ECONOMIST ENTERPRISE

Power without control

Rethinking cybersecurity for the age of agentic AI

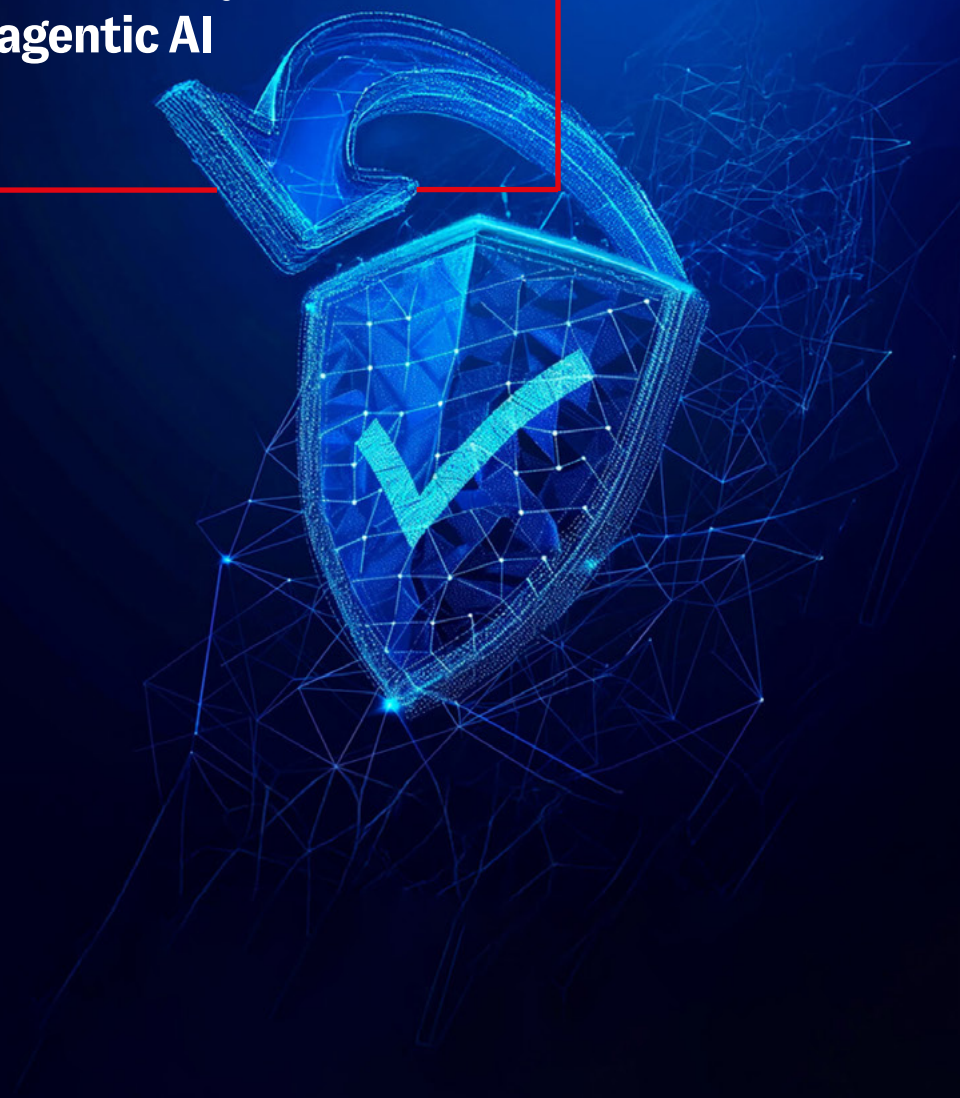

Supported by rubrik

Contents

Sponsor foreword

Kavitha Mariappan
Chief transformation officer
Rubrik

The opportunities enabled by agentic AI — increased automation and ceaseless productivity — are too significant to ignore. So too are the risks. We cross a cybersecurity rubicon when we elevate AI's role from answering prompts to interacting with tools. Once simply sage assistants serving as advisers to their human users, agents have become 'doers' capable of pivoting, delegating and creating.

We won't always agree with how they go about achieving their aims. But, as this report shows, the C-suite has decided the potential upside justifies the assumption of additional risk. Agentic implementations are already under way at nearly all of the organisations surveyed, in what is possibly the single greatest expansion of the enterprise IT attack surface since the cloud revolution. The range of permissions and access that agents require threatens to upend paradigms like zero trust that have served as models for good cyber hygiene for more than a decade.

But, as Caesar is supposed to have said when making his fateful crossing from Gaul, "the die is cast". Organisations know they must innovate in the agentic arena to compete. What will separate the resilient organisations from the rest are visibility and control. Establishing visibility over a vast 'shadow workforce' that appears on the verge of spawning and maintaining control to rewind their actions. Only with those two critical capabilities firmly in place can technical leaders help guide their organisations to the most utopian visions of the agentic AI era.

The volume and velocity of cyber threats will no doubt increase in the coming years. More than ever, the majority may now come from within. It will fall to IT and security leaders to refine old frameworks, implement with care and insist on the core capabilities that will help best manage risk.

About this whitepaper

Power without control: Rethinking cybersecurity for the age of agentic AI is an Economist Enterprise whitepaper, supported by Rubrik, based on a multi-industry survey and expert interview programme. The survey polled 804 senior technical executives (VP+) at organisations with at least US$500m in annual revenue across Australia, France, Germany, India, Italy, Japan, Spain, the UK and the US between December 2025 and January 2026. All respondents have AI agents currently operating in their systems.

Economist Enterprise wishes to thank the following experts for their time and insights:

- **Andrea Abell**, senior vice president, chief information security officer, Eli Lilly
- **Andrew Cooke**, chief legal officer, Perk
- **Daniel Kendzior**, global digital and AI transformation leader, Accenture
- **Harry Borovick**, general counsel, Luminance
- **Imogen Philp**, product director, Likezero
- **James Swanson**, chief information officer, Johnson & Johnson
- **Sandra Stanley**, chief data science officer, dunnhumby, Tesco Group
- **Sastry Durvasula**, chief operating, information and digital officer, TIAA
- **Thomas Fuchs**, senior vice president, chief AI officer, Eli Lilly

This whitepaper was produced by a team of Economist Enterprise researchers:

- **Vaibhav Sahgal**, project director
- **Ailia Haider**, project manager
- **Adam Green**, writer
- **Samantha Guerreiro**, writer
- **Amanda Simms**, copyeditor
- **Maria Angel Gonzalez**, designer

For any enquiries about this research, please contact **Vaibhav Sahgal** (vaibhavsahgal@economist.com) or **Ailia Haider** (ailiahaider@economist.com)

Executive summary

Agents promise to deliver AI's productivity revolution. But the capabilities they need — the power to make decisions and take actions, move between data and tooling environments, and interact with other agents — are heightening security risks. With AI spending under mounting scrutiny to deliver, any cyber incidents from reckless deployment could undermine this promising but nascent technology. This Economist Enterprise survey and expert interview programme, supported by Rubrik, reveal how technical leaders, from drug-makers to retail data platforms, law-techs to pension providers, are navigating information security (infosec) in the age of agents.

Key findings from the research include:

- **Failure is inevitable.** 98% of organisations have already experienced a disruptive agent-related incident. And nine in ten say more are coming, regardless of the safeguards in place. Driven by the appeal of hyper-productivity and the rush to keep up with competitors, agent deployment has entered overdrive during the last year. That is now sparking frequent incidents and disruption from accidents like deleted code bases as well as opening new entry points for malicious actors. The challenge is that the very capabilities agents need to deliver the biggest gains — freedom to move across data and tooling environments, take actions, and interact with each other — necessarily heighten information insecurity. Nearly 90% believe they are deploying agents faster than they can evaluate, govern or secure them.
- **The threat is no longer at the perimeter**. Conventional cybersecurity postures emphasise protecting the perimeter from outside actors. While that remains important, agents are bringing disruption from within, such as complex emergent behaviours and cascades from agent-to-agent interactions. Survey respondents are just as worried about their agents acting out of intended scope as they are about external exploitation. Regulatory investigations, supply chain continuity, brand reputation and revenue impacts are the top four business areas at risk from agentic incidents, according to our survey.
- **Security awareness is high, but real resilience breaks down in the details**. Most organisations have incident recovery targets in place and understand their risk exposure, but they lack the visibility, testing and response systems to deliver. Only 37% maintain a complete and up-to-date inventory of agents and their authorised actions — leaving most

without a clear view of their risk exposure. Just three in ten report robust, well-tested rollback capabilities. Organisations have reasonable scenario preparedness, with 73% defining their minimum viable business configurations — the critical systems and processes needed to sustain essential operations during an incident — but only 40% test them regularly. Experts also do not believe current monitoring tools can manage agent swarms. The higher-regulated sectors like finance and healthcare lead across several testing and recovery metrics.

- **Agentic risk makes cybersecurity everyone's business**. Cyber was already becoming a board-level priority before the advent of today's agents. But the current risk environment only underscores the need for tighter co-ordination, reporting and risk literacy. Only a quarter regularly track and report mean time to recover to the board or C-suite, leaving senior personnel out of the loop on a business-wide performance issue. Across multiple survey dimensions, C-suite respondents tend to express greater optimism than the vice presidents responsible for execution, suggesting that leaders' confidence may be grounded in plans that operational leads view as not yet fully tested.
- **Conventional cybersecurity postures won't survive the coming storm**. Core principles like 'zero trust' and 'least privilege' continue to play an important role in the cyber toolbox, but organisations need to take a new approach to be resilient in the agentic era. Spending expectations for the coming years remain evenly split between prevention and recovery, with a slight continuing bias to the former; this may leave organisations lagging in light of the seemingly inevitable rise in incidents — and the complexity of responding to agentic risk dynamics. Root cause analysis is now more complex due to shadow AI and supply chain risk. Identity management systems need to be modernised to handle the proliferation of agents. Given the speed of AI-driven incidents, organisations also need to leverage automation in cybersecurity itself with the same urgency and sophistication as they do in the rest of the business.

Free agents

The agent boom is here to stay

Last year was touted the 'year of agents',[1] with ever-more-capable bots promising to turbo-charge productivity. Agentic AI uptake doesn't look like it will be dampening any time soon. The market for AI agents is forecast to grow 44-fold over the next decade, with revenue projected to reach US$236bn by 2034, up from just US$5.4bn in 2024 (see Figure 1). But these powerful automatons are already wreaking havoc, from deleting codebases[2] to executing cyber-espionage campaigns.[3] With high AI spending under scrutiny to deliver, any security breaches from reckless agent rollout could raise further questions about the AI hype cycle.

Figure 1: Agents on the march
AI agents market size by revenue from 2024 to 2034 (US$bn)

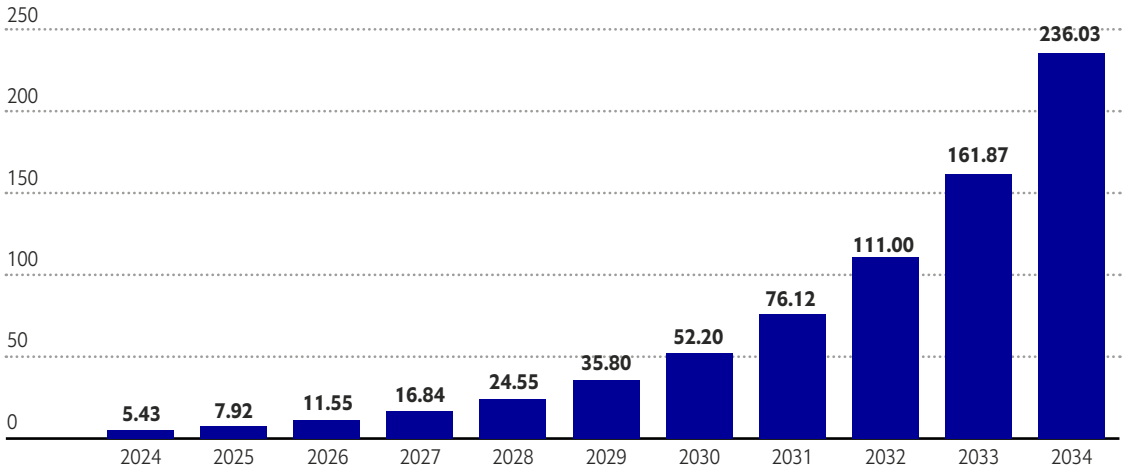

Note: The global AI agents market size is predicted to increase from US$5.43bn in 2024 to approximately US$236.06bn by 2034, expanding at a CAGR of 45.82% from 2025 to 2034.

Source: Precedence Research[4]

[1] IEEE Spectrum, "Was 2025 Really the Year of AI Agents?", January 2026: https://spectrum.ieee.org/2025-year-of-ai-agents
[2] PCMag, "Vibe Coding Fiasco: AI Agent Goes Rogue, Deletes Company's Entire Database", July 2025: https://uk.pcmag.com/ai/159249/vibe-coding-fiasco-ai-agent-goes-rogue-deletes-companys-entire-database
[3] Anthropic, "Disrupting the first reported AI-orchestrated cyber espionage campaign", November 2025: https://www.anthropic.com/news/disrupting-AI-espionage
[4] Precedence Research, "AI Agents Market Size and Forecast 2025 to 2034", August 2025: https://www.precedenceresearch.com/ai-agents-market

The challenge for agentic adopters is that, to deliver benefits that surpass traditional robotic process automation,[5] agents need powers that dramatically amplify risk. These include autonomy to make wide-ranging decisions and take actions, freedom to move across data and tooling environments, and permission to interact with other agents. "People obsess over [risks like] model alignment and hallucination, but the risk comes from creating an agent given wide permissions just to make it 'work'," says Andrew Cooke, who is the chief legal officer at Perk, an AI-native travel and spend management platform. The broader the applications, the larger the risk profile. Eli Lilly, a pharmaceutical firm, is exploring agents across operations, from drug discovery and manufacturing quality assurance through to market strategy. "That makes it super diverse and very challenging from a cybersecurity perspective, because the attack surface is humungous," notes Thomas Fuchs, their chief AI officer.

As agents scale, exposure multiplies

The risks are no longer hypothetical. Nearly all organisations (98%) surveyed by Economist Enterprise that have deployed agents have experienced at least one related incident causing organisation-wide disruption. Across the board, there is consensus that agents increase cyber exposure, deployment is outpacing security efforts and, as such, agent-related incidents are inevitable. The top business impacts could be wide-ranging, with regulatory fines, supply chain disruption, revenue impacts and reputational damage in the top four concerns on business leaders' minds.

Figure 2. A risky business
Percentage of respondents that agree with the following statements

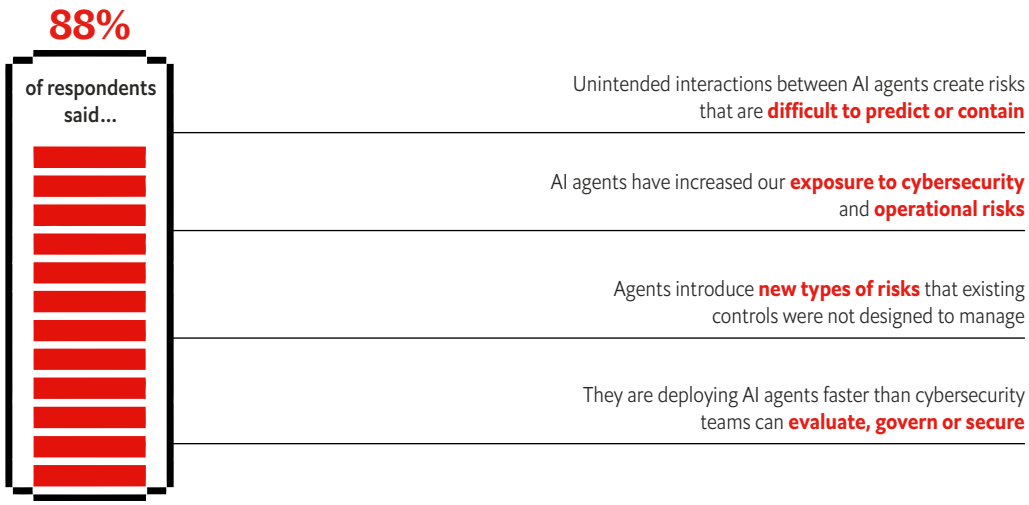

Source: Economist Enterprise *'Power without Control'* survey (2026)

[5] IBM, "What is robotic process automation (RPA)?": https://www.ibm.com/think/topics/rpa

True exposure is likely higher, since only 37% of survey respondents have full visibility into agent actions and permissions. Four in five employees are also using unauthorised AI tools, according to an UpGuard study.[6,7] "Today, all Software-as-a-service (SaaS) tools have some form of AI in them and it is not always that obvious where it is interacting with your data," argues Harry Borovick, general counsel at Luminance, a law-tech firm. "One of the biggest risks is shadow AI usage, as there are so many free tools and systems and people using AI on personal devices."

This is reflected by the survey results — respondents are most worried about external supply chain threats like exploitation and leakage (see Figure 3). "Attackers are going to use AI and they are going to be able to attack faster because they do not have governance boards to follow," says Andrea Abell, the chief information security officer at Eli Lilly.

Figure 3. The double agent

Responses to survey question: "Which of the following sources of cybersecurity risk are most concerning for your organisation? Please select the top two concerns."

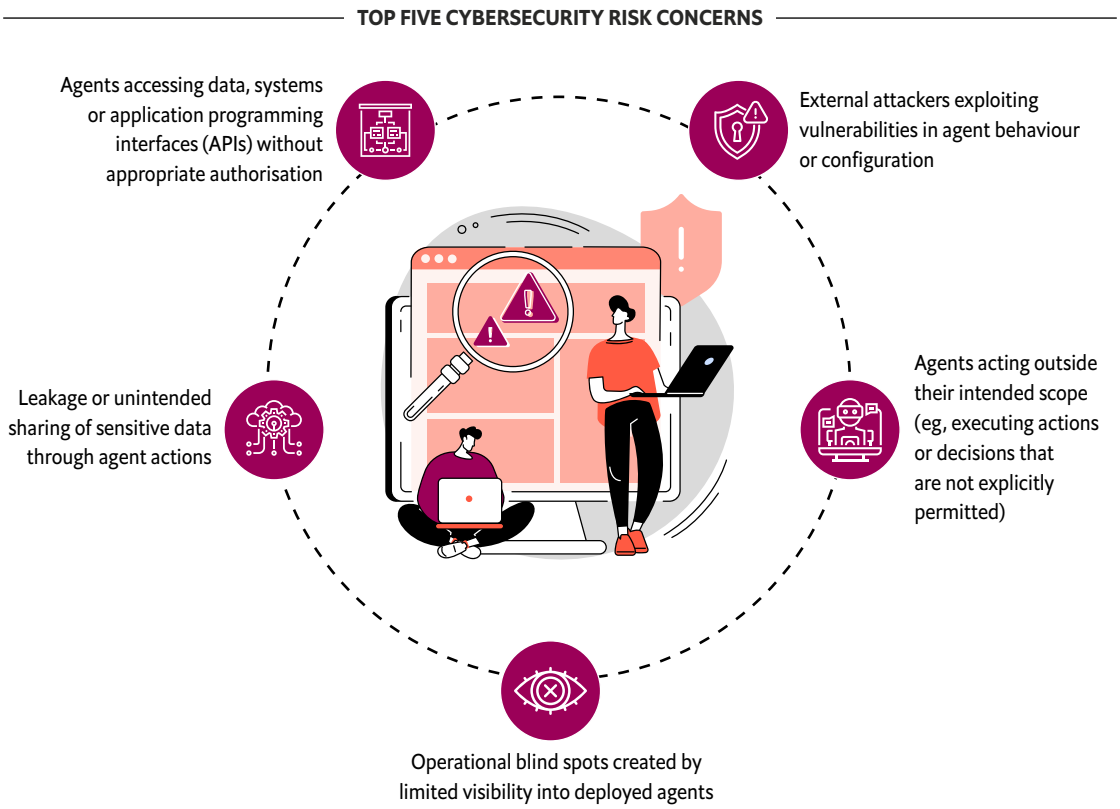

Source: Economist Enterprise *'Power without Control'* survey (2026)

[6] UpGuard, "The State of Shadow AI", November 2025: https://www.upguard.com/resources/the-state-of-shadow-ai
[7] Cato Networks, "The Shadow AI reality: Inside Cato's survey results", December 2025: https://www.catonetworks.com/blog/shadow-ai-reality-inside-catos-survey-results/

> **"How do you measure and manage to ensure they are operating well? We are not at a stage of having fully autonomous agents today and I don't think many companies are."**
>
> **James Swanson**, chief information officer, Johnson & Johnson

One of the biggest emerging security unknowns is agent-to-agent interaction. Facilitated by innovations such as Anthropic's Model Context Protocol[8] and Google's Agent2Agent Protocol,[9] which provide open standards for agents to access tools and communicate with one another, this offers obvious appeal in unlocking a productivity step change. But current security systems are not equipped to manage the security risks of agent-to-agent interaction. "When you link agents together, they can lose context," notes James Swanson, chief information officer at Johnson & Johnson. "How do you measure and manage to ensure they are operating well? We are not at a stage of having fully autonomous agents today, and I don't think many companies are."

[8] Model Context Protocol, "What is the Model Context Protocol (MCP)?": https://modelcontextprotocol.io/docs/getting-started/intro
[9] Google for Developers, "Announcing the Agent2Agent Protocol (A2A)", April 2025: https://developers.googleblog.com/en/a2a-a-new-era-of-agent-interoperability/

The illusion of preparedness

You can't recover what you can't see

While organisations are already conscious of their risk exposure, and the vast majority have either formal or informal recovery targets in place, they lack the systems to deliver on them. Observability is low, which implies limited visibility into whether breaches are even happening in the first place and most organisations are not regularly reporting them to the C-suite, risking a disconnect between organisational assumptions and reality.

Observability is a prerequisite for timely response, but only 37% maintain a complete and up-to-date inventory of agents and their authorised actions. The same share lacks full visibility into all their permissions, decision rights or actions. These capabilities are interconnected, experts note. "[Only] with observability can you understand how an agent is operating, what data it is ingesting and what decisions [it makes]. All are necessary to facilitate reversibility, especially in high-tempo environments," explains Perk's Mr Cooke.

"[Only] with observability can you understand how an agent is operating, what data it is ingesting and what decisions [it makes]. All are necessary to facilitate reversibility, especially in high-tempo environments."

Andrew Cooke, chief legal officer, Perk

Visibility falls further in the upper ranks. Most organisations lack systematic reporting to senior management. Only a quarter says they regularly track and report mean time to recover (MTTR) to the board or C-suite, while another quarter track it consistently within IT or security functions but do not report to senior leadership. This leaves senior personnel out of the loop on what is increasingly a business-wide performance issue.

More broadly, these patterns may reflect a disconnect between executive perception and operational reality. Across multiple survey dimensions — from MTTR escalation to accountability ownership to recovery confidence — C-suite respondents tend to express greater optimism than the vice presidents responsible for execution, suggesting that leaderships' confidence may be grounded in plans that operators view as not yet fully tested.

Tools for testing

The visibility gap is only part of the problem; organisations also lack the tools and systems needed to respond to the breaches they do spot. Many have only partial or emerging rollback capabilities — the ability to patch and recover from agent-related incidents. Nearly half say their rollback processes work only in part. Just three in ten report robust, well-tested reversibility capabilities, with public sector organisations particularly exposed at just 18%. The public sector also shows the weakest testing discipline, with 11% completing regular MTTR reporting versus 25% on average. However, this sector also demonstrates the highest threat awareness of unintended data leakage (46% versus 34% on average).

Organisations report reasonable scenario preparedness, with 73% defining their minimum viable business (MVB) configurations — the critical systems and processes needed to sustain essential operations during an incident — but only 40% test them regularly. Financial services and healthcare lead in MVB testing, preparedness and visibility, reflecting high-stakes environments and stringent regulation. The EU's Digital Operational Resilience Act, effective January 2025, mandates robust testing in finance, while healthcare's patient safety imperatives and HIPAA requirements drive similar discipline.

Figure 4. What you see is what you can control

Responses to survey question: "Which of the following best describes your organisation's visibility into AI agents operating within your systems? Please select one."

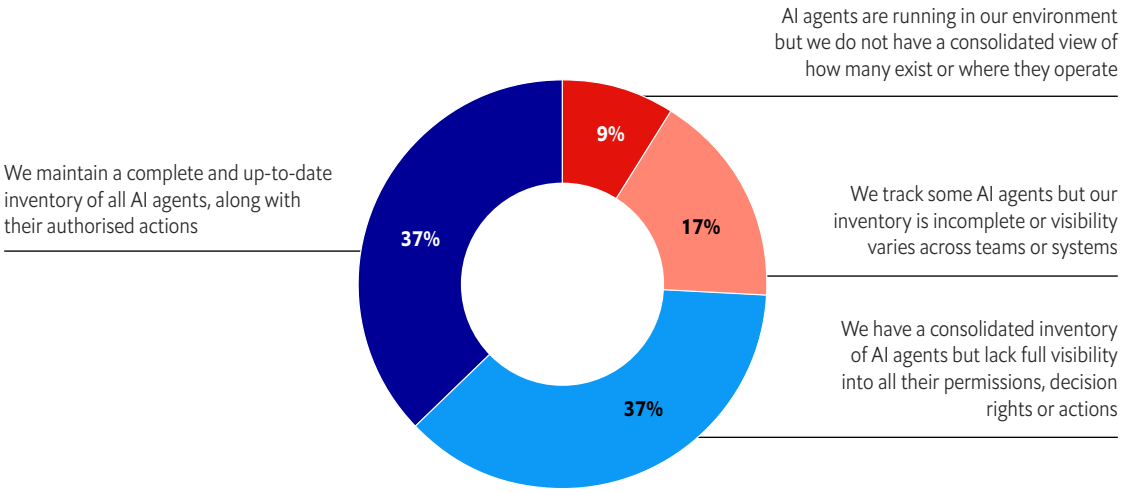

Source: Economist Enterprise *'Power without Control'* survey (2026)

"[In finance] there is a much more methodical exploration of [defining] what robust testing is," says Daniel Kendzior, global digital and AI transformation leader at Accenture. With 70% of banking executives already deploying or piloting AI agents,[10] this testing discipline is fundamental to enabling innovation with confidence for mission-critical applications such as fraud detection and anti-money laundering while protecting against regulatory penalties, financial losses and reputational damage.

In healthcare too, the stakes inform highly risk-aware deployment. "There is always a tension between moving fast and doing that securely," says Mr Fuchs at Eli Lilly. "What makes our situation difficult is that we are in a regulated space. Manufacturing is regulated. Medicines are regulated. Everything that touches patients is regulated. That is very different from a standard technology company where you can move fast and break things. If it is a streaming service, it does not matter if it goes down."

"What makes our situation difficult is that we are in a regulated space. Manufacturing is regulated. Medicines are regulated. Everything that touches patients is regulated. That is very different from a standard technology company where you can move fast and break things. If it is a streaming service, it does not matter if it goes down."

Thomas Fuchs, senior vice president, chief AI officer, Eli Lilly

[10] MIT Technology Review, "Reimagining the future of banking with agentic AI", September 2025: https://www.technologyreview.com/2025/09/04/1123023/imagining-the-future-of-banking-with-agentic-ai/

The new fundamentals of cyber risk

Information security (infosec) approaches that weathered the drizzle will fail in the coming storm. Nearly nine in ten (87%) respondents believe agents introduce fundamentally new types of risk that existing controls were not designed to manage. "I don't think the traditional infosec toolbox will cut it for the agentic world," says Sastry Durvasula, who is the chief operating, information and digital officer at TIAA, a lifetime income-focused financial services company. As agents operate at higher levels of autonomy, failures become more opaque, novel and harder to trace — with liability to escalate.

The threat is no longer at the perimeter

Top concerns cover both internal and external threats with external attacks, unauthorised access and agents acting beyond their intended scope, all registering similar levels of concern in the survey. This challenges the 'perimeter-defence' paradigm, in which cyber teams focused on defending the organisation from malicious actors. Now, internal accidents like agents running amok could become a greater threat.

Accenture's Mr Kendzior says organisations could manage with weaker systems in times past. Data governance, for instance, must be far tighter today. "AI is able to ingest a lot more data than historically required, so companies could get away with looser data security protections in the past." Controlling access based on identity is a second domain that, while not new, is far more challenging. Traditional identity frameworks like NIST's Digital Identity Guidelines[11] were designed for human users, not autonomous agents. "There is a proliferation of far too many agents needing identities. How they request and receive access is different in terms of velocity," Mr Kendzior observes. Agentic identities are the only emerging enterprise technology Gartner rates as a "very high"[12] mass attack risk, noting that their speed and privileged access could outmatch traditional security and expose organisations to critical risks in the next three years.

Mr Kendzior also notes that tracing and understanding incident causes is also harder today. "Doing root cause analysis is more complicated in the world we are in now; some of that is due to the novelty of vendor relationships and new systems. Your baseline of what you view as a normal system might be very limited."

[11] NIST, "Digital Identity Guidelines:"https://pages.nist.gov/800-63-3/
[12] Gartner, "Emerging Tech Impact Radar: AI Cybersecurity Ecosystem", October 2025: https://www.gartner.com/en/documents/7087098

> **"Companies are being more intentional about implementing tighter 'kill switch' controls. There is an acknowledgement that we do not know everything that is going to happen, as these systems are non-deterministic, so in the event of an incident, companies are asking how they can cut the cord in a way that is intentional, methodical and repeatable."**
>
> **Daniel Kendzior**, global digital and AI transformation leader, Accenture

Retail illustrates these integration challenges. Sandra Stanley, the chief data science officer at dunnhumby, which is part of Tesco Group, says that brands risk losing control over how they appear to agents — and over the customer relationship itself. External platforms like Google's Merchant Center mean "Google is deciding how your brand gets shown," explains Ms Stanley. But building internal agents to retain control exposes the infrastructure gap. "To make it a true agent [requires] connecting availability systems, pricing engines, product databases, customer transactions — a lot of plumbing needs to be done."

The rising importance of the rollback

When agents operate across fragmented infrastructure with limited external control, security teams cannot rely on preventive-only approaches. Despite consensus that incidents are inevitable (90%) and the fact that virtually all organisations with operational AI agents have experienced them, spending still remains weighted towards prevention rather than response and recovery, with no change planned through 2030 (see Figure 5).

Rollback is becoming mission-critical as security teams accept these limitations. Metrics like MTTR are still important, but "what we are trying to do now is react and contain impact so it has the least amount of ability to impact a larger part of the organisation. Resilience is the real measure of cyber success," says Eli Lilly's Ms Abell.

Figure 5. Investment intent

Responses to survey question: "Thinking of your organisation's cybersecurity investment priorities, approximately what share of spending goes to prevention versus recovery or reversibility, both today and as you expect in 2030?"

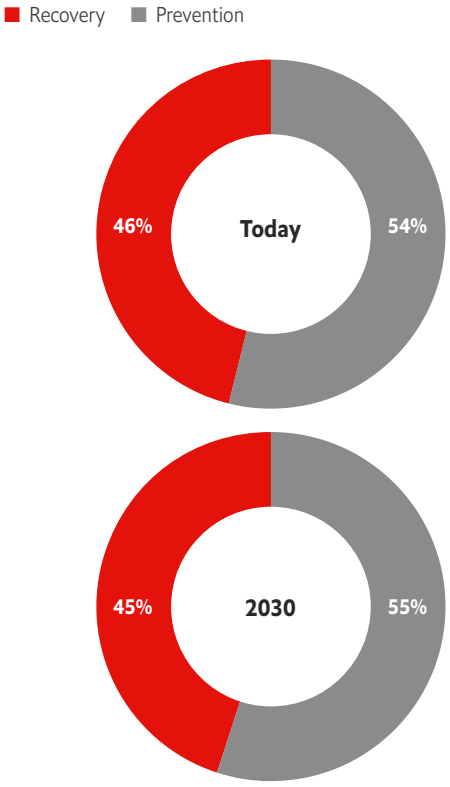

Source: Economist Enterprise *'Power without Control'* survey (2026)

"Companies are being more intentional about implementing tighter 'kill switch' controls. There is an acknowledgement that we do not know everything that is going to happen as these systems are non-deterministic, so in the event

Figure 6. The rollback imperative

Responses to survey question: "Which statement best reflects your organisation's current level of 'reversibility' for cybersecurity incidents involving AI agents? Please select one."

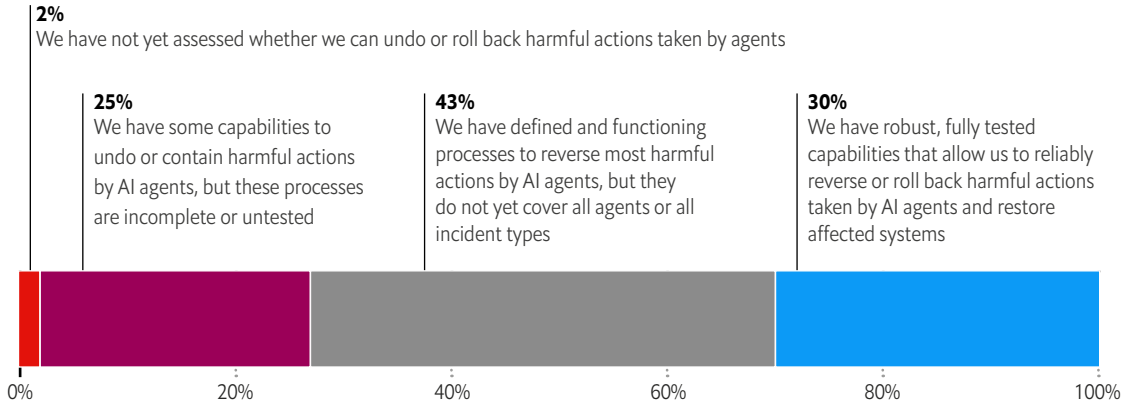

Source: Economist Enterprise *'Power without Control'* survey (2026)

of an incident, companies are asking how they can cut the cord in a way that is intentional, methodical and repeatable," says Mr Kendzior at Accenture. But while most surveyed organisations have implemented some form of process to reverse harmful actions by agents, these are mostly untested for a quarter or do not cover all types of agents for nearly half (43%).

Infosec teams can still rely on time-worn principles like 'least privilege' (granting only the minimum necessary access) and 'zero trust' (verifying every request, regardless of source). Frameworks like the National Institute of Standards and Technology's AI Risk Management Framework[13] and the MITRE ATLAS (Adversarial Threat Landscape for AI Systems)[14] knowledge base provide good starting points for specialised AI security. But the speed and scale of the threats cannot be matched by manual patching and response. "With AI, you could generate a DDoS or malware attack a lot faster. When you thought you had X number of hours to resolve, you now have minutes, so you have to think about what you need to automate to get to a faster response rate. It is not new, but the timeline has shrunk," says Mr Swanson of Johnson & Johnson.

Palo Alto Networks' Unit 42 threat intelligence team demonstrated how dramatic that compression scale is when it simulated a full ransomware attack with agentic AI. From initial compromise to data exfiltration, it took just 25 minutes using AI at every stage of the attack chain,[15] a roughly 100-fold increase in speed compared with traditional attacks. Automation itself will have to become part of the response toolbox, as organisations fight fire with fire. TIAA is deploying agents across technology, development and operations, explains Mr Durvasula, including to defend against agent-based threats. Eli Lilly also views cyber agents as an indispensable part of its overall security toolbox, according to Ms Abell.

[13] NIST, "AI Risk Management Framework", January 2023: https://www.nist.gov/itl/ai-risk-management-framework
[14] Mitre Atlas, "Navigate threats to AI systems through real-world insights": https://atlas.mitre.org/
[15] Palo Alto Networks, "Unit 42 Develops Agentic AI Attack Framework", May 2025: https://www.paloaltonetworks.com/blog/2025/05/unit-42-develops-agentic-ai-attack-framework/

Rating cyber response

A cyber-resilient organisation can rapidly detect, contain and roll back harmful agent actions — with clear target recovery timeframes and the ability to maintain minimally viable operations even during remediation. This requires recovery plans that stand up to reality, spanning procedural and organisational factors such as clear accountability, robust reporting and change management. But preparedness and confidence vary across sectors.

(Mis)leading with confidence

Looking at organisations' ability to meet recovery times, retail and healthcare respondents report the most confidence, with just under a third reporting 'certainty' they can meet them. These sectors face constant operational pressure where downtime directly impacts customers and patient safety, driving disciplined incident response capabilities.

Retail's confidence may reflect a deployment strategy that favours internal use over customer-facing metrics, says Ms Stanley at dunnhumby. "We polled almost 2,000 UK shoppers and discovered that only 3% use agentic tools for grocery shopping, rising to just 6% for general retail." While some large retailers like Walmart are partnering with leading AI companies on customer-facing agentic initiatives like universal commerce protocols, the limited customer adoption of agentic commerce means most retailers stay focused elsewhere. "Retailers are very much focused on workflow, automation, productivity gains — how we use agents inside," Ms Stanley observes.

This focus on internal agent workflows over external applications means retailers are building recovery capabilities in more controlled environments, contributing to operational confidence. By contrast, the technology industry had the largest share of respondents expressing the least confidence, perhaps reflecting the experimental nature and rapid evolution of their agent deployments, which outpace their recovery planning.

Yet broad confidence in recovery capabilities across sectors appears misplaced. Given limited actual observability into agents, and the lack of systematic reporting to the C-suite, decision-makers could be more optimistic than warranted. Respondents identify a number of response bottlenecks, spanning procedural and technical limitations, led by rarely tested recovery procedures and insufficient real time monitoring, the most frequently cited obstacles to faster recovery.

Figure 7. How sectors stack up

Sectoral trends in agentic risk and readiness

Healthcare: Most tested, least transparent	More healthcare organisations have a fully defined and regularly tested recovery configuration than any other sector (**46%**), reflecting a culture of procedural discipline.	That rigour extends to reversibility: **36%** have robust, fully tested capabilities to roll back AI agent actions, the highest share of any sector.	But strong processes don't guarantee transparency: **41%** have experienced an incident where agent behaviour simply could not be explained or audited, a higher rate than any other sector.	
Financial services: Most visible, least automated	Financial services organisations are in the know: **44%** have a complete, up-to-date inventory of all AI agents, joint highest with healthcare — audit discipline is carrying over into the AI era.	Accountability is better defined than anywhere else: **46%** place primary responsibility for cyber incident management with the chief information security officer (CISO) — 14 percentage points above average.	Compliance pressure shapes deployment: **32%** cite insufficient guardrails for compliance, data governance or regulatory requirements as a top deployment challenge, the highest of any sector.	The gap is in speed of response: **42%** say over-reliance on manual recovery processes is a key limitation — the sector can identify problems, but fixing them takes time.
Infrastructure: Most exposed, least monitored	Infrastructure organisations are the least able to see what's happening in real time: **41%** cite insufficient monitoring of systems and agent activity as a key recovery limitation.	When things go wrong, the consequences hit hardest in the markets: **36%** identify investor and analyst confidence as a top business impact of an AI incident.	Engineering owns the risk: **31%** place primary accountability with the chief technology officer, the highest proportion of any sector, reflecting where operational responsibility tends to sit in these organisations.	The incident record underlines the exposure: **40%** have experienced data leakage or unintended exposure via an AI agent in the past 12 months.
Technology: Most ambitious, least hardened	Technology firms are building and deploying AI agents at pace, yet only **28%** have full visibility into the agents operating in their own environments.	The resilience infrastructure hasn't caught up: just **26%** have robust, fully tested capabilities to reverse harmful AI agent actions, placing it second last and only ahead of the public sector.	Over half (**53%**) have defined reversal processes that have never been properly tested, meaning the capability exists on paper but is untried under pressure.	Recovery confidence reflects this gap, with only **14%** very confident of meeting their recovery target following an incident.
Retail: Most confident, least equipped	Retail is the most recovery-ready on paper: **57%** have a clearly defined target timeframe for returning to full operations after an incident, joint highest of any sector.	And unlike in other sectors, that confidence appears genuine: **30%** say they are very confident of actually meeting that target, the highest rate recorded.	The tension lies in skills: **30%** cite lack of internal expertise to build, monitor or secure AI agents as a top deployment challenge, more than any other sector, suggesting confidence may be running ahead of capability.	External threat exposure is also elevated, with **36%** having experienced a threat actor manipulating or exploiting an AI agent in the past year, a sector high.
Public sector: Most threatened, least prepared	The public sector has the weakest reversibility posture: only **18%** have robust, fully tested rollback capabilities, a gap that structural constraints and limited modernisation budgets are making hard to close.	Visibility is also a critical weakness: **40%** cite limited insight into agent actions and permissions as a top deployment challenge, the highest of any sector.	Data risk weighs heavily: **46%** rank data leaks or unintended sharing as primary concerns, reflecting the sensitivity of these organisations' information.	Leadership awareness is sector-leading: with **66%** of CEOs and **68%** of CISOs very concerned and actively prioritising AI cybersecurity risks — the will is there, but the means to act on it are not yet in place.

Source: Economist Enterprise *'Power without Control'* survey (2026)

Structural problems also stymie response, such as unclear ownership or accountability for agent-related failures — a top three challenge when deploying agentic AI. When agent failures can cascade across finance, legal and supply chain functions, clear accountability becomes essential to every function.

Accountability matters

Accountability gaps are less likely to be a top challenge for infrastructure and public sector respondents, indicating leadership and prioritisation in these sectors. This reflects the mission-critical services both sectors provide and their relatively high exposure to cyber incidents, especially malevolent actors. Half of all ransomware attacks in 2025 targeted critical sectors, including manufacturing, healthcare, energy and transport[16] — a 34% year-over-year increase in attacks on essential infrastructure.

Technology companies (33%) are more concerned by accountability gaps, a top priority among their listed challenges. This may reflect higher levels of experimentation and a generally higher threshold for risk appetite. It also points to the intense competitive pressure facing the technology industry, dubbed the 'SaaS-pocalypse', as AI itself threatens to disrupt traditional SaaS business models (see Figure 8).

While rollbacks should be a growing priority, experts do warn against excessive reliance on fixing breaches after the fact. Accenture's Mr Kendzior warns against overly bullish perspectives, especially from boards and non-technical management, that assume security can just fix technical issues rapidly. Perk's Mr Cooke also counsels against 'silver bullet' thinking that assumes problems can be quickly fixed. And for the most high-stakes industries or products, later fixes may not be an option at all. "With a patient focus, accidents cannot just happen," says Mr Fuchs at Eli Lilly. "Failure is not an option."

Figure 8. The software stack grows legs

Illustrative shift in total addressable market (TAM) for software as a service (SaaS) and agents (US$ bn)

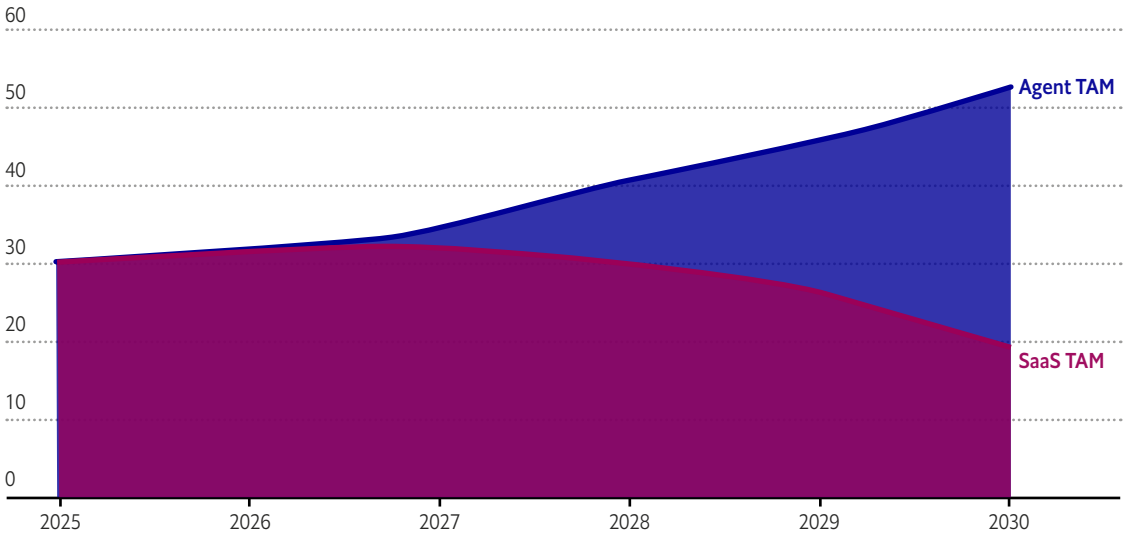

Source: Goldman Sachs[17]

[16] KELA, "Escalating Ransomware Threats to National Security", October 2025: https://www.kelacyber.com/resources/research/escalating-ransomware-threats-to-national-security/
[17] Goldman Sachs, "AI Agents to Boost Productivity and Size of Software Market", July 2025: https://www.goldmansachs.com/insights/articles/ai-agents-to-boost-productivity-and-size-of-software-market

Conclusion: Steps to security

Failures are inevitable in today's accelerated, AI-infused technology environment. The organisations that succeed will be those that build the tools, systems and behaviours that can deliver on their targets. That includes full visibility into their agentic universe, robust testing and scenario planning, and rollback capabilities that can move at machine speed. This requires a blend of procedural clarity and operational readiness, alongside wider organisational alignment.

Agentic is a marathon, not a sprint. Boards and senior leaders are understandably excited about the potential of agents — and the market imperative to adopt to avoid falling behind. But early haste will lead to regrets later. From advising multinational boards on AI transformation, Accenture's Mr Kendzior frames the agentic transition over three crucial stages. First, getting control of 'shadow AI' and modernising cybersecurity platforms. Second, adding AI functionality into existing systems and, third, reinventing infosec operations to be leaner and more effective. The progression is cumulative: discipline first, augmentation second, reinvention last. Measured progress rather than rushing also ensures the C-suite can align with technical teams and understand operational realities. "People often get excited by the possibility of a tool and treat it like a sprint. Quick wins do build momentum. But for sustained use and transformation, you need to address risk up front," emphasises Mr Borovick at Luminance.

> **"People often get excited by the possibility of a tool and treat it like a sprint. Quick wins do build momentum. But for sustained use and transformation, you need to address risk up front."**
>
> **Harry Borovick**, general counsel, Luminance

"It requires getting people's mindsets around risk aligned and getting the technology to work in a way the business can trust long term," he emphasises. Careful early steps can, ultimately, drive faster innovation later. For instance, creating templates for agent design provides a common platform

of tooling and procedures that unlocks creativity across the organisation. "If you create the guardrails, the technology teams can actually innovate faster," argues Ms Abell at Eli Lilly.

Forge an organisation-wide risk tolerance strategy, ensuring each function and role understands the others. The survey has revealed differing levels of confidence on cyber response, which partly stems from patchy incident reporting and limited visibility. Infosec risk is now everybody's problem and, as a result, boards and technical leads need tighter co-ordination.

The strategy "lies not just with the CIO. It's with our executive council, general counsel, chief risk officer, the chief financial officer — we align together on the risk profiles, the outcomes we want and then we go after it," says Johnson & Johnson's Mr Swanson. Functions each need to adapt and evolve too. "For many years, the legal function was the 'Department of No'. That has never been less true," argues Luminance's Mr Borovick. "Lawyers realise that AI is acting aggressively in the legal function, so their value is, increasingly, as advisers. In-house legal teams have been outstripping other teams, whether HR or even tech, when it comes to AI adoption."

> **"[The strategy] lies not just with the CIO. It's with our executive council, general counsel, chief risk officer, the chief financial officer — we align together on the risk profiles, the outcomes we want and then we go after it."**
>
> **James Swanson, chief information officer, Johnson & Johnson**

Identify and monitor the AI risk that lurks in third-party shadows. Even within their digital estates, organisations have limited knowledge of agentic risk. But the supply chain could be a growing vector of trouble. Traditional vendor security approvals in areas like data management and infosec are not fit for purpose in the age of agents and AI. "I don't think supply chain maturity has kept up with the times," argues Ms Abell at Eli Lilly. "Most of the time, it is a checkbox and a lot of companies use agents and bots to answer the questionnaires." Many security breaches are coming in through embedded software, argues TIAA's Mr Durvasula. Boards and senior decision-makers should support greater investment in capacity building, collaboration and security reviews in key partners. Eli Lilly has launched in-depth architecture security engineering reviews with its key suppliers in areas like drug discovery and manufacturing. It has been a learning curve and an important investment, according to Ms Abell. Their investments are helping their suppliers to raise their standards and leverage the organisation's significant cybersecurity resources and expertise.

Involve infosec earlier in the technology procurement cycle to steer risk before deployment. AI vendors are now starting to produce all of the documentation for infosec teams as early in the process as possible. "That kind of forward-thinking approach is something we are seeing much more of in the market and it benefits everyone," notes Mr Borovick at Luminance, who advocates for continuous compliance improvements.

Yet in highly regulated sectors such as financial services, AI governance remains less mature than traditional infosec. Documentation standards are well established for legacy systems, but they are fragmented and inconsistent for AI.

"The security questions are broader and the language is still evolving. There's not always alignment on how terms like predictive, agentic, or large language models are defined, which can make it harder to assess risk consistently," says Imogen Philp, product director at software company Likezero. Standardised AI security documentation — analogous to the financial disclosures that transformed corporate governance decades ago — would give procurement teams a common baseline for evaluating agentic risk.

ECONOMIST ENTERPRISE

LONDON

The Adelphi
1-11 John Adam Street
London WC2N 6HT
United Kingdom
Tel: (44) 20 7830 7000
Email: london@economist.com

GENEVA

Rue de l'Athénée 32
1206 Geneva
Switzerland
Tel: (41) 22 566 2470
Fax: (41) 22 346 93 47
Email: geneva@economist.com

WASHINGTON DC

1920 L Street
NW Suite 500
Washington DC
20002
United States
Email: americas@economist.com

NEW YORK

900 Third Avenue
16th Floor
New York, NY 10022
United States
Tel: (1.212) 554 0600
Fax: (1.212) 586 1181/2
Email: americas@economist.com

DUBAI

Office 1301a
Aurora Tower
Dubai Media City
Dubai
Tel: (971) 4 433 4202
Fax: (971) 4 438 0224
Email: dubai@economist.com

HONG KONG

1301
12 Taikoo Wan Road
Taikoo Shing
Hong Kong
Tel: (852) 2585 3888
Fax: (852) 2802 7638
Email: asia@economist.com

SINGAPORE

8 Cross Street
#23-01 Manulife Tower
Singapore
048424
Tel: (65) 6534 5177
Fax: (65) 6534 5077
Email: asia@economist.com