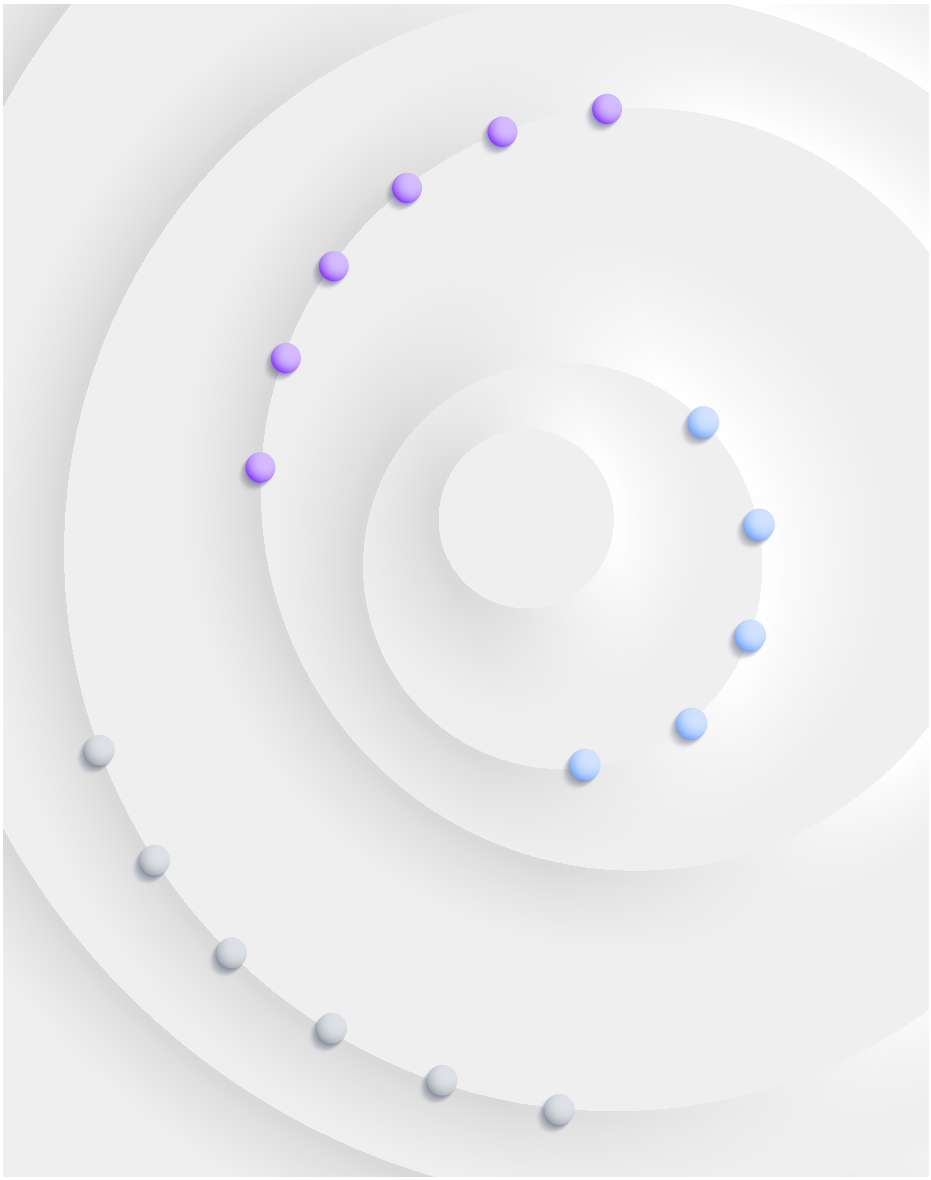

Cybersecurity 2028

Your workforce, built for the AI frontier

Contents

About the AWS + IBM partnership

IBM Consulting brings business and IT transformation experience together with deep industry expertise, methods, frameworks, and generative AI-powered assets and assistants to accelerate your hybrid cloud and AI journey on the AWS Cloud. IBM's expertise in security, enterprise scalability, and open innovation with ecosystem partners help your business to operate with speed, scale, and trust.

As an AWS Premier Tier partner, IBM Security Services helps secure complex workloads running on AWS and accelerate hybrid cloud strategy by using critical security capabilities from both IBM and AWS.

Our partnership empowers businesses to innovate, optimize, and scale with confidence in a rapidly evolving digital landscape. By combining trusted expertise with leading software and services, IBM and AWS enable organizations to seize unprecedented opportunities to achieve transformative, business-driven outcomes at scale.

For more information, please visit:
ibm.com/aws/security

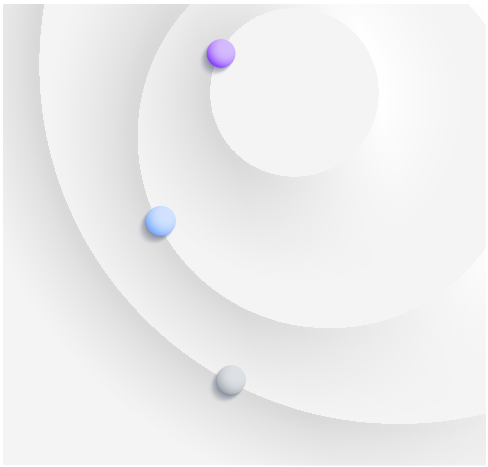

Key takeaways

"One of the biggest threats is not any adversary, but our own lack of imagination."

Koos Lodewijkx, Vice President, CISO, IBM

Security's new core is AI.

AI is now a strategic capability at the center of the operating model, transforming how tech and security teams work together. Over the next three years, AI augmentation is expected to increase 50% while the use of gen AI security capabilities is expected to grow 63% Nearly two-thirds of executives (64%) expect every employee in their IT/IS organization to be using AI agents within the next two years.

A 36-month sprint will separate the AI resilient from the AI disrupted.

Organizations are in one of three states at the moment vis-à-vis AI-centric security operations: crawl, walk, or run. 18% are crawling, 52% are walking, and 30% of organizations are running on the fast track.

AI-first transformation can pay for itself.

Not only in terms of budget but also in terms of effort. Executives estimate they can save an average of 10-20% of their total cybersecurity budget through the adoption of advanced AI use cases. See page 25 for more details.

A look at our study

7 sectors and industries

Average annual revenue of surveyed organizations **~$21B**

17 countries across Asia Pacific, Europe, Middle East, Latin America, and North America

Over
550 CISOs, CTOs, and CIOs
450 CEOs, CFOs, COOs, and CHROs

Chief information security officers (CISOs) have spearheaded the evolution of digital security, converting every technological upheaval into an opportunity to strengthen their organizations' defenses. The playbooks for most security and technology leaders were written for the growth opportunities of cloud migrations and digital services—tempered by the expanding attack surfaces these created. And yet, many of the assumptions that underlie these playbooks are becoming obsolete. Why?

First, human-intermediated operations are being replaced by AI-intermediated technologies.[1] And second, frontier AI models are evolving so rapidly leaders can't foresee with any confidence where capabilities will be in the near to mid-term—complicating their planning and investment decisions.[2]

What is clear: advanced AI modalities are changing how organizations work in fundamental ways—creating opportunities but also introducing uncertainties and risks.[3] For instance, more than one in four AI initiatives have been cancelled, postponed, or failed to scale because of security concerns, and more than one in three organizations indicate their AI capabilities have already been compromised by cyberattacks.[4] The unchecked growth of nonhuman identities fueled by generative AI actually expands the attack surface, creates new targets for threat actors, and leads to a loss of visibility into who or what is accessing critical systems.[5]

Fittingly, AI technology can help CISOs with AI threats. However, instead of traditional operations enhanced with AI, AI is moving to the center of the operating model—evolving into a strategic capability that is redefining how technology and security teams work together and collapsing the distance between business operations and business outcomes. Yet our research reveals security and operations leaders are living in a house divided, with sentiment split on whether their organizational culture is inhibiting (44%) rather than enabling change (56%).

One conclusion from our analysis: the next 36 months represent a critical window of opportunity for CISOs. We believe how leaders respond during this time will separate the AI-resilient from the AI-disrupted.

Beating the ticking clock: The 36-month sprint

In this report, we map out the transformation journey toward AI-first operations. Hint: cybersecurity workforce development and changes to the IT/IS operating model will play a big role. Organizations are in one of three states at the moment: crawl, walk, or run.

Roughly one in five enterprises are in "crawl" mode

Based on insights from over 1,000 security, technology, and business executives across 17 geographies and seven industry sectors (see research methodology on page 29), a minority of organizations are still in the early stages of modernizing their cybersecurity delivery and support model. Our analysis reveals approximately one in five organizations (18%) are still in a "crawl" state when it comes to AI transformation. These enterprises are largely focused on planning and incremental change, with little to show for their transformation efforts to date.

The majority are in "walk" mode

Most organizations (52%) are finding their way, typically following one of two paths: prioritizing AI-first workforce transformation or enhancing IT/IS integration. We characterize these organizations as being in the "walk" phase of their transformation journey, where leaders are activating their workforce and building momentum for change. The difference for these organizations is that AI is positioned to become the connective tissue running throughout their entire organization.

The future is all "run"

30% of organizations have moved the fastest toward building their AI-first foundation. These organizations are entering the "run" phase of their transformation journey, where cybersecurity capabilities achieve a new level of visibility, reach, and autonomy. In this state, IT/IS operations have evolved to become more self-regulating, self-correcting, and self-healing, with the ability to deliver not only automated (scripted) responses but effectively self-initiated and goal-oriented cybersecurity outcomes. In the run state, cybersecurity capabilities are largely autonomous—formulating decisions, testing their own hypotheses, and leveraging agentic AI to both enrich decisioning and automate inference-driven actions. These organizations are adapting rapidly to changes in the operations environment, opening the door to greater resilience, innovation, and growth—safely and at scale.

> "We recently had a conversation with one of the leaders of the new frontier AI firms. He said that over the past two years, the capabilities of models have doubled every 10 months. If that growth curve continues, how long will it take for these capabilities to be 10 times what they are today? It's about three years. That puts us in 2028."
>
> Koos Lodewijkx, Vice President, CISO, IBM

Perspective

Cybersecurity's AI boom is changing all the rules

This isn't just another tech tremor; the AI revolution is a seismic event, a fundamental rewriting of the rules of digital engagement.[6] This presents both opportunities and risks. According to one study, the complexity of tasks "generalist AI" can complete has been doubling roughly every seven months for the last six years.[7]

Executives anticipate new AI modalities will grow significantly over the next three years: AI augmentation is expected to increase 50%, the use of gen AI security capabilities is expected to grow 63%, with AI agents (48%) and workflow automation and orchestration (45%) expected to play a more prominent role in day-to-day operations.

But to fully seize these opportunities, CISOs must manage three distinct paradigm shifts—and all at the same time. Tackling these demands simultaneously—not sequentially—is essential to securing the organization for 2028 and beyond.

AI: New uses, new growth

Qs: Estimate the percentage of cybersecurity workloads that rely upon the following capabilities—Today (n=1013).
Estimate the percentage of cybersecurity workloads that will rely upon the following capabilities—In 3 years (n=1013).
Note: Items are not mutually exclusive so percentages do not sum to 100%

First, organizations are racing to get ahead of AI-enabled threat actors who are applying new capabilities to learn, adapt, and evolve with relentless speed and sophistication.[8] Forget simply securing *where* data lives or *how* applications are secured; AI is rewriting *who* organizations are defending against and *what* makes the split-second decisions that determine the trustworthiness and resilience of their digital operations. Indeed, cybercriminals are already leveraging AI tools to outmaneuver and outpace enterprise defenses, using adaptive malware, deep fakes, and automated attacks to create and exploit new vulnerabilities.

Second, CISOs must address the demand for new AI services across the enterprise.[9] As AI becomes the connective tissue in every process and function, how they enable secure and trusted AI needs to change. The next wave of AI innovation—multiagent systems, large action models, and synthetic data—promises to redefine what's possible, pushing AI deeper into mission-critical operations.

And third, CISOs are grappling with how to use AI for their own operations.[10] From automation's efficiency, to generative AI's creative problem-solving, to agentic AI's autonomous actions, AI is rapidly becoming the cybersecurity team's most powerful ally.

While AI capabilities are surging, many enterprises are stalled by ungoverned sprawl and tangled clouds.[11] The core question for security and tech leaders: can they pivot fast enough to build AI operations at scale, or will legacy infrastructure and operations derail their AI transformation?

> "A few weeks ago, I had a meeting with a bunch of start-ups. I asked how many of them were training their own large language models. There were a few doing so in focused areas. But what surprised me was how many of them were not. They told me large language models are evolving so fast that when they pause to do their own fine tuning or training, they find that their capability falls behind the general purpose models. The evolution is happening amazingly fast. I don't think that this growth will continue forever but that's the phase we're in now. We're just seeing massive changes month over month."
>
> Chris Betz, Chief Information Security Officer, AWS

Perspective

A litmus test for cyber resilience in 2028

> "We know that things are going to be very different in 2028. Are businesses going to take advantage of this? I'm going to give you an answer that I expect to be entirely right: some are and some aren't."
>
> Chris Betz, Chief Information Security Officer, AWS

Three questions to ask yourself.

Over the next three years, cybersecurity operations will reorient around new AI capabilities that span the IT/IS portfolio. Below are three areas where change is likely to be significant.

Does your AI provide total visibility?

IT/IS observability and resilience go beyond traditional monitoring. They are about leveraging AI to provide a deep, real-time understanding of your entire IT/IS environment. By continuously analyzing vast streams of logs, metrics, and traces from every system, AI can deliver unprecedented visibility into system behavior and performance. This enables proactive identification of anomalies, predictive troubleshooting, and the ability to optimize systems for maximum availability and reliability. It's the foundation for truly understanding and strengthening your digital posture. Observability enables IT and IS teams to proactively manage and optimize their systems, helping ensure high availability, reliability, and performance.

Does your AI deliver smarter operations?

AIOps isn't just about automation; it's about infusing intelligence into your IT and security operations. It involves using AI and machine learning to analyze massive data sets from diverse IT systems—logs, metrics, events, and more—to automatically detect complex anomalies, predict potential issues before they impact operations, and automate critical incident responses. This significantly improves operational efficiency, reduces mean time to resolution (MTTR), and frees your teams to focus on strategic initiatives rather than reactive firefighting.

Can your AI self-direct for successful outcomes?

Autonomous cybersecurity solutions represent the evolution to self-managing security—using AI, machine learning, and advanced automation to detect, respond to, and mitigate cyber threats with minimal or no human intervention. Beyond simple threat blocking, these systems can autonomously identify vulnerabilities, analyze threats in real time, and take decisive actions—such as containing an intrusion or enforcing policy controls. This dramatically enhances the speed, efficiency, and scalability of your security operations, allowing for continuous defense against rapidly evolving threats.

AI overload: Finding signal in the noise

For all the interest and investment in AI, there is an equal amount of confusion and indecision. Security concerns have emerged as a primary obstacle to widescale enterprise AI adoption.[12]

One hurdle is a perception that innovation and security are somehow at odds—creating some strategic paralysis among technology executives who may be thinking they need to choose between the two.[13] Opinion is evenly divided on whether new AI governance frameworks create excessive friction, with 49% saying they do. Similarly, leaders are split on whether organizational culture is impeding transformation, with 44% seeing their culture as a barrier rather than an enabler.

Further skepticism is evident from the 62% of executives who say their organizations are placing excessive faith in AI's transformative potential. This finding underscores a growing impatience among stakeholders: it's time for theoretical benefits to translate into measurable business outcomes.

Agentic AI: The next frontier

Looking ahead, autonomous AI agents represent the most promising development on the near-term horizon:

76% of executives anticipate AI agents will fundamentally improve operations within two years.

72% view AI agents as key catalysts for organizational innovation.

67% anticipate agents will significantly enhance ROI on existing AI investments.

64% expect universal adoption of AI agents among IT/IS staff in the same timeframe.

These insights suggest that while the AI security landscape remains challenging, executives are identifying specific technologies—particularly autonomous agents—that can justify accelerating investment in AI infrastructure despite the hurdles. In other words, some IT/IS, business, and operations leaders are finding signal in the noise.

FIGURE 1

Leaders express optimism about AI agents

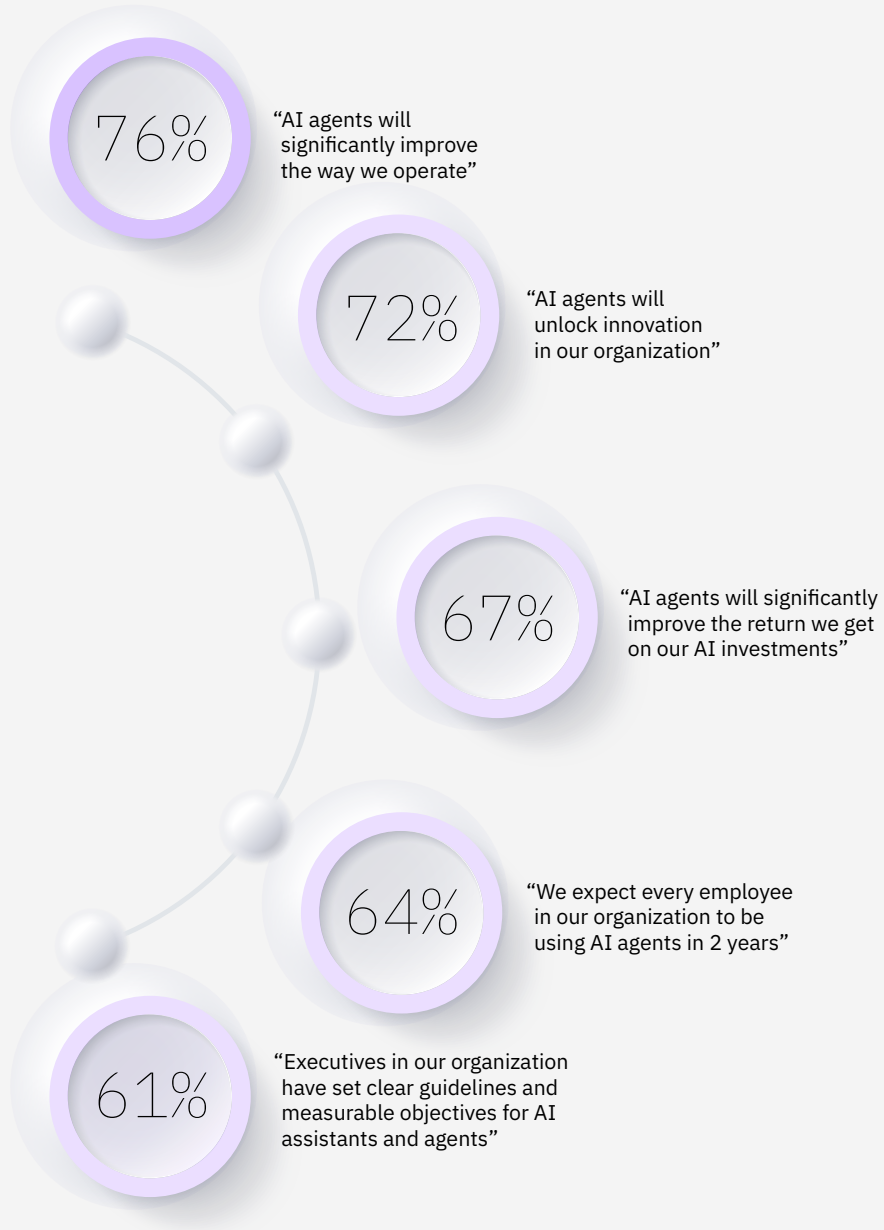

Q: To what extent do you agree with following statements about the use of AI agents within your IT/ IS organization? (n=1013)

"We need everybody to be thinking about this and be very deeply involved in what AI is capable of and where it's going. Where we use it today, where we can't use it today, where we can use it in six months when a whole new set of use cases become available to us. Because I think the capabilities are advancing faster than people's imagination."

Koos Lodewijkx, Vice President, CISO, IBM

Bridging the AI knowledge gap

Many organizations are in a crawl state not only because of issues around governance and ROI, but also because they need to prepare their security workforce. Some are making significant strides in workforce preparation—but a troubling disparity between executive and employee readiness threatens long-term success.

The talent equation has emerged as an equally critical factor in AI implementation. According to a recent IBM IBV study of C-level executives, talent, skills, and cybersecurity are the leading challenges organizations will face over the next three years.[14]

FIGURE 2

Executives point to talent, skills, and cybersecurity as their biggest challenges

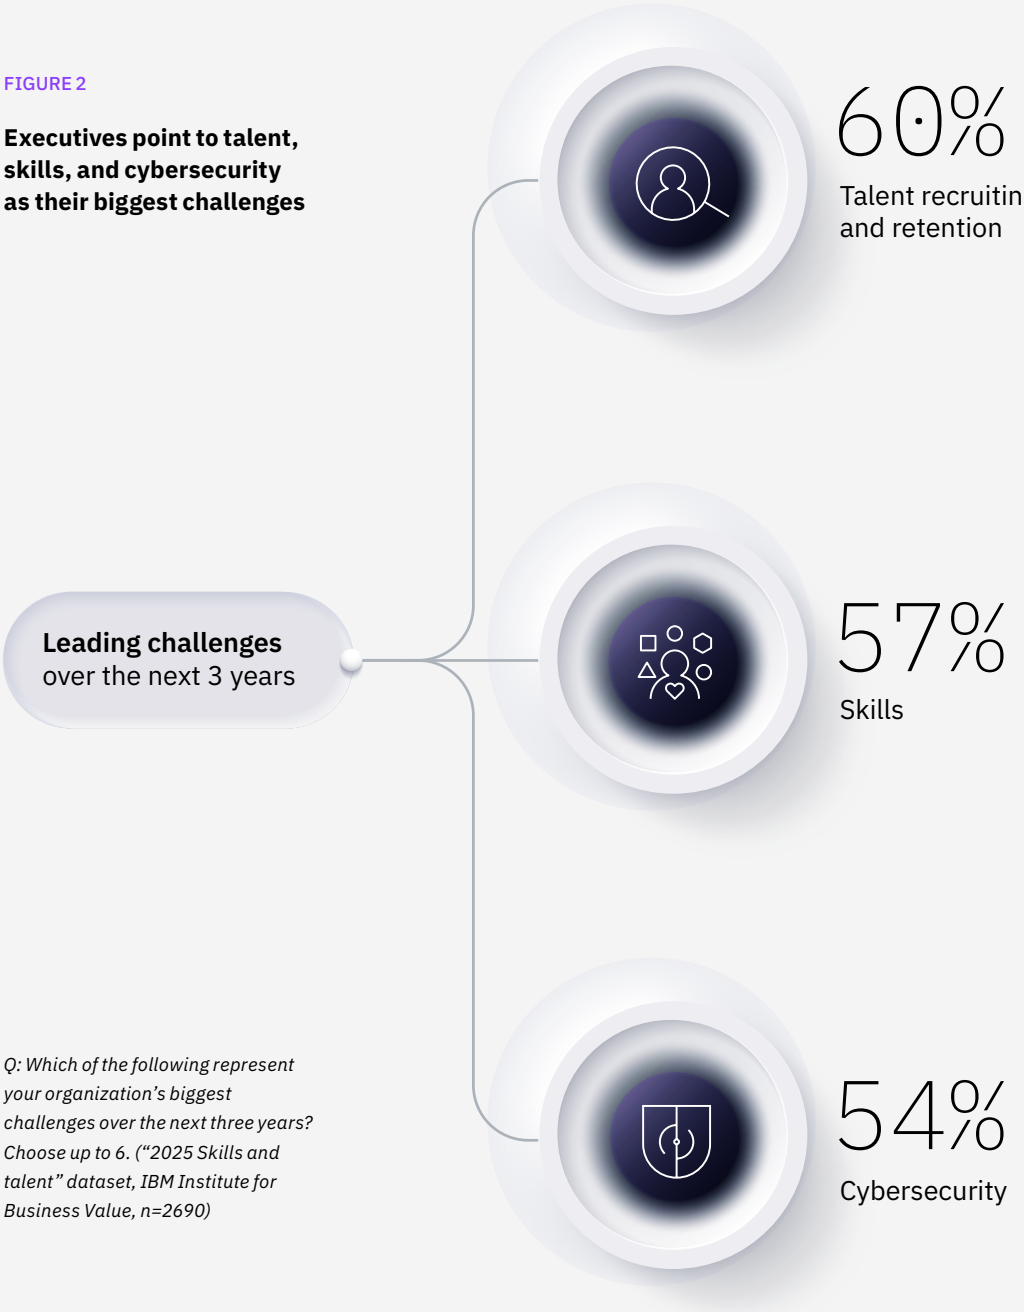

Q: Which of the following represent your organization's biggest challenges over the next three years? Choose up to 6. ("2025 Skills and talent" dataset, IBM Institute for Business Value, n=2690)

Organizations have moved swiftly to address security workforce needs, with 65% establishing formal upskilling plans to support their AI strategy. Nearly as many (63%) have launched structured change management programs to integrate AI assistants and agents, while 60% have delivered formal training on how AI will reshape daily work.

Yet beneath these encouraging metrics lies a concerning reality: AI readiness remains disproportionately concentrated at the executive level. The data reveals a significant disparity—82% of executives demonstrate fluency in AI capabilities, limitations, and responsible use practices, compared to just 53% of frontline employees. Similarly, while 52% of organizations have embedded AI skills into executive development pathways, only 38% have extended these opportunities to their broader workforce.

This asymmetry isn't just a talent issue—it represents an existential threat to AI transformation itself. When the people actually using these tools lack fundamental understanding, even the most sophisticated AI implementations will inevitably falter in practice. Closing the executive-employee AI fluency gap over the next 36 months is an essential brick in the foundation upon which sustainable AI transformation must be built.

> "We know that there are going to be right ways and wrong ways and we don't know what all those hidden traps are yet."
>
> Chris Betz, Chief Information Security Officer, AWS

FIGURE 3

Top-heavy AI: Readiness skewed toward executives, not employees

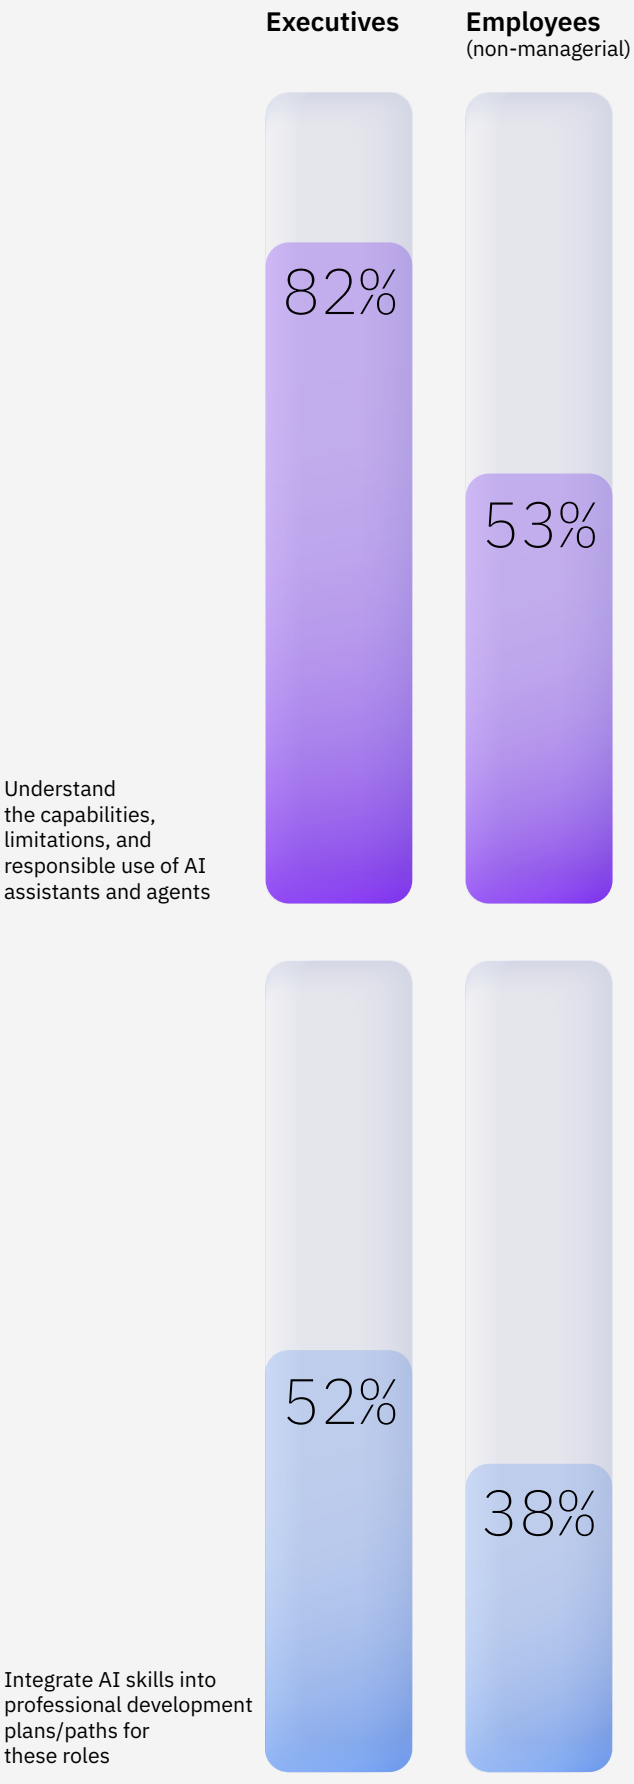

Perspective

Adopting security ABCs speeds operating model transformation

Security awareness, behaviors, and culture practices (security ABCs) are the underpinning for AI-centric operating model transformation. Many CISOs still have work to do in this area. For the most part, only a minority of CISOs describe their organizations as having robust capabilities in AI-related security ABCs.

Q: Considering your IT/IS organization's ability to counter AI-related threats, asses your current capability across the following factors: [Percentages represent responses for "Wide-scale capability (most use cases)" and "Comprehensive capability (virtually all use cases)"] (n=1013). Note: To assess the relative security maturity of the organization, the above factors were compiled into a "Security ABCs" index. Index scores were combined with other factors to assess the organization's AI transformation progress, as described in "Research methodology" on page 29.

Perspective

Adopting security ABCs speeds operating model transformation

Feeding the security flywheel

Security awareness, behaviors, and culture reinforce each other, building momentum for change.

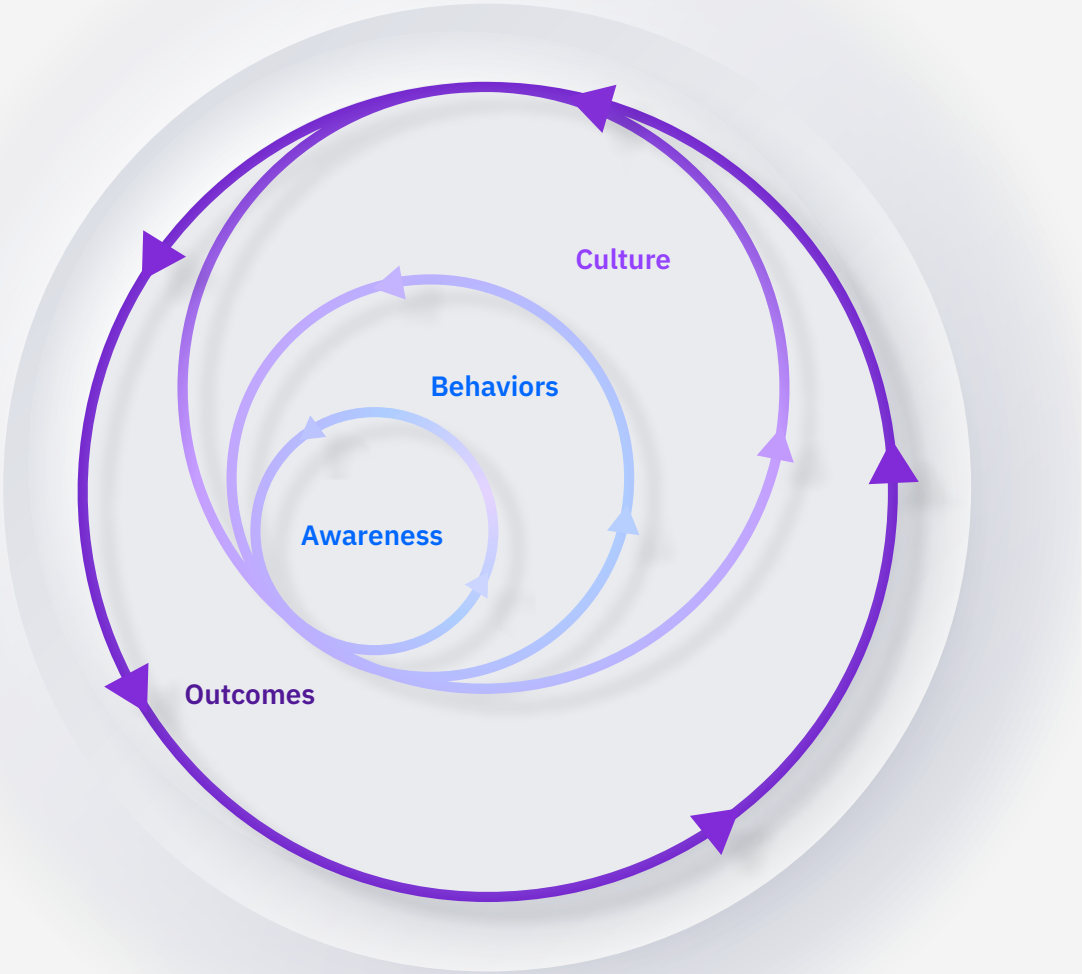

> "You've got to keep feeding the flywheel, but once you get it going, it keeps momentum that will help you deal with the bumps you encounter along the way. I think of security ABCs as the flywheel that helps power things forward, that helps us adapt."
>
> Chris Betz, Chief Information Security Officer, AWS

Case study

Singapore's strategy for public sector gen AI adoption

Singapore's Government Technology Agency (GovTech), responsible for the nation's digital public services, has created a model for enterprise-scale generative AI adoption.[15] GovTech saw generative AI's potential to transform public services, from streamlining processes to personalizing citizen interactions. A major challenge was the high cost of running large language models (LLMs) at a national level.

GovTech's solution (aka MAESTRO) is designed to deliver cost-efficient, pre-built, and production-ready generative AI capabilities across government agencies.

Jeffrey Chai, MAESTRO Product Manager at GovTech, noted the strategic importance: "Making generative AI more accessible and sustainable for our agencies is critical. This allows us to use its potential while managing resources responsibly."

GovTech's platform and adoption model offer a blueprint for CISOs and executives globally. Their experience shows that focusing on optimized cost structures and accessible platforms facilitates scaling of generative AI, delivering significant operational improvements and better citizen services efficiently.

Key outcomes and strategic approaches

- 75% improved cost performance for gen AI workloads. GovTech optimized its computing infrastructure for generative AI's demands. They selected high-performing foundation models and used techniques such as model quantization. Amazon Bedrock and SageMaker JumpStart allowed them to select right-sized models, and then deploy multiple smaller, specialized models instead of large, general ones. This approach significantly cut resource consumption and costs without losing capability.
- Accelerated use cases and broad accessibility. The platform features a no-code, unified, web-based interface that simplifies ML model building, training, and deployment. This made complex machine learning accessible to nontechnical staff. By removing technical barriers, GovTech rapidly increased generative AI use across the government. Within nine months, the platform was adopted by 20 public sector organizations, involving over 45 project teams and more than 300 data scientists and ML engineers.

Real-world impact and operational transformation

GovTech's platform investment is delivering clear benefits:

- Ministry of Manpower (MOM). Used the platform to build an AI "sensemaker" tool. It processed over one million documents in three months, boosted insights extraction by 60%, and cut sensemaking time by 50%, saving over 2,000 work hours. MOM also deployed an automated job classification tool that processed 10 million job postings with 92% accuracy in three months.
- Central Provident Fund Board (CPFB). Used the platform to summarize transcripts from about 600,000 annual citizen calls. These AI-generated summaries help with follow-ups and identifying emerging public issues, improving service quality and operations.

Action guide

Security in motion

Break out by committing
to an AI-first security strategy

AI-first transformation leaders (CISOs, CTOs, CIOs)

Commit. Convert AI experimentation into AI innovation and execution. Understand the priorities of your business and operations counterparts and discuss how distributing technology and security expertise could improve efficiency.

Use AI as a catalyst. Create a holistic view of AI capabilities across infrastructure and operations. "Solve once" by architecting AI, cloud, and security in concert and at scale, then modernize your technology and security operations according to AI-first principles.

Connect everything. Build bridges across the organization by connecting security ABCs to goals and outcomes. Work on building momentum to make faster decisions and to power change.

Become a bridge builder. Focus on creating a better value proposition by connecting core security principles such as trust, integrity, and resilience with business objectives such as innovation, speed, and growth. Rally around these to create a flywheel effect.

Leaders critical to AI-first success (CEOs, CFOs, COOs, CHROs)

Make security your Rosetta Stone. Use security ABCs to align business, operations, technology, and security stakeholders. Articulate priorities for each domain and reconcile how innovation and governance mechanisms should reinforce each other. Stop thinking of risk, governance, innovation, and growth as separate things.

Walk in my shoes. Assemble a working group of business, operations, technology, and security leaders to understand where domain AI strategies intersect. Identify some objectives for a cross-functional secure and resilient AI roadmap. Instead of trying to cover all bets, start with a few, high-impact investments to understand how to support AI use cases across the operations lifecycle.

Find your groove. Create a flywheel effect by inviting technology and security experts onto your product and service teams. Work with your IT/IS counterparts to identify ways to remove friction, accelerate delivery, and improve customer outcomes.

Everything connects: Ambient AI across the enterprise

AI isn't just another technology layer—it's becoming the central nervous system of modern enterprise operations. On the horizon, multiagent systems, large action models, and synthetic data will extend AI capabilities into new domains, increasing reliance on AI-driven outcomes. Yet many organizations find themselves swimming in solutions and managing complex cloud infrastructure, unprepared for delivering intensive AI-centric operations at speed.

This mirrors what on-premises infrastructure organizations experienced during the cloud transition. Just as cloud services introduced the shared responsibility model as a fundamentally new operational approach, organizations today must adapt to an emerging cooperative work model centered on AI-workforce integration—or risk being left behind. Successful organizations will be those that learn how to use, build, deliver, and scale AI responsibly.

The hard data on current performance

Our research reveals a sobering reality: current operating models are delivering mostly mixed outcomes. Despite 24% of organizations claiming alignment across technology, security, and talent strategies, even these leaders show uneven results—stronger threat management but disappointing returns on cybersecurity investments. Across common operational metrics, different strategies reflect different trade-offs, with no clear correlation between approach and performance.

> "I want to be surrounded by smart people solving hard problems and let computers do the rest. This generation of gen AI is going to accelerate that. The kinds of things AI can do is going to increase. The people who are going to be most successful over the next few years are going to be those that use AI and automation across the lifecycle so they can focus on using their human judgment, experience, and knowledge to solve the really, really hard problems that AI solutions cannot."
>
> Chris Betz, Chief Information Security Officer, AWS

The talent crisis meets AI opportunity

The well-documented cybersecurity staffing gap is now colliding with fierce competition for AI expertise.[16] In many cases, tasks are being redirected to AI solutions and roles are being redefined. Recruiting just one of these specialized professionals takes an average of 99 days, according to executives—nearly a full business quarter spent searching for AI security expertise. Meanwhile, survey participants shared that an alarming 21% of security team members leave annually, taking valuable training and institutional knowledge with them.

What keeps CISOs awake at night? Finding professionals who can master "the mission-critical trifecta": navigating the regulatory labyrinth, speaking fluent business, and battling the ever-evolving threat landscape—each cited by 54% of our respondents. These complex, judgment-intensive areas are precisely where the augmented workforce—humans and AI working in concert—will create competitive advantage.

Forward-thinking organizations are already reimagining talent development. IBM has pioneered "Cyber Academy" to bring specialized security talent into specific industries and markets, while its "Cyber Campus" features virtual cyber range experiences to accelerate practical skills.[17]

FIGURE 4

AI transformation is driving an increase in talent budgets

Organizations are spending more (not less) on talent as a result of AI. And leaders are expecting to spend a growing percentage of their talent budgets on reskilling for AI.

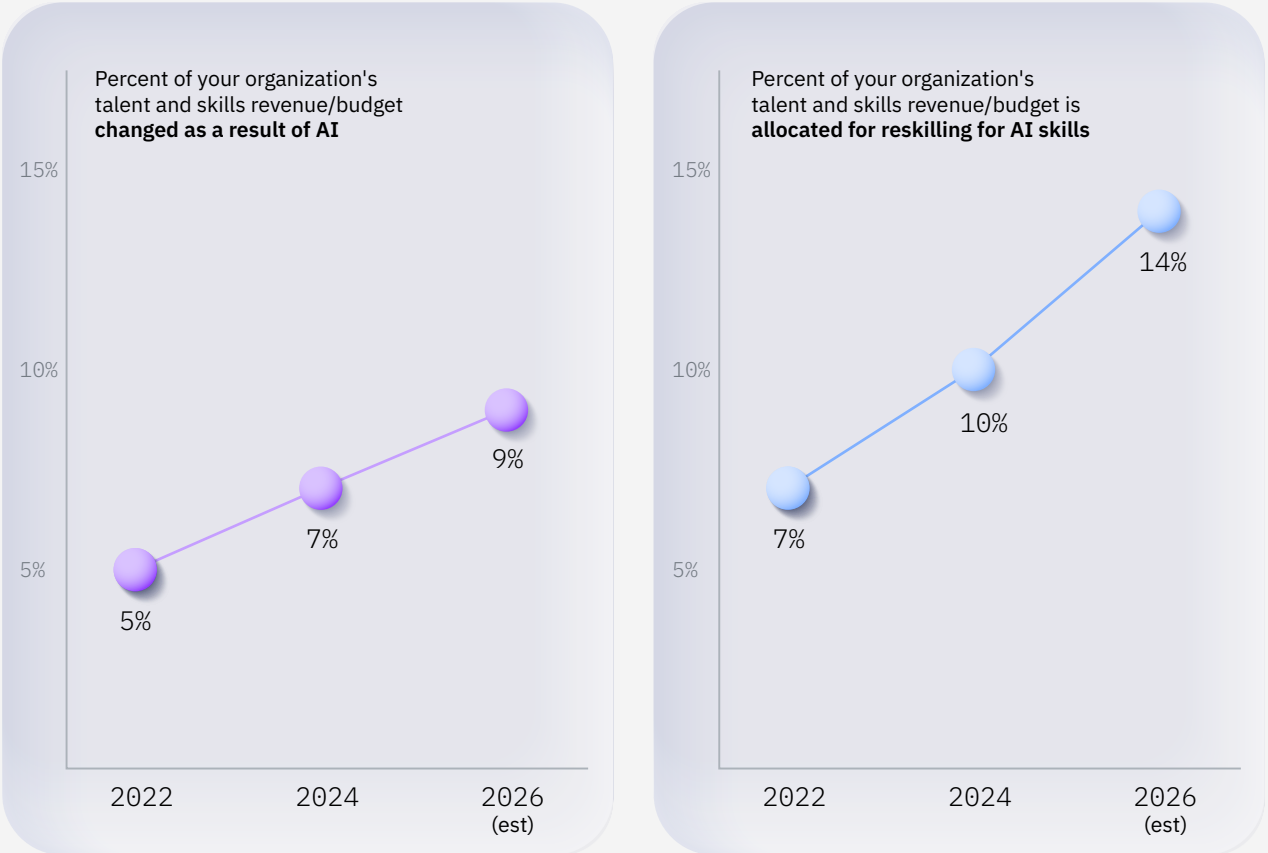

Qs: What percent of your organization's talent and skills revenue/budget changed as a result of AI? ("2025 Skills & talent" dataset, IBM Institute for Business Value, n=2579. What percent of your organization's talent and skills revenue/budget is allocated for reskilling for AI skills? ("2025 Skills & talent" dataset, IBM Institute for Business Value, n=2579).

> "Over the past five years, we've seen significant upskilling because of the way we've adopted automation. Our people are better and our tools are more effective. We've gotten to a place where over a 24-hour period, we're starting to see 100% of our threat responses be completely automated, with no human intervention. That's the direction we want to go."
>
> Koos Lodewijkx, Vice President, CISO, IBM

The evolution of security operations

Just as today's cutting-edge aircraft rely on intricate autopilot systems while still demanding human pilots for critical maneuvers, cybersecurity is undergoing a similar transformation. Once painstakingly manual threat management and incident response tasks are evolving into increasingly autonomous capabilities, with human analysts strategically positioned to handle complex escalations—especially as digital perimeters explode across sprawling partner ecosystems.

The numbers tell the story: 60% of executives describe administering their organization's risk, security, and compliance posture as "highly effort-intensive," while an even more striking 78% characterize it as "highly expertise-intensive." Human expertise is considered essential in 67% of cybersecurity workloads today, but this reliance is expected to drop by 34% over the next three years.

Decentralizing expertise: The security culture solution

So the dilemma for leaders is: how to concentrate resources and maximize efforts, given the talent and skills crunch? The most forward-thinking organizations are tackling the AI readiness challenge with a counterintuitive approach: rather than centralizing AI security expertise and responsibilities, they're deliberately dispersing them throughout operations. They're distributing security responsibilities out into the business and sharing responsibilities with product owners and practice leaders across the organization (see *Perspective: At AWS, a vibrant security culture powers business speed, scale, and growth*).

This distributed responsibility model represents a radical departure from traditional cybersecurity approaches. Instead of building fortress walls around AI systems, these organizations are cultivating security awareness as a shared responsibility, embedding AI governance principles into the daily workflows of product owners and practice leaders across departments.

What makes this approach particularly effective is its recognition that AI transformation isn't merely a technological shift—it's fundamentally about enhancing human decision-making. By democratizing both AI capabilities and security responsibilities, these organizations are creating ecosystems where employees become active participants and owners rather than passive recipients of AI transformation.

Perspective

At AWS, a vibrant security culture powers business speed, scale, and growth

> This initiative is more than scaling a security team; it's about scaling a culture of security across the entire organization.

Forget the outdated notion of a centralized security team playing catch-up. AWS has pioneered a fundamentally different approach with its Security Guardians program,[18] effectively distributing security ownership deep within its business and development units. As AWS CISO, Chris Betz, explained: "Traditionally, we've seen business, operations, technology, and security as distinct and separate things. We're now seeing an opportunity where that has the potential to change in some tremendous ways." This isn't just a minor tweak; it's a strategic evolution in how security scales in a high-velocity environment. "We can inject security expertise deep into the organization," Betz said. "The reason this works is because the organization is asking for it, values it, desired it. This is not about security driving the business."

AWS Security Guardians' 5 steps to success:

1. Start with the "why." AWS grounds this program in tangible business challenges and opportunities, working backward to define a compelling security vision directly relevant to core objectives. This helps ensure the program isn't a theoretical exercise but a practical solution driving meaningful business outcomes. They also understand the power of branding, giving the initiative names such as "Security Guardians" that resonate and foster a sense of shared identity.

2. Identify natural leaders. Instead of mandating participation, AWS strategically identifies individuals and teams already exhibiting an interest in security. Leveraging existing relationships, particularly between application security engineers and product teams, allows them to tap into inherent enthusiasm. These "Security Guardians" often emerge organically within business units, requested by their own leadership, making them highly effective liaisons between technical security requirements and business realities. "One of the things that makes our Security Guardians program so powerful is that it comes from the business. We haven't forced this on the business. The business has seen the need; they came to us. We've found the Guardians to be a great way to build and deliver trust faster and more efficiently."

3. Establish clear expectations. AWS doesn't leave the program's impact to chance. They clearly define the desired security-conscious behaviors for everyone involved—the Guardians, the developers, and the central security teams. This helps ensure a shared understanding of responsibilities and fosters a collaborative environment focused on achieving specific security goals.

4. Prioritize engagement. Recognizing these security champions often volunteer their time on top of existing demanding roles, AWS actively focuses on maintaining their engagement. This understanding that sustained participation is crucial for long-term success differentiates their approach.

5. Measure for continuous improvement. AWS emphasizes the importance of quantifiable metrics to assess the Security Guardians program's effectiveness. This data-driven approach not only validates the program's impact on security outcomes but also provides valuable insights for ongoing refinement and optimization.

The result? AWS is demonstrating that by strategically distributing security responsibility and fostering a culture of security ownership, organizations can achieve both rapid innovation and robust security. This initiative is more than scaling a security team; iit's about scaling a culture of security across the entire organization—making security a top priority and a shared responsibility in how the business operates.

The AI edge

Supercharging security and productivity

The undeniable momentum of AI is translating into tangible benefits for security teams. AI augmentation is effectively tackling labor-intensive and expertise-draining security functions, leading to significant operational improvements. This powerful upside is driving aggressive adoption strategies, with organizational leaders forecasting:

- A substantial 50% increase in AI augmentation adoption within the next three years
- A remarkable 63% expansion in generative AI security capabilities
- A significant 45% rise in workflow automation and orchestration.

Crucially, this technological shift is enhancing, not eliminating, human intellect. A strong majority of executives (65%) agree that AI and automation are cultivating a more productive landscape for their IT and security professionals, and close to two-thirds (62%) are already realizing considerable value from embedded AI functionalities.

The path forward

Reinventing the operating model

The next three years will be defined by fundamentally new ways of working. How do organizations integrate digital labor into everyday operations? How do they work with AI agents as cooperators, collaborators, and tireless 24x7 extensions of themselves? What are the security implications of machine identities proliferating throughout environments?

While this vision remains a work-in-progress, organizations that effectively answer these questions—integrating them into strategy and operations—can create significant competitive advantages.

For CISOs and security leaders, success requires not just new technology but a fundamental reimagining of the security operating model itself. While there's no consensus yet on the ideal approach, the race for competitive advantage in AI-powered security has already begun.

> "As leaders, we need to ask ourselves what if we are truly looking at 10X the most advanced AI capabilities we have today—and that's generally available to us. What does that mean for our business? For our operations? For the adversaries we are up against? It's our job as leaders to be thinking about this. We can use this inflection point to become even more useful to the business, to help us change and grow for the next set of opportunities.
>
> Koos Lodewijkx, Vice President, CISO, IBM

FIGURE 5

Organizations are in transition, and operating models are showing their age

Currently, there's little correlation between strategy and performance. As AI moves to the center of operations, operating models must be modernized to reflect fundamentally new ways of working.

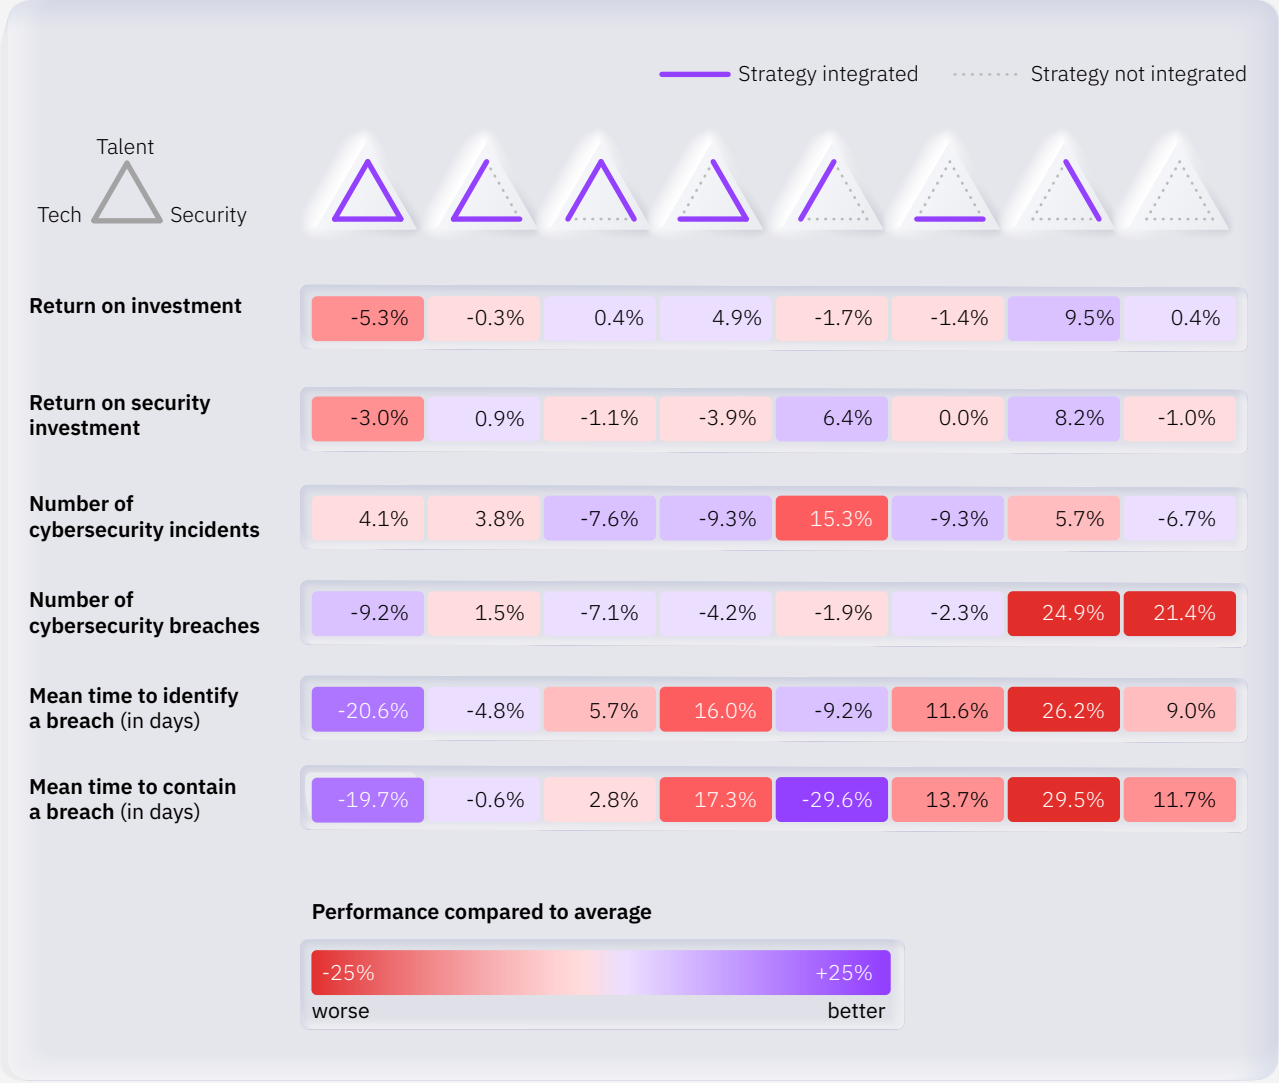

The columns represent the extent of integration between the organization's talent, technology, and security strategies. For the questions below, there are a total of 8 possible response combinations, each of which is represented by 1 column. Qs: Our tech and talent strategies are well integrated [COMBINE 3, Agree & 4, Strongly agree]. Our tech and security strategies are well integrated [COMBINE 3, Agree & 4, Strongly agree]. Our talent and security strategies are well integrated [COMBINE 3, Agree & 4, Strongly agree]. The rows indicate the organization's performance on the following measures. Qs: Over the past two years (2023 to 2024), what is your organization's ROI on its cyber risk and cybersecurity investments? Over the past two years (2023 to 2024), what is your organization's return on security investment (ROSI)? Estimate your organization's performance in the last 12 months for the following metrics—Number of cybersecurity incidents. Estimate your organization's performance in the last 12 months for the following metrics—Number of cybersecurity breaches. Estimate your organization's performance in the last 12 months for the following metric—Mean time to identify a breach (MTTI) in days. Estimate your organization's performance in the last 12 months for the following metrics—Mean time to contain a breach (MTTC) in days

Action guide

Security accelerated

Build momentum by modernizing your operating model for AI

AI-first transformation leaders (CISOs, CTOs, CIOs)

Build the foundation. Improve hygiene, improve effectiveness, improve ROI, improve resilience. Because agentic AI will require humans and AI working in a more cooperative, more interdependent way, leaders should look to improve alignment across their tech, security, and talent strategies.

Excellence everywhere. Make your work environment world-class. Rationalize your technology and security toolsets to improve efficiency, visibility, and governance. Standardize around a few core platforms with comprehensive data integration and workflow orchestration capabilities. Pair high-skill, subject-matter experts with comprehensive AI and automation services across the IT/IS lifecycle.

Hyphenate. Move beyond established roles and cultivate a team of business-minded security specialists. Think "cyber curators" and "AI forensics strategists"—individuals who combine technical expertise with sharp cyber instincts. These are the people who will guide, interpret, and challenge increasingly autonomous AI systems. It's not just about filling seats; it's about building a team that can anticipate risks and outpace threat actors' rapidly evolving AI and automation tactics.

Leaders critical to AI-first success (CEOs, CFOs, COOs, CHROs)

Make AI your operating model. Create a balanced scorecard to assess the level of alignment between your talent, technology, and security strategies. Articulate high-impact user stories that span all three.

Shift from top down to middle out. Rather than route every decision through a centralized security function, distribute security capabilities throughout the organization. Then champion security as a core value. Make sure AI advocacy and enthusiasm are shared by executives and employees. Ensure that AI benefits are not flowing to some but not others. Start with culture and work backward. Avoid prioritizing efficiency to the detriment of your work environment.

Turn guardians into champions. Forge an elite team of cyber curators, innovators, and creators—the core of your AI security strategy—and land these in high-impact business and operations teams. Your frontline against the escalating AI threat landscape extends beyond technology to a carefully constructed team.

Double-down on expertise. Consider where human skills, expertise, and judgment make the greatest impact. Use AI augmentation to enhance performance. Assess effort and expertise-intensive tasks. Shift to automation where it makes sense to reduce errors, improve responsiveness, and recover cycle time.

Autonomous AI as the enterprise guardian

AI everywhere, all at once

Organizations are being outmaneuvered in the AI race by fighting the wrong battle. While threat actors deploy AI to craft sophisticated deepfakes and bypass authentication, enterprises remain fragmented—security teams operate in silos, blind to the bigger picture.

The vulnerability isn't technical—it's organizational. Threat actors see entire enterprises as one connected system. Organizations see functions and responsibilities. Attackers deploy coordinated tactics that evolve from a routine privilege escalation to lateral network movement. Defenders analyze data in isolation, often failing to appreciate its significance because they don't see corresponding side-channel tactics or because the real-time context is missing.

This fundamental misalignment leaves enterprises exposed. Current access controls fail because attackers exploit the gaps between business, operations, technology, and security— gaps that have been baked into how the organization operates.

The solution requires more than better tools. Organizations need an AI-centric security operating model that mirrors how attackers think: holistically. When IT and security functions truly integrate, AI becomes transformative—processing threats across time and context, not just individual incidents.

Trust powers institutions, but the mechanisms that create trust are under siege. Organizations must rebuild them with the same sophistication adversaries use to undermine them.

The good news: a majority of organizations (59%) are already pursuing deep integration across security, infrastructure, and application operations, with a significant 67% anticipating full consolidation of IT, application, and security observability within the next three years.

> "Defenders today have some interesting advantages. First, they have all the data on their environment. They see all those connections, all the cloud trail logs for all the API calls. Attackers are dying for that data. Instead, they're coming into a black box. Second, generative AI creates some huge opportunities for defenders. It can be used in so many ways—to secure your workloads, your accounts, your applications, your data. This generation of genAI is going to accelerate all those capabilities."
>
> Chris Betz, Chief Information Security Officer, AWS

FIGURE 6

AI emerges as a differentiator

Approximately one in three organizations (30%) has established a strong AI-first foundation. These organizations are best prepared to succeed with their AI transformation efforts.

Source: IBM Institute for Business Value data analysis. The rows represent the organization's performance across 4 dimensions. The performance for each dimension is calculated using an index score, with higher scores indicating greater proficiency across that dimension (See "Research methodology" on p.29 for details). The columns represent a distinct grouping of organization's based on their responses to the items above. Organizations on the left show the least amount of progress with their AI transformation efforts. Organizations on the right show the greatest.

Three years, three steps

Over the next three years, organizations will learn to scale AI while learning to trust it. Three foundational use cases reveal how this transformation unfolds: IT/IS observability, AIOps, and autonomous cybersecurity.

These capabilities unlock advanced use cases across the operations lifecycle. Organizations such as AWS have automated security policies so extensively they've eliminated traditional, highly centralized security operations centers.[19] This evolution required moving from discrete tools to mutual accountability to a culture where security and success reinforce each other—the security flywheel in action.

The future lies in AI that sparks the imagination. For example, using foundation models to learn the sequence of threat actor behaviors; using log files to identify anomalous activity, flagging API calls with unconventional syntax, or running spot checks when code check-ins aren't consistent with the expected order of operations —the sum total of which transforms threat detection and response capabilities into a new, far more proactive way of working.

As we move beyond scripted automation and orchestration, our AI support teams will become extensions of ourselves. With greater visibility and reach, we will shift our focus to learning—developing new, more efficient ways to connect our risk posture to security policies to business outcomes. New security solutions will shift from capabilities to characteristics, such as:

1. Self-correcting	AI captures and collates data to resolve issues before they become threats. Like a smart thermostat, it proactively identifies cloud misconfigurations and unusual traffic patterns, automatically adjusting settings while reducing alert fatigue for human experts.
2. Self-healing	When disruptions occur—DDoS attacks, data breaches—AI orchestrate sautomated recovery. It reroutes traffic, isolates affected systems, and restores services with minimal human intervention, freeing specialists to tackle novel threats.
3. Self-directing	AI learns from incidents and threat intelligence to autonomously refine policies, update defenses, and anticipate attack vectors. It acts as an intelligent guardian while humans shift to strategic oversight, supervisory learning support, and innovation activities.[20]

Perspective

How AI-first transformation pays for itself

Estimating potential savings

Let's take an example organization with $20B in annual revenue and the following budget outlay:

IT budget as % of organization's annual revenue	IS budget as % of organization's IT budget	Average personnel costs (as % of total cybersecurity budget)
7%	8%	38%
$1.4 B	$112 M	$42.56 M

Next, let's estimate their effort and budget savings when using advanced AI and automation

	Low	Average	High
Estimated effort saved	32%	53%	74%
Estimated budget freed for augmented workforce re-investment			
Scenarios as % current cybersecurity budget in USD			
Moderate use (supporting 50% of cybersecurity workloads)	**6%** $6,809,600	**10%** $11,278,400	**14%** $15,747,200
Extensive use (supporting 75% of cybersecurity workloads)	**9%** $10,214,400	**15%** $16,917,600	**21%** $23,620,800
Comprehensive use (supporting 100% of cybersecurity workloads)	**12%** $13,619,200	**20%** $22,556,800	**28%** $31,6494,400

Source: IBM Institute for Business Value data analysis.
Note: the above scenario is based on an organization with $20 BN in annual revenue. Model inputs and assumptions are based on responses to specific survey questions, including firmographic questions. Please see "Research methodology" on page 29 for more details.

Action guide

AI overdrive

Full throttle in the “Run” stage

AI-first transformation leaders (CISOs, CTOs, CIOs)

Make data your strength. Your institutional data is your advantage—both with competitors and adversaries. Think historical ticket data, policy configurations, design documents. Make the most of retrieval augmented generation (RAG) capabilities to turn these into AI model inputs. Shift from running playbooks to supervising AI learning.

Turn AI risks into AI lessons. Prepare for an AI-meets-AI world. Get ahead of your adversary doppelganger by training your AI models for operational resilience. Assess your IT, OT, and IS footprint to determine how you can make baseline risk management, threat identification, and incident response automatic—that is, fully automated and autonomous.

Find yourself. Focus on making your security operations more self-regulating and self-correcting through IT/IS observability and resilience capabilities. Make them more self-sustaining and self-healing through AI operations capabilities. Make them more self-directed, outcome oriented, and autonomous by freeing specialists to focus on higher-value supervisory tasks, where context and complexity require interpretation and judgment.

Follow the money. Run the numbers to understand how you will fund your AI-first transformation program. Use advanced AI modalities such as AI agents and multiagent systems to recover cycle time and expand capacity. Shift personnel expenditures from effort and expertise-intensive tasks to new AI-augmented roles.

Leaders critical to AI-first success (CEOs, CFOs, COOs, CHROs)

Turn security personas into business personas. Think of a day in the life of various personas in your org: what they want to achieve, what data sources and tools they use, what outcomes they create and deliver. Can security be improved in ways that improve the user experience? Can friction be reduced with AI agents? Can elapsed time be reduced via autonomous decision-making?

All things are difficult before they are easy. Work with your IT/IS counterparts to shift workloads to AI-moderated solutions in stages—providing feedback and improving outcomes over time. The mark of success is when your team embraces AI as an essential complement, companion, and extension of their own work.

Take smarter risks. Consider AI, cloud, and security investments together—as a whole. The key to risk management and better resilience is thinking holistically. While cloud and AI are evolving together, too often security lags. The real advantage comes when AI, cloud, and security are fused together, from the start.

The power of starting where you are

The AI revolution doesn't wait for perfect plans or complete transformations. It is a work-in-progress and will be for the foreseeable future. Every month spent deliberating is a month competitors pull ahead—and threat actors sharpen their tools.

The organizations that emerge as AI-resilient won't be the ones that had all the answers from day one. They'll be the ones that started moving, failed fast, learned faster, and adapted better.

This isn't about perfection—it's about momentum. The "crawl, walk, run" journey isn't linear, and it doesn't require starting from scratch. Start from where you are, with what you have, right now.

Over the next 36 months, the cybersecurity landscape will be redrawn entirely. The question isn't whether AI will transform how your organization defends itself—it's whether you'll be driving that transformation or scrambling to catch up. The enterprises that thrive won't just be AI-enabled; they'll be AI-native, with security woven into the fabric of every automated decision, every synthetic data set, every autonomous action.

Your future awaits. The question is: will you crawl, walk, or run toward it?

> "I think the number one most important thing is that when I do my job best, I make it faster and easier for the business to achieve their goals in a trustworthy way. That means I need to go on that journey with them. It's not a separate security journey. It's a security journey with the business."
>
> Chris Betz, Chief Information Security Officer, AWS

About
the authors

Leonard Bernstein

Global Cybersecurity Leader, AWS
linkedin.com/in/leonardbernstein/

Leonard is Global Cybersecurity Leader at AWS, helping enterprises strengthen cyber resilience, modernize security operations, and transform security into a foundation for business scale, growth, and innovation.

Srinivas Tummalapenta

CTO, IBM Security Services
IBM Distinguished Engineer
linkedin.com/in/srinivastummalapenta/

As CTO of IBM Security Services, Srini partners with product partners, solution architects, and strategy leaders to define and deliver security solutions across the NIST Cybersecurity Framework.

Abhi Chakravorty

Partner, Cloud & Infrastructure Security Offering Leader
IBM Consulting
linkedin.com/in/achakrav/

Abhi leads strategy and offering management for IBM's cloud & infrastructure security services, helping global clients transform their cloud platforms and modernize their networks.

Michael Massimi

Global Principal, Cloud Security Services for AWS
IBM Consulting
linkedin.com/in/michael-massimi-b4b3b21/

Michael is responsible for managing the worldwide strategy and execution for IBM Security Services digital transformation and modernization programs on Amazon Web Services (AWS).

Gerald Parham

Global Research Leader, Security & CIO
IBM Institute for Business Value
linkedin.com/in/gerryparham

Gerry's research insights have appeared in publications such as *The Wall Street Journal, Forbes, CIO, Cyber,* and *Infosecurity Magazine*. His papers have been recognized as among the leading examples of thought leadership in the world.

Contributors

We would like to thank the following individuals for their significant contributions to these materials:

Chris Betz
Chief Information Security Officer, AWS

Koos Lodewijkx
Chief Information Security Officer, IBM

Dimple Ahluwalia
Global Offering Leader, CyberDefend
IBM Consulting

Mark Hughes
Global Managing Partner, IBM Security Services
IBM Consulting

Liam Cleaver
Research Director, Open Innovation and Ecosystems,
IBM Institute for Business Value

Cathy Fillare
Global Research Leader, HR and Talent Transformation,
IBM Institute for Business Value

Kristine Rodriguez
Editor-in-Chief, IBM Institute for Business Value

Heba Nashaat
Business analytics consultant, IBM Institute for Business Value Research Hub

Sara Aboulhosn
Associate Creative Director, IBM Institute for Business Value

Angela Finley
Design Lead, IBM Institute for Business Value

Andrew Womack
Creative Director, IBM Institute for Business Value

Bhuvana Chandar
Global Partner Marketing Leader, AWS Strategic Partnership, IBM Consulting

Research methodology

To understand how new AI modalities are impacting cybersecurity operations, the IBM Institute for Business Value, in collaboration with Oxford Economics, surveyed 1,013 C-level executives across security, technology, operations, and business domains in Q1 and Q2 2025. This included 250 CISOs, 300 CTOs and CIOs, and over 450 CEOs, CFOs, COOs, CHROs (or their functional equivalents).

Respondents were screened for familiarity with their organization's cybersecurity strategy and operations and their organization's use of AI. Only respondents who indicated they were "Very familiar" completed the survey.

Respondent organizations have an annual revenue between $500 million USD to more than $20 billion USD. Respondents represented the following sectors/industries: banking and financial markets, government/public sector, energy and resources (including oil and gas), industrial/manufacturing (including automotive), telecommunications, retail and consumer goods, and transportation (including airlines).

Respondents came from the following countries: Australia, Brazil, Canada, China, France, Germany, India, Italy, Japan, Mexico, Saudi Arabia, Singapore, South Korea, Spain, United Arab Emirates, United Kingdom, United States.

Performance was assessed across financial and operational factors, while transformation practices were evaluated across multiple domains including business operations, security, technology, and talent. To uncover data patterns and relationships, we deployed various analytical techniques including correlation analysis, regression modelling, trend analysis, and K-means clustering. These methods helped identify key associations between operational practices and business outcomes.

For operating model analysis, we compared responses to three questions assessing the degree of integration across talent, tech, and security strategies. This resulted in eight possible response combinations. Each of these was then assessed against six performance measures: Two financial performance measures (ROI and ROSI) and four operational performance measures (number of cybersecurity incidents, number of cybersecurity breaches, mean-time-to-identify [MTTI] a breach in days, and mean-time-to-contain [MTTC] a breach in days).

To understand the organization's AI-transformation progress, we analyzed IT/IS organizations across four dimensions: security awareness, behavior, and cultural practices; AI workforce activation practices; AI operations maturity; and the degree of AI integration across the organization's IT/IS portfolio. Organizations were then scored along each dimension and grouped into a cluster solution for subsequent analysis. Organizations fall into four distinct groups. These are described as stages to represent the organization's relative progress toward enabling an AI-centric operating model. It should be noted, this depiction as discrete stages represents a snapshot of AI transformation programs at this current point in time. Organizations wouldn't necessarily proceed sequentially from one stage to the next and transformation objectives (and relative progress) are likely to change as organizations mature.

The cost savings scenario models a $20B-annual-revenue organization. The following items were used as inputs for the financial modeling scenario: average IT (technology) budget as a percentage of organizational revenue, average IS (cybersecurity) budget as a percentage of IT budget, average personnel costs as a percentage of current IS budget. Respondents were asked to estimate the effort saved from three future state use cases: IT/IS observability and resilience, AI operations, and autonomous cybersecurity. These resulted in three outputs: the average for all respondents, a low scenario (-2 standard deviations from average), and a high scenario (+2 standard deviations from average). Collectively, the range of estimates covers more than 95% of respondents. Finally, these inputs were modelled across three scenarios, representing moderate (50%) adoption of advanced AI capabilities, extensive (75%) adoption of advanced AI capabilities, and comprehensive (100%) adoption of advanced AI capabilities. The model outputs are represented both as a percentage of total cybersecurity budget and in financial terms (in USD).

About Research Insights

Research Insights are fact-based strategic insights for business executives on critical public- and private-sector issues. They are based on findings from analysis of our own primary research studies. For more information, contact the IBM Institute for Business Value at iibv@us.ibm.com.

IBM Institute for Business Value

For two decades, the IBM Institute for Business Value has served as the thought leadership think tank for IBM. What inspires us is producing research-backed, technology-informed strategic insights that help leaders make smarter business decisions.

From our unique position at the intersection of business, technology, and society, we survey, interview, and engage with thousands of executives, consumers, and experts each year, synthesizing their perspectives into credible, inspiring, and actionable insights.

To stay connected and informed, sign up to receive IBV's email newsletter at ibm.com/ibv. You can also find us on LinkedIn at ibm.co/ibv-linkedin.

The right partner for a changing world

At IBM, we collaborate with our clients, bringing together business insight, advanced research, and technology to give them a distinct advantage in today's rapidly changing environment.

About AWS

For over 15 years, Amazon Web Services has been the world's most comprehensive and broadly adopted cloud offering. Today, we serve millions of customers, from the fastest growing startups to the largest enterprises, across a myriad of industries in practically every corner of the globe. We've had the opportunity to help these customers grow their businesses through digital transformation efforts enabled by the cloud. In doing so, we have worked closely with the C-suite, providing a unique vantage point to see the diverse ways executives approach digital transformation—the distinct thought processes across C-suite roles, their attitudes and priorities, obstacles to progress, and best practices that have resulted in the most success. For more information, please visit: aws.amazon.com

About IBM Security

IBM Security Services works with you to help protect your business with an advanced and integrated portfolio of enterprise cybersecurity solutions and services infused with AI. Our modern approach to security strategy uses zero-trust principles to help you thrive in the face of uncertainty and cyberthreats. For more information, please visit: ibm.com/services/security

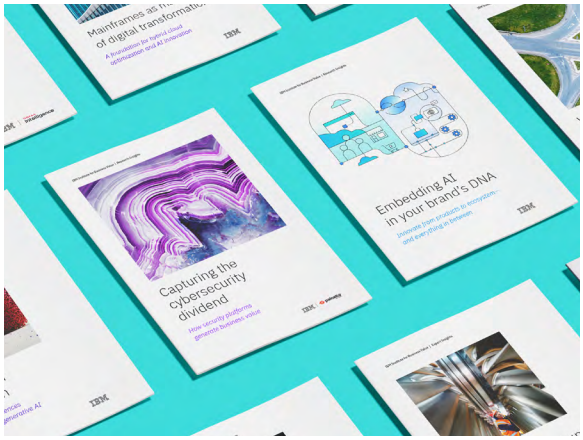

Subscribe to our IdeaWatch newsletter

Just the insights. At your fingertips. Delivered monthly.

Brought to you by the IBM Institute for Business Value, ranked #1 in thought leadership quality by Source Global Research for the second consecutive year.

Research-based thought leadership insights, data, and analysis to help you make smarter business decisions and more informed technology investments.

Subscribe now: ibm.co/ideawatch

Related Reports

Securing generative AI: What matters now

ibm.com/thought-leadership/institute-business-value/en-us/report/securing-generative-ai

6 blind spots tech leaders must reveal

ibm.com/thought-leadership/institute-business-value/en-us/report/cxo

Reimagine human potential in the generative AI era

ibm.com/thought-leadership/institute-business-value/en-us/report/human-potential-genai

Notes and sources

1 Elgan, Mike. “Is AI saving jobs… or taking them?,” IBM. October 8, 2024. https://www.ibm.com/think/insights/is-ai-saving-jobs-or-taking-them

2 Rattner, Nate, and Deepa Seetharaman. “Here’s How Big the AI Revolution Really Is, in Four Charts,” The *Wall Street Journal*. April 23, 2025. https://www.wsj.com/tech/ai/ai-boom-companies-afb8c7e0

3 Bousquette, Isabelle. “Why Moderna merged its tech and HR departments,” *The Wall Street Journal*. May 12, 2025. https://www.wsj.com/articles/why-moderna-merged-its-tech-and-hr-departments-95318c2a; “Shadow AI breach risks escalate,” *CybersecurityHQ Newsletter*, https://newsletter.cybersecurityhq.com/p/shadow-ai-breach-risks-escalate

4 Hughes, Mark, and Karim Temsamani. *Capturing the cybersecurity dividend: How security platforms generate business value.* IBM Institute for Business Value in partnership with Palo Alto Networks. January 2025. https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/unified-cybersecurity-platform

5 Boyle, John D. “The Hidden Cybersecurity Crisis: How GenAI is Fueling the Growth of Unchecked Non-Human Identities”. Security Boulevard. February 15, 2025. https://securityboulevard.com/2025/02/the-hidden-cybersecurity-crisis-how-genai-is-fueling-the-growth-of-unchecked-non-human-identities/

6 Stanley, Edward. “Investors are still underestimating the long-term impact of AI.” *The Financial Times.* March 31, 2024. https://www.ft.com/content/a171bd7c-7f70-4145-acd5-beeb0b8da732

7 Moore-Colyer, Roland. “AI can handle tasks twice as complex every few months. What does this exponential growth mean for how we use it?” *Live Science*. April 27, 2025. https://www.livescience.com/technology/artificial-intelligence/ai-can-handle-tasks-twice-as-complex-every-few-months-what-does-this-exponential-growth-mean-for-how-we-use-it

8 Rodgers, Clarke, Moumita Saha, Dimple Ahluwalia, Kevin Skapinetz, and Gerald Parham. *Securing generative AI: What matters now*, IBM Institute for Business Value. May 2024. https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/securing-generative-ai

9 Tully, Tim, Jeff Redfern, and Derek Xiao with Claude Sonnet 3.5. “2024: The State of Generative AI in the Enterprise.” Menlo Ventures, November 20, 2024. https://menlovc.com/2024-the-state-of-generative-ai-in-the-enterprise/

10 “CISA Artificial Intelligence Use Cases.” CISA. Accessed June 4, 2025. https://www.cisa.gov/ai/cisa-use-cases

11 Law, Marcus. “CIO AI Spending to Rise, but Many Concerned by App Sprawl.” *Technology Magazine.* February 1, 2024. https://technologymagazine.com/articles/cio-ai-spend-set-to-rise; Security Staff. “Report shows nearly 600% annual growth in vulnerable cloud attack surface.” *Security Magazine.* May 1, 2023. https://www.securitymagazine.com/articles/99276-report-shows-nearly-600-annualgrowth-in-vulnerable-cloud-attack-surface; JupiterOne Research. “State of Cyber Assets Report (2023).” JupiterOne. Accessed June 4, 2025. https://info.jupiterone.com/hubfs/SCAR%202023/jupiterone_2023-state-of-cyber-assets-report_scar.pdf

12 Wiz Experts Team. “7 Serious AI Security Risks and How to Mitigate Them.” Wiz. March 28, 2025. https://www.wiz.io/academy/ai-security-risks

13 Rodgers, Clarke, Moumita Saha, Dimple Ahluwalia, Kevin Skapinetz, and Gerald Parham. *Securing generative AI: What matters now*, IBM Institute for Business Value. May 2024. https://www.ibm.com/thought-leadership/institute-business-value/en-us/report/securing-generative-ai

14 2025 Skills and Talent study, IBM Institute for Business Value. February 27, 2025. n=2690. Unpublished data.

15 “GovTech Collaborates with AWS Generative AI Innovation Center to Scale Generative AI Adoption Across Public Sector Organizations.” AWS. Accessed June 4, 2025. https://aws.amazon.com/solutions/case-studies/govtech/

16 Morgan, Steve. “Cybersecurity Jobs Report: 3.5 Million Unfilled Positions In 2025.” *Cybercrime Magazine*. April 14, 2023. https://cybersecurityventures.com/jobs/

17 “IBM Cyber Campus.” IBM. Accessed June 4, 2025. https://www.ibm.com/services/consulting-cyber-campus

18 Beaumont, Mitch, and Ana Malhotra. “How to build a Security Guardians program to distribute security ownership.” AWS. https://aws.amazon.com/blogs/security/how-to-build-your-own-security-guardians-program/

19 Swinhoe, Dan. “Automating security at AWS: How Amazon Web Services operates with no SOC,” CSO. com. November 1, 2018. https://www.csoonline.com/article/566501/automating-securityat-aws-how-amazon-web-services-operates-withno-soc.html

20 Boston Consulting Group (BCG). “AI Oversight Needs to Be Designed, Not Delegated,”Up Next from BCG blog, LinkedIn. April 18, 2025. https://www.linkedin.com/pulse/ai-oversightneeds-designed-delegated-boston-consulting-group-8daqe/; Mills, Steven, Noah Broestl, and Anne Kleppe. “You Won’t Get GenAI Right If You Get Human Oversight Wrong.” BCG. Accessed June 4, 2025. https://www.bcg.com/publications/2025/wont-get-gen-ai-right-if-human-oversight-wrong

IBM Corporation
New Orchard Road
Armonk, NY 10504

Produced in the United States of America | June 2025

131cf8a39db32790-USEN-03