance policy or are still developing one, and among those that do have a policy, only 34% audit regularly for unsanctioned AI. One in five organizations reported a breach traceable to shadow AI, and those with high levels of it saw roughly $670,000 in higher breach costs than organizations with little or none. This is the phase that creates the demand behind Models 3 and 4.

The operating-model progression also pressures legacy pricing. Per-analyst and per-alert pricing made sense when labor was the unit of value. As autonomy increases, pricing migrates toward outcomes, coverage, and accountability: what the provider is answerable for, rather than how many people it assigns.

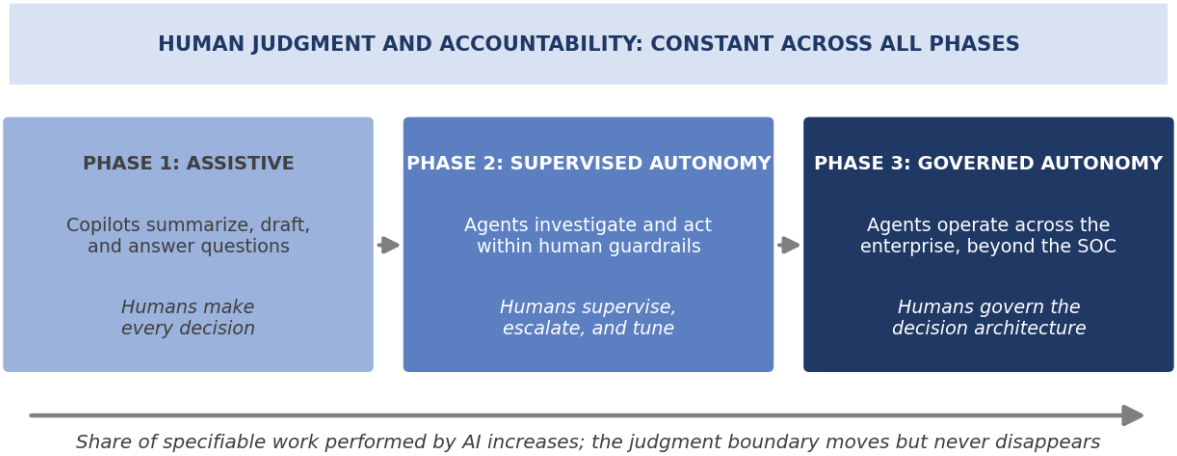

Exhibit 7. The AI security operating model. AI absorbs specifiable work; the judgment layer remains human.

The SOAR Precedent: What the Last Automation Wave Actually Delivered

Before accepting any forecast about the agentic SOC, executives should study what happened the last time the industry declared that automation would remake security operations. The SOAR story is the closest available precedent, and it is instructive from beginning to end.

In 2017, Gartner coined the term SOAR, for security orchestration, automation, and response, to describe a new product class that would integrate security tooling, encode response procedures as playbooks, and progressively automate the work of incident response. The promise was explicit: playbooks would absorb the repetitive work of the SOC, ease the analyst shortage, and put organizations on a path toward automated response. The market responded as markets do. A wave of acquisitions followed, with Rapid7 buying Komand in 2017, Palo Alto Networks buying Demisto in 2019, and Sumo Logic acquiring DFLabs in 2021, as platform vendors raced to own the automation layer. Adoption forecasts were confident: Gartner projected at the time that organizations with security teams of five or more would move from roughly 1% SOAR adoption toward 15% by 2020, with expectations climbing from there. That forecast figure is as reported in contemporaneous trade coverage rather than verified against the original Gartner document, and is offered here to characterize the mood of the period rather than as a precise benchmark.

The reality was different in three specific ways. First, adoption ran well behind the forecasts, and satisfaction ran behind adoption. Playbooks proved brittle: they automated the predictable minority of incidents well, but every environment change broke them, and maintaining them became its own engineering discipline that most security teams could not staff. Second, the technology did not fail; the category did. Automation

succeeded so thoroughly as a feature that it stopped making sense as a product. SOAR capabilities were absorbed into SIEM, XDR, and MDR offerings, and in 2024 Gartner placed the standalone category in the "obsolete before plateau" position on its hype cycle, the designation it applies when a category is overtaken by a competing approach before it ever reaches broad productive use. Eric Ahlm, senior director analyst at Gartner, speaking to Dark Reading at the time, drew the distinction carefully: what was ending was not automation, and not the concept, but the market for buying automation as a separate product. In his assessment the vendors selling nothing but dedicated automation platforms did not have "a very lively future" in front of them. Eric Parizo, managing principal analyst at Omdia, quoted in the same coverage, located the cause on the demand side, observing that SOC decision-makers were never shopping for orchestration and automation as such but for a faster and more consistent threat detection, investigation, and response lifecycle. They took the automation from whichever platform delivered that outcome.

Third, and most important for the present moment: the binding constraint was never the technology. It was judgment. Playbooks could execute anything that could be specified in advance; they could not decide what mattered, weigh ambiguity, or absorb accountability for a consequential call. A decade into the automation era, the SANS 2025 SOC Survey still finds 69% of SOCs reporting metrics through manual or mostly manual processes and AI and machine learning tools ranked at the bottom of the technology satisfaction list, with generative language tools scoring two out of four. The survey's own commentary is pointed. Seth Misner, a SANS faculty fellow, observes that AI and automation are the most commonly planned expansions despite ranking lowest in value delivered and warns that "AI should augment analysts, not replace them," his concern being that leadership will treat AI as a shortcut to fill staffing gaps rather than investing in the talent and integration that actually improve a SOC. The work that automated was the work that never required judgment; the work that required judgment did not automate.

The parallels to today's agentic AI wave are hard to miss (Exhibit 8). The same pattern is visible in sequence: a named category, a burst of acquisitions and venture funding, confident adoption curves, and early practitioner reports of brittleness and oversight burden. This does not mean agentic AI will fail; SOAR did not fail either, in the sense that automation is now table stakes in every serious security product. It means executives should expect four specific rhymes. Agentic capability will be absorbed into platforms and services rather than persisting as a standalone category. Adoption will run at perhaps half the speed of the forecasts. The gap between deployment and operational integration, already visible in the SANS data, will persist for years. The judgment boundary will move but not disappear: agents will absorb more of the specifiable work, and humans will retain the ambiguous, consequential decisions. And the durable winners will be the operators and advisors who fold the new capability into an accountable service, just as the SOAR era's value ultimately accrued to the platforms and MDR providers that absorbed it, not to the standalone automation vendors.

For buyers, the lesson is patience without complacency: adopt deliberately, demand operational proof, and do not pay a premium for a category label. For providers, the lesson is sharper: the money in the last automation wave went to those who sold outcomes that happened to use automation, not to those who sold automation itself. There is no reason to expect the agentic wave to distribute its rewards differently.

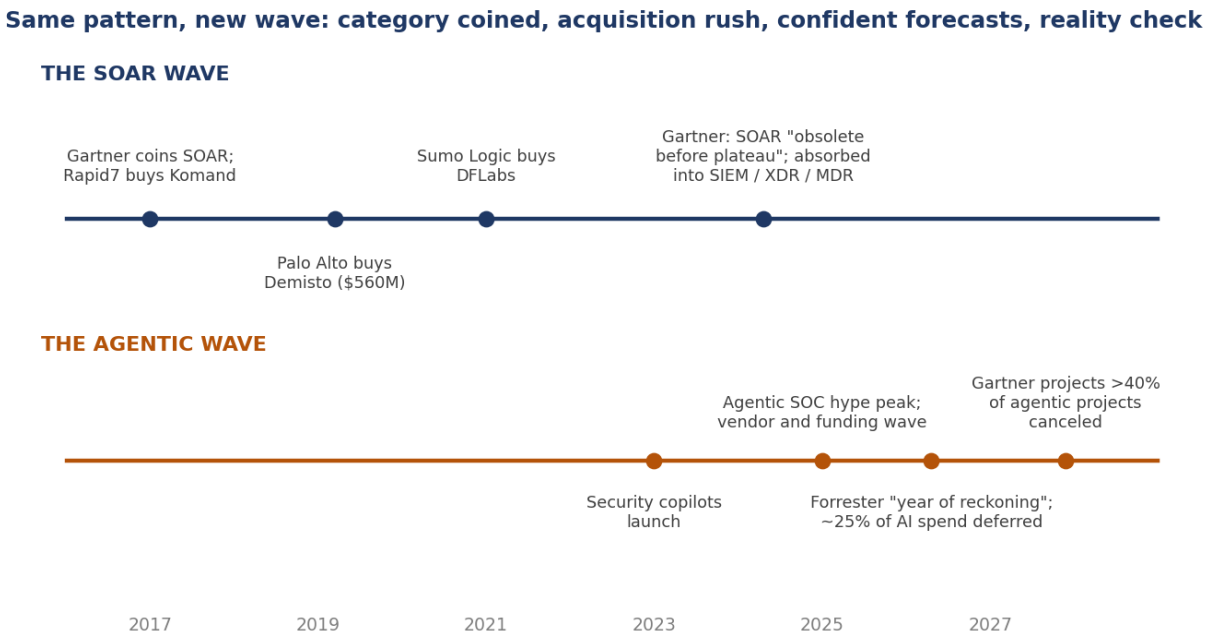

Exhibit 8. The SOAR precedent mapped against the agentic wave. Sources: Dark Reading, Gartner, Forrester; author's analysis. 2027 entry is a projection.

How CISOs Are Becoming AI Business Leaders

The four buying models are mirrored inside the customer organization by the evolution of the CISO role. The CISO who buys in Model 1 is a technology owner. The CISO who buys in Model 2 is an operations leader. The questions in Models 3 and 4 pull the role somewhere new: the CISO becomes the executive responsible for making AI adoption safe enough to be fast.

This is a fundamental repositioning. For twenty years the security function was structurally a brake: it reviewed, restricted, and remediated. AI inverts the demand. Boards and CEOs want to adopt AI aggressively, competitors are moving, and the executive who can say "here is how we adopt this quickly and defensibly" becomes an accelerant rather than a brake. The vacuum above is real: Deloitte found that while only 8% of boards fail to engage management on AI at all, two-thirds of directors say their boards lack sufficient AI knowledge, and NACD's 2026 Governance Outlook ranks AI among the trends directors expect to most affect organizational performance. Someone must fill that gap. CISOs who master AI governance, vendor evaluation, model risk, and regulatory readiness are being pulled into enterprise AI steering committees and board conversations their predecessors never attended. CISOs who remain purely operational will find that an AI risk officer, a chief AI officer, or the CFO's office takes those seats instead.

For providers, this evolution defines the Model 3 opportunity precisely: the product is not a service catalog; it is making the customer's security leadership successful in a role that is expanding faster than most incumbents' skills. The providers that help CISOs make this transition will own relationships that outlast any technology cycle.

Measuring AI ROI Beyond Cost Reduction

Every conversation about AI in security eventually arrives at return on investment, and most of those conversations are conducted in the vocabulary of Model 1: hours saved, alerts closed, headcount avoided.

Those metrics are real but they systematically understate the value at stake, and executives evaluating AI strategy need a broader frame. Four layers of value are worth distinguishing.

Efficiency is the first layer: the same work done cheaper and faster. It is the easiest to measure and the least durable as a differentiator, because every competitor gets the same efficiency from the same tools. It is also the layer most exposed to the tokenomics risk described earlier: efficiency gains denominated in token-based consumption are savings priced in a volatile input, and a sustained move in inference costs can erode them without any change in the security program itself. Efficiency cases should therefore be underwritten net of input-cost assumptions, with the vendor-portability and contractual spend constraints discussed above treated as part of the investment, not as overhead on it.

Effectiveness is the second: better security outcomes, meaning intrusions detected earlier, investigations of higher quality, and incidents contained before they become material. IBM's data quantifies both of these first two layers: extensive AI and automation use correlates with a $1.9 million reduction in average breach cost and an 80-day reduction in breach lifecycle, against a global average lifecycle that fell to a record-low 241 days.

Execution is the third, and the most undercounted by common security metrics today: the value of the AI initiatives the business can pursue because security made them defensible, governance was implemented, and the workforce was ready to act on opportunities. McKinsey's finding that only about 6% of AI adopters achieve meaningful earnings impact is usually read as an AI problem; it is equally a governance problem, because ungoverned pilots cannot be scaled into production. When governance is mature, product teams ship AI features quarters earlier, procurement approves AI vendors faster, and the enterprise captures first-mover economics its cautious competitors forfeit. This value never appears in a security budget line, which is why it so rarely enters the investment case.

Enterprise value is the fourth: trust as a market position. Accenture's research gives this layer empirical shape: the 10% of organizations it classifies as reinvention-ready are 69% less likely to face advanced attacks, 1.5 times more effective at blocking them, and see a 15% higher level of customer trust. Organizations that can demonstrate governed, auditable AI operations win regulated customers, survive diligence in mergers, acquisitions and capital raises, and carry lower regulatory risk premiums. For boards and private equity owners, this is the layer that ultimately matters, and it is the layer that only Models 3 and 4 address.

The practical guidance for executives is to insist that any AI security investment case state which layer it claims to affect and how that claim will be observed. Efficiency claims should show up in unit costs within quarters, and should be stated net of AI input-cost assumptions. Execution and enterprise value claims take longer and require different evidence, but they are where the strategic returns concentrate, and they have a property the efficiency layer lacks: they are largely insulated from tokenomics, because governance maturity and earned trust do not reprice when a model vendor does.

The Counterargument: What If AI Collapses the Models Instead?

The strongest objection to this thesis deserves a direct answer. It runs as follows: if AI makes Tier-1 and Tier-2 security operations nearly free, platform vendors will simply absorb Model 2 into their products, and the market will bifurcate into "products with embedded AI" on one side and "elite advisory" on the other. The middle hollows out, and the four models collapse into two.

Parts of this objection are likely correct. The pure labor-arbitrage MDR business, differentiated mainly by analyst headcount and price, is structurally endangered, and the SOAR precedent shows exactly how platform absorption of an automation layer plays out. But the full collapse thesis fails on three points.

First, accountability does not automate. When an autonomous system takes a containment action that halts a production line, or fails to take one that would have prevented a breach, someone must answer for the decision architecture. Platform vendors have consistently declined to accept operational accountability for customer outcomes, for sound liability reasons, and there is little sign that embedding more AI in their products changes that posture. The demand for an accountable operating partner survives the automation of the operating work; what changes is that the partner is accountable for supervising machines rather than managing analysts. It is telling that even amid the automation wave, analyst forecasts still project the MDR market growing at better than 20% annually through 2030: what is being repriced is the labor inside the service, not the demand for an accountable operator.

Second, a large share of Model 2 and Model 3 demand comes from organizations that have no SOC and no security organization to begin with and never intend to build one. Platform vendors cannot absorb this demand with embedded AI, because these customers are not buying tooling they will operate themselves under any scenario. They are buying an operator. AI changes that provider's cost structure and delivery model; it does not eliminate the customer's need for the provider.

Third, judgment demand is growing, not shrinking. Every increase in AI autonomy generates new governance questions, new regulatory obligations, and new board anxiety. Regulation is compounding this: Gartner expects fragmented AI regulation to extend to three-quarters of the world's economies by 2030, driving roughly $1 billion in compliance spending, and the EU AI Act's shifting compliance deadlines (high-risk obligations recently deferred from August 2026 to December 2027) illustrate that even the timeline itself now requires expert interpretation. Automation commoditizes the work of operations while simultaneously expanding the market for judgment about automation. That asymmetry is why the spending gravity moves toward Models 3 and 4 even as Model 2 compresses.

The Next Five Years

Extending these dynamics forward suggests five developments executives should plan around, offered with the SOAR-era discipline in mind: right directionally, uncertain in timing.

Repricing, then consolidation, in the operational middle. Mid-market MDR providers without a distinctive AI operating model or a credible advisory motion will face margin pressure first and consolidation second. The survivors in Model 2 will be those that turn AI supervision itself into a demonstrable discipline, and those serving the large population of customers who outsource because they will never build in-house.

The supervised SOC, not the autonomous SOC, becomes the reference architecture. The default expectation for security operations will be AI agents performing bounded investigation and response work under human-defined guardrails, with analysts owning judgment, escalation, and accountability. If the SOAR precedent holds, this transition will take longer than vendor roadmaps suggest, and providers will compete on the quality, transparency, and auditability of supervision rather than on claims of full autonomy.

AI governance becomes a standing board obligation. With board-level AI oversight disclosures already up more than 150% since 2022, board committees will treat AI risk the way they treat financial controls: as a recurring agenda item with defined ownership, reporting, and external attestation. This creates a durable advisory and assurance market, and an emerging category of AI assurance services analogous to financial audit.

Pricing shifts from headcount to accountability. Contracts anchored to analyst counts and alert volumes give way to contracts anchored to outcomes, coverage, and accepted responsibility. This favors providers confident enough in their AI operations to warrant them.

Security leadership converges or at least collaborates with AI leadership. In many enterprises the CISO role, the AI governance role, and portions of enterprise risk will merge or report jointly. The executives and providers who prepare for that convergence will define it; those who wait will inherit whatever structure others design.

What This Means for Buyers

For executives buying in this market, the practical implication is to stop asking one relationship to answer both kinds of question. Operational excellence and independent evaluation are separate purchases with separate success criteria, and bundling them costs the buyer the thing hardest to replace, which is a candid answer. The structure that works is straightforward: send the operational work to the best available operator, hold that operator to outcome-based commitments, and place the judgment about whether those commitments are being met, whether the AI inside them is performing, and whether the spend is defensible to a board somewhere that has no revenue riding on the answer.

A simple test identifies whether that independence actually exists. Ask who in the relationship loses money if the honest recommendation is to spend less, cancel an initiative, or leave an existing arrangement in place. If the answer is the party giving the advice, the advice is not independent, however competent and well-intentioned the people giving it are. Beyond that test, the questions worth asking of any advisory relationship are concrete: what AI is being used in the delivery of the service, what it decides without a human, who reviews those decisions and how often, what evidence exists that outcomes improved rather than merely accelerated, how input costs are contracted so a change in AI pricing does not silently reprice the engagement, and what specifically would be presented to an audit committee to demonstrate the program is governed.

One caution against overcorrecting. Judgment without operational literacy is theater. An advisor who cannot read a detection engineering backlog, interrogate a mean-time-to-respond figure, or tell the difference between an AI agent that closed a ticket and one that resolved an incident will produce governance documents and no risk reduction. The judgment layer earns its position by being credible about operations, not by being distant from them. Equally, the operational relationship is the one that produces the outcome, and it should be strengthened rather than destabilized by independent evaluation. The goal is not to introduce doubt into a working partnership. It is to give the executives accountable for the result a defensible basis for the decisions only they can make.

Conclusion

The market has spent the past decade optimizing security operations. The next decade will be defined by something different: helping organizations govern, operationalize, and capitalize on AI itself.

The four buying models offer a more value-driven map of that market than the category labels the industry inherited. They clarify what each customer is actually purchasing, who signs for it, and where the money is moving. The buying-center analysis carries the harder implication: because the decision migrates from the security organization to the finance, legal, and board table as the question shifts from operations to judgment, no provider inherits the next engagement by performing well on the last one. Position has to be earned separately at each level.

The SOAR precedent supplies the tempo. Automation waves in security arrive slower than forecast, get absorbed into platforms and services rather than surviving as categories, and consistently reward those who sell accountable outcomes over those who sell the automation itself. Applied to agentic AI, that history argues for adopting deliberately, demanding operational proof, and declining to pay a premium for a category label.

And through every wave the scarce input has remained the same: human judgment about what matters, what is acceptable, and who answers for the consequences.

Two developments could change this picture, and executives should watch both. The first is input economics. This analysis assumes AI's marginal cost stays low enough that operational compression continues; a sustained move in inference pricing would reprice every model described here. The second is the adversary. This paper examines what customers buy and why, which is a demand-side question. If AI-enabled attackers compress defensive timelines faster than governed autonomy matures, pressure to move judgment out of the loop will intensify, and where that boundary finally settles is a question this paper does not attempt to answer.

What the analysis does support is narrower and more useful. For buyers, the discipline is to stop asking a single relationship both to deliver the work and to grade it, and to require that any AI investment case name which layer of value it claims to move and how that claim will be observed. For providers, it is to choose terrain deliberately rather than drifting across it, and to build the capability every model increasingly demands, which is making AI trustworthy enough to move quickly.

Security spent two decades being measured by what it prevented. The next decade will measure it by what it made possible. That is a harder thing to demonstrate, and the organizations that learn to demonstrate it will set the terms for everyone else.

Sources

A note on method. Every figure in this paper is attributed to the organization that produced it. The World Economic Forum, IBM, Accenture, McKinsey, and SANS findings were verified directly against the primary reports or the publishing organization's own release, as was the trade coverage cited for the SOAR period. Gartner and Forrester projections are drawn from those firms' public press releases and from contemporaneous trade reporting rather than from subscription research, and should be read as directional rather than precise; where a specific figure could not be verified against an original document, the text says so at the point of use. Forward-looking projections are identified as projections throughout.

World Economic Forum with Accenture, Global Cybersecurity Outlook 2026, fifth edition (January 2026), consulted in full: reports.weforum.org/docs/WEF_Global_Cybersecurity_Outlook_2026.pdf

Gartner, Forecasts Worldwide End-User Spending on Information Security to Total $213 Billion in 2025, including the 2026 outlook (July 2025): gartner.com/en/newsroom/press-releases/2025-07-29

IBM, Cost of a Data Breach Report 2025, conducted by the Ponemon Institute across 600 organizations breached between March 2024 and February 2025: ibm.com/reports/data-breach

IBM Newsroom, 13% of Organizations Reported Breaches of AI Models or Applications (July 2025): newsroom.ibm.com

Mordor Intelligence, Managed Detection and Response Market, 2025-2030: mordorintelligence.com/industry-reports/managed-detection-and-response-market

McKinsey, The State of AI 2025 Global Survey: mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai

Accenture, State of Cybersecurity Resilience 2025: accenture.com/us-en/insights/security/state-cybersecurity-2025

Accenture Newsroom, Only One in 10 Organizations Globally Are Ready to Protect Against AI-Augmented Cyber Threats (June 2025): newsroom.accenture.com

Deloitte Global Boardroom Program, Governance of AI: A Critical Imperative for Today's Boards, 2nd edition (2025): deloitte.com/global/en/issues/trust/progress-on-ai-in-the-boardroom-but-room-to-accelerate.html

ISC2, 2025 Cybersecurity Workforce Study (December 2025): isc2.org/Insights/2025/12/2025-ISC2-Cybersecurity-Workforce-Study

SANS Institute, SANS SOC Survey 2025, written by Christopher Crowley (July 2025), consulted in full: sans.org/white-papers/sans-2025-soc-survey. Quotations from Seth Misner and all SOC statistics are taken directly from that report.

Fortune, Uber burned through its entire 2026 AI budget in four months (May 2026): fortune.com/2026/05/26/uber-coo-ai-spending-tokens-claude-code

Robert Lemos, SOAR Is Dead, Long Live SOAR, Dark Reading (September 2024), consulted in full; source of the quoted remarks by Eric Ahlm of Gartner and Eric Parizo of Omdia: darkreading.com/cybersecurity-operations/soar-is-dead-long-live-soar

Stanford HAI, Artificial Intelligence Index Report 2025 (inference cost trends): hai.stanford.edu/ai-index/2025-ai-index-report

Forrester, Predictions 2026: Technology and Security (October 2025): forrester.com/press-newsroom/forrester-tech-security-2026-predictions

Harvard Business Review, AI Is Reshaping Cyber Risk. Boards Need to Manage the Threat (April 2026): hbr.org/2026/04/ai-is-reshaping-cyber-risk-boards-need-to-manage-the-threat

Harvard Law School Forum on Corporate Governance, Cyber and AI Oversight Disclosures: What Companies Shared in 2025 (October 2025): corpgov.law.harvard.edu/2025/10/28

Gartner Newsroom, Global AI Regulations Fuel Billion-Dollar Market for AI Governance Platforms (February 2026): gartner.com/en/newsroom/press-releases/2026-02-17

NACD, 2026 Governance Outlook and AI in Cybersecurity director resources: nacdonline.org

European Commission AI Act Service Desk, Implementation Timeline; Gibson Dunn analysis of the 2026 Digital Omnibus deferrals: ai-act-service-desk.ec.europa.eu and gibsondunn.com