ECONOMIST ENTERPRISE

Ctrl+Z

Keeping control in the age of agentic AI

Survey insights deck

Supported by rubrik

Contents

These slides contain supplementary charts and data points from our 2026 survey of 750 C-suite and vice presidents.

Details about the survey can be found in our report, ***Ctrl+Z: Keeping control in the age of agentic AI.***

ECONOMIST ENTERPRISE

About this research

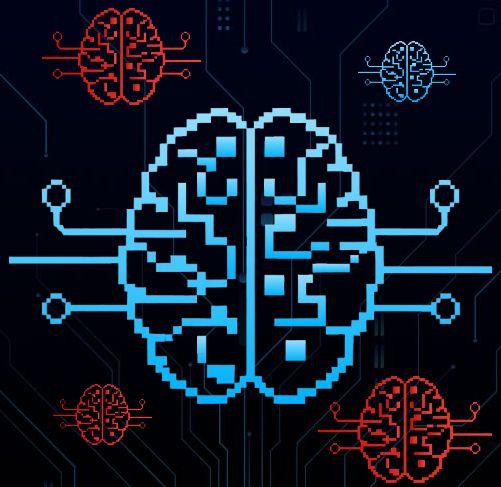

This study's global survey, developed and conducted by Economist Enterprise and supported by Rubrik, polled 800+ business leaders. It was fielded between December 2025 and February 2026.

The survey respondents are business leaders with decision-making power / high degree of influence or visibility over agentic AI, cyber risk, security and/or resilience. They work in large enterprises (annual revenue of US$500m or greater) or public sector organisations, located in nine countries across North America, Europe and Asia-Pacific.

Respondents' job titles:

- Chief Information Officer CIO) or equivalent
- Chief Data Officer or equivalent
- Chief Digital Officer or equivalent
- Chief Information Security Officer CISO) or equivalent
- Chief Technology Officer CTO) or equivalent
- Chief Risk Officer or equivalent
- Chief AI Officer or equivalent
- Chief Transformation Officer or equivalent
- Chief Privacy Officer or equivalent
- Chief Compliance Officer or equivalent
- Head of department or equivalent
- Managing Director / Executive
- Vice-President / Senior Vice-President or equivalent
- Vice-President or equivalent

ECONOMIST ENTERPRISE

About this research

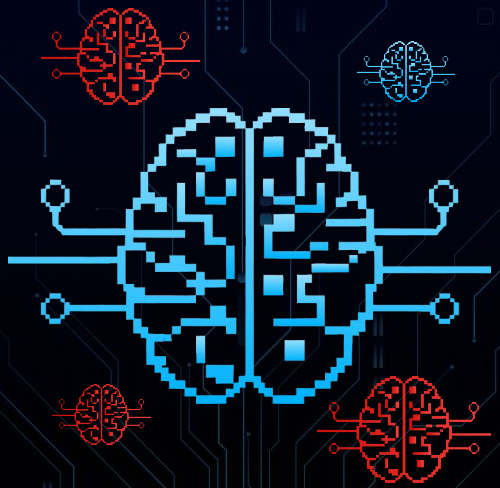

Respondents' locations:

Americas

- United States

Asia-Pacific

- Australia
- India
- Japan

Europe

- France
- Germany
- Italy
- Spain
- United Kingdom

Respondents' industries:

- Construction
- Infrastructure
- Financial services / banking / insurance
- Public sector
- Healthcare and life sciences
- Pharmaceuticals
- Retail and consumer goods
- Technology

ECONOMIST ENTERPRISE

ECONOMIST ENTERPRISE

1. The agentic boom

As agents scale, exposure multiplies

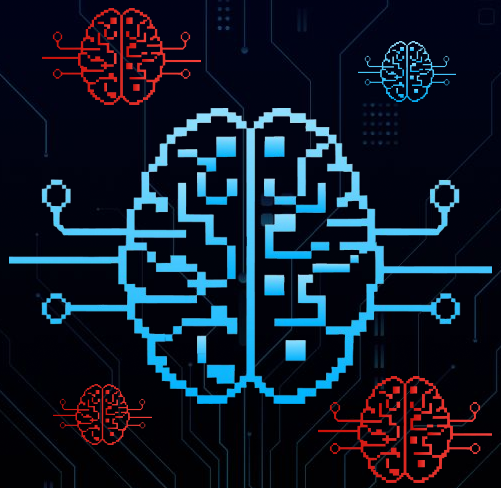

Agents, with their ability to act, decide and move across systems independently, promise a step change in productivity.

But the same autonomy that drives value is creating new forms of cyber risk, challenging organisations' ability to maintain control when things go wrong.

The market for AI agents is forecast to grow **44-fold** over the next decade. As agents spread across systems, workflows and data environments, they create more pathways for **disruption to operations, information and decision-making.**

AI-agent market revenue is set to reach

US$236bn

2034

US$5.4bn

2024

Source: Precedence Research, "AI Agents Market Size and Forecast 2025 to 2034", August 2025: https://www.precedenceresearch.com/ai-agents-market

ECONOMIST ENTERPRISE

ECONOMIST ENTERPRISE

2. Defining agentic risk

As AI agents expose critical business priorities, nearly 90% of organisations recognise the risks they bring

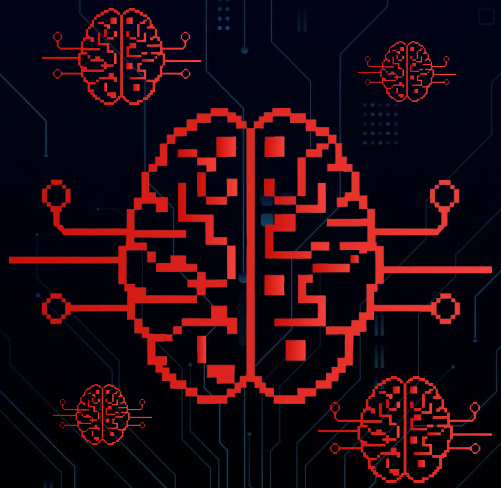

Disagree Neutral Agree Don't know

"Cyber incidents involving AI agents are **inevitable, even with strong preventative security measures**"

"AI agents have **increased our exposure** to cybersecurity and operational risks"

"Agents introduce **fundamentally new types of risks** that our existing controls were not designed to manage"

"We are deploying AI agents **faster** than our cybersecurity teams can evaluate, govern or secure them"

"Unintended interactions between AI agents create risks that are **difficult to predict or contain**"

Question: To what extent do you agree with the following statements? Please select one option per row.

ECONOMIST ENTERPRISE

Failure is inevitable

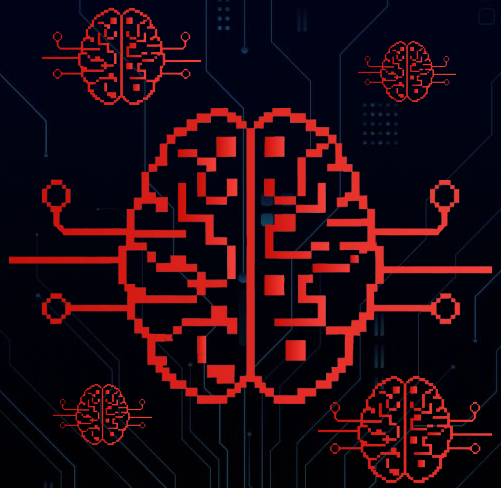

Nearly **all** organisations globally have deployed agents that have experienced at least one related incident causing organisation-wide disruption...

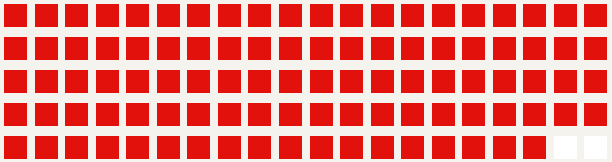

98% **have experienced at least one related incident causing organisation-wide disruption.**

...and business leaders are aware of the risks.

88%

of respondents said...

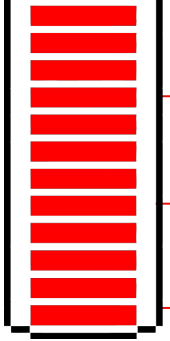

Unintended interactions between AI agents create risks that are **difficult to predict or contain**

AI agents have increased our **exposure to cybersecurity** and **operational risks**

Agents introduce **new types of risks** that existing controls were not designed to manage

They are deploying AI agents faster than cybersecurity teams can **evaluate, govern or secure**

Questions: In the past 12 months, has your organisation experienced any of the following incidents involving AI agents which resulted in organisation-wide disruption? Please select all that apply; and To what extent do you agree with the following statement - "Cyber incidents involving AI agents are inevitable, even with strong preventative security measures."

ECONOMIST ENTERPRISE

Risks are wide-ranging, with no clear frontrunner

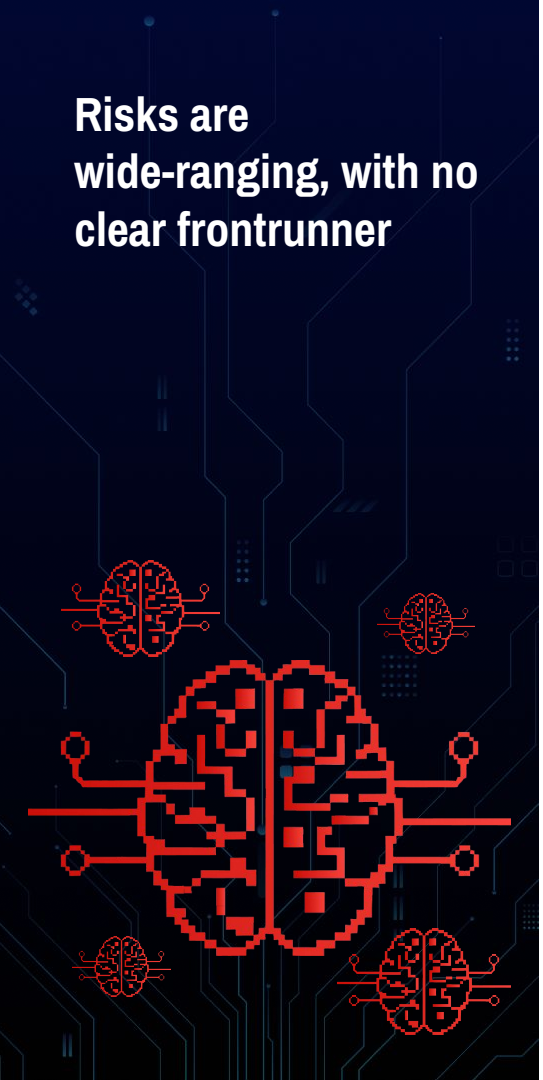

In the age of agentic, risks are widespread and diverse. When asked about their leading cybersecurity concerns, business leaders across all industries are split. This suggests concern is spread across the full range of agentic AI risks.

The proliferation of agentic risks challenges the **traditional 'perimeter-defence' paradigm**, in which cyber teams focus on defending the organisation from malicious external actors.

Top sources of cybersecurity risk most concerning for organisations

(All respondents)

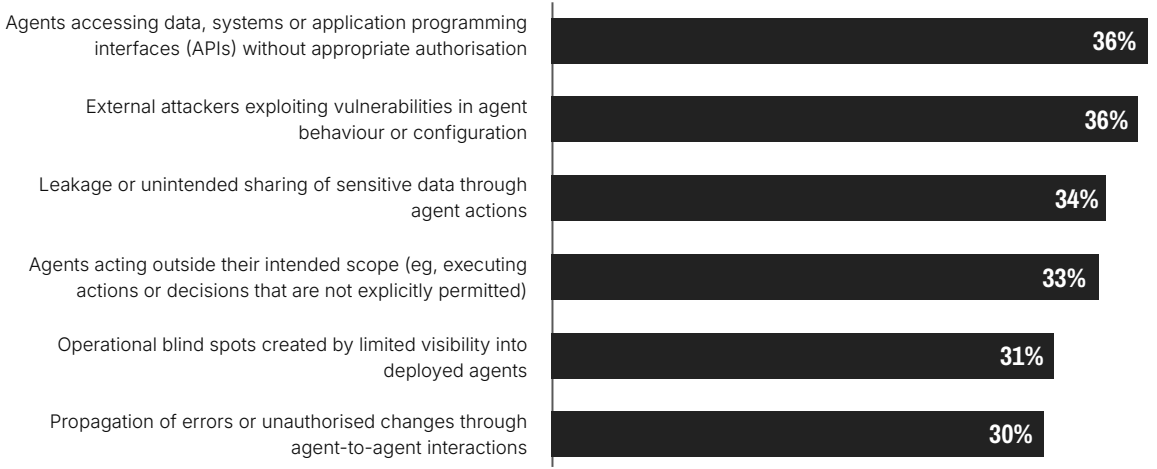

Question: Which of the following sources of cybersecurity risk are most concerning for your organisation? Please select the top two concerns.

ECONOMIST ENTERPRISE

Risks are wide-ranging, with no clear frontrunner

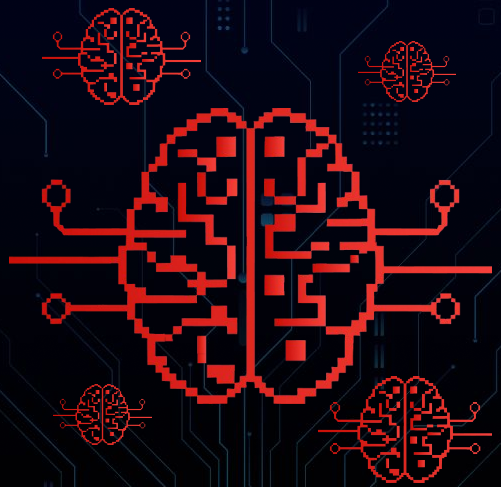

When asked about their leading cybersecurity concerns, business leaders in the **healthcare / pharma / life sciences** industry are relatively split.

Leading concerns are agents accessing data, systems or application programming interfaces (APIs) without appropriate authorisation and leakage or unintended sharing of sensitive data through agent actions

Top sources of cybersecurity risk most concerning for organisations

(Healthcare / pharma / life sciences respondents)

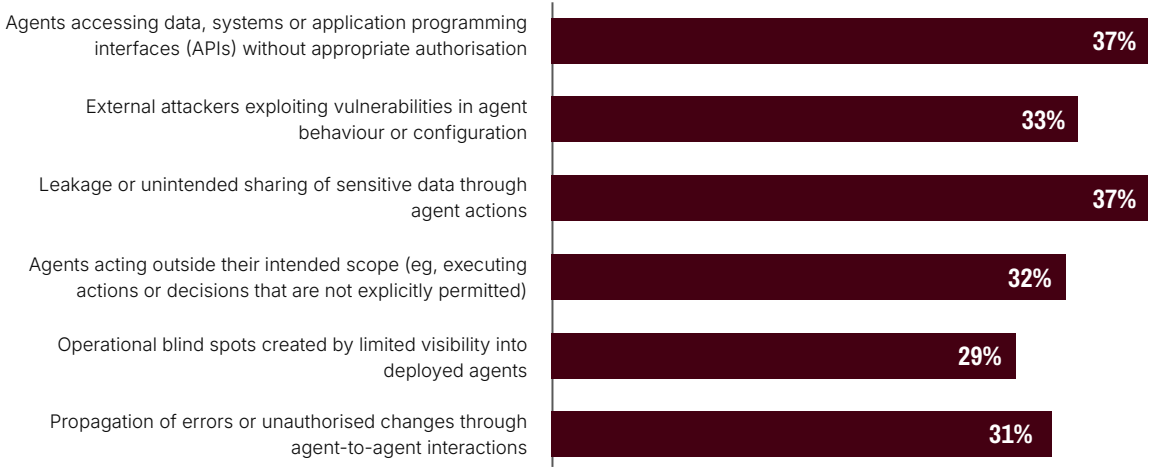

Question: Which of the following sources of cybersecurity risk are most concerning for your organisation? Please select the top two concerns.

ECONOMIST ENTERPRISE

Risks are wide-ranging, with no clear frontrunner

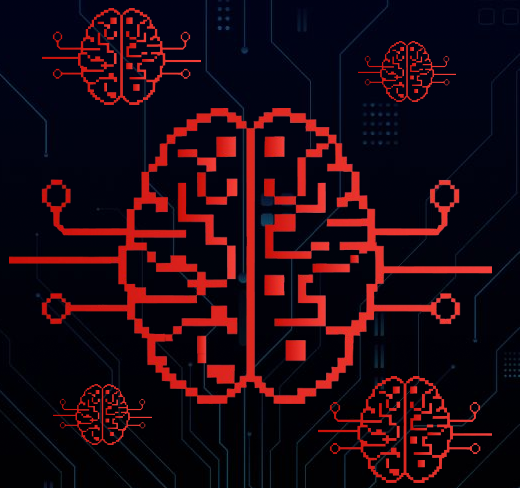

When asked about their leading cybersecurity concerns, business leaders in the **infrastructure / construction** industry are relatively split.

The leading concern is external attackers exploiting vulnerabilities in agent behaviour or configuration.

Top sources of cybersecurity risk most concerning for organisations

(Infrastructure / construction respondents)

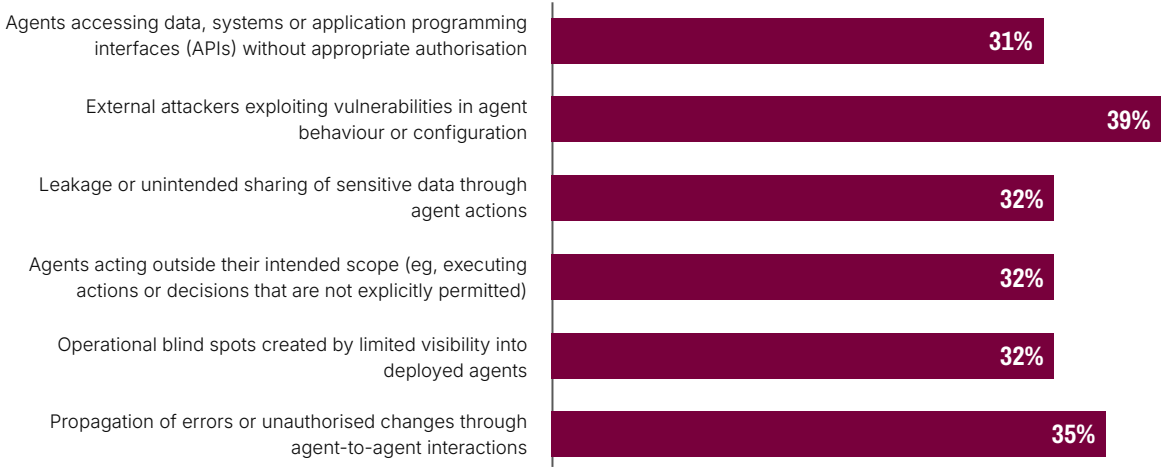

Question: Which of the following sources of cybersecurity risk are most concerning for your organisation? Please select the top two concerns.

ECONOMIST ENTERPRISE

Risks are wide-ranging, with no clear frontrunner

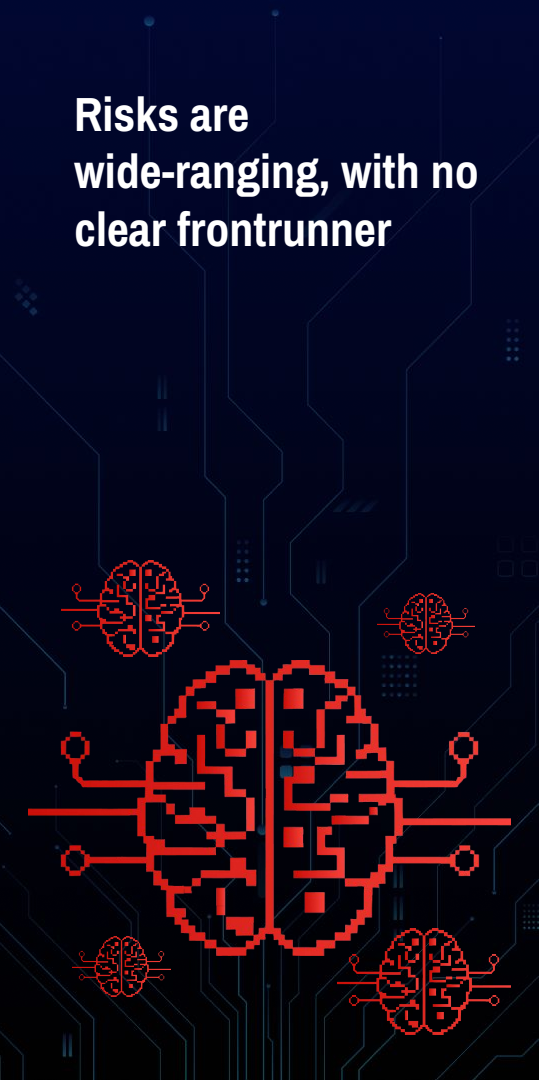

When asked about their leading cybersecurity concerns, business leaders in the **technology** industry are relatively split.

The leading concern is agents accessing data, systems or application programming interfaces (APIs) without appropriate authorisation.

Top sources of cybersecurity risk most concerning for organisations

(Technology respondents)

Question: Which of the following sources of cybersecurity risk are most concerning for your organisation? Please select the top two concerns.

Risks are wide-ranging, with no clear frontrunner

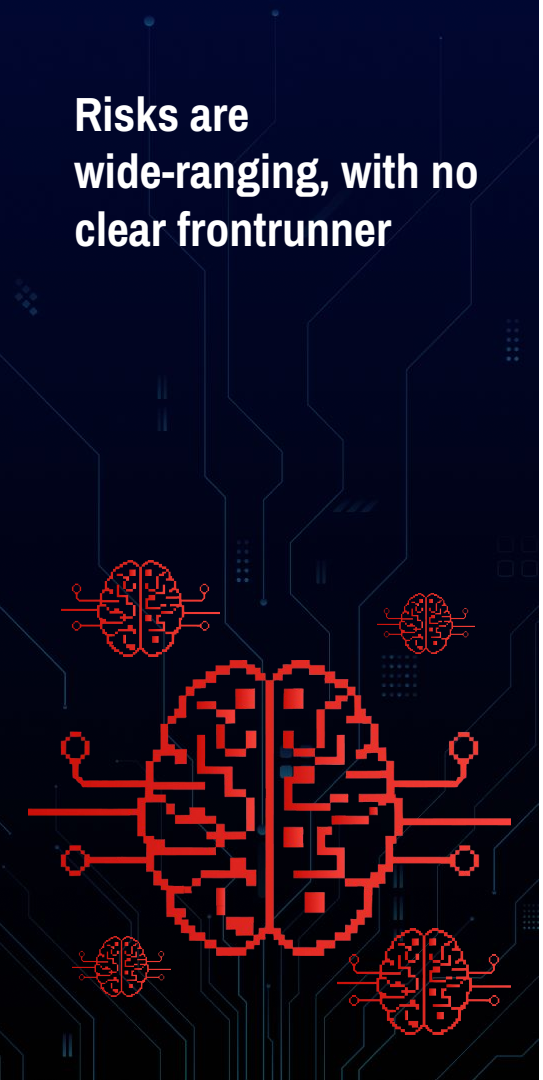

When asked about their leading cybersecurity concerns, business leaders in the **financial services / banking / insurance** industry are relatively split.

Leading concerns are agents accessing data, systems or application programming interfaces (APIs) without appropriate authorisation; external attackers exploiting vulnerabilities in agent behaviour or configuration; and agents acting outside their intended scope (eg, executing actions or decisions that are not explicitly permitted)

Top sources of cybersecurity risk most concerning for organisations

(Financial services / banking / insurance respondents)

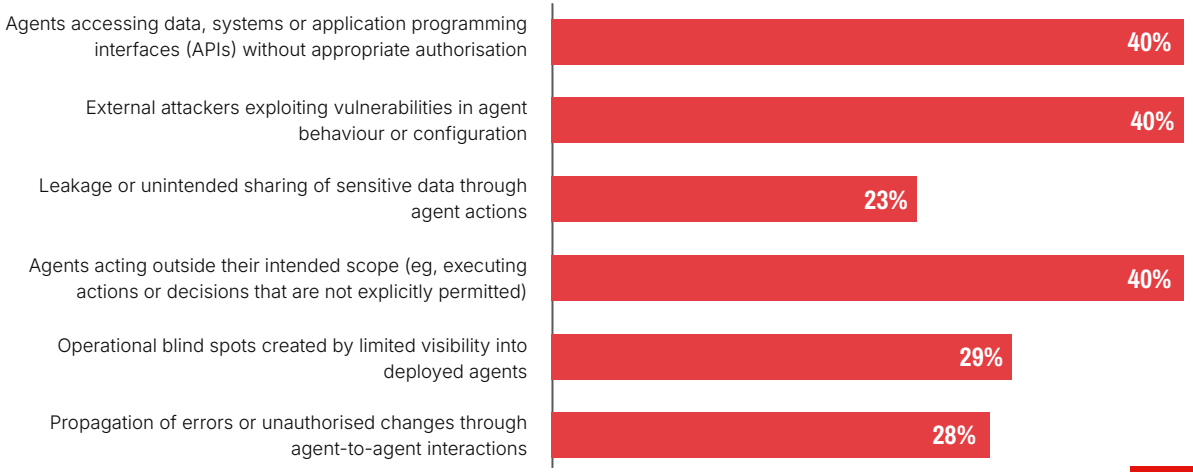

Question: Which of the following sources of cybersecurity risk are most concerning for your organisation? Please select the top two concerns.

ECONOMIST ENTERPRISE

Risks are wide-ranging, with no clear frontrunner

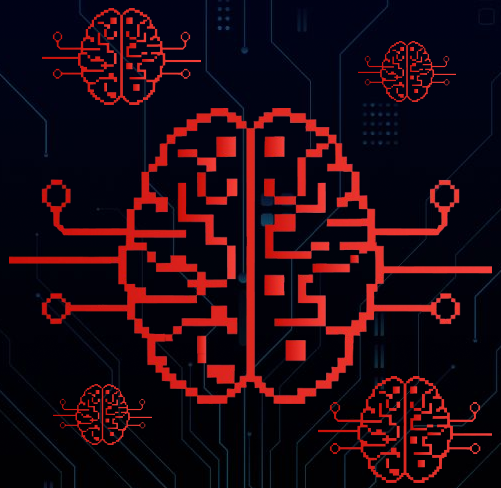

When asked about their leading cybersecurity concerns, leaders in the **public sector** are relatively split.

The leading concern is leakage or unintended sharing of sensitive data through agent actions.

Top sources of cybersecurity risk most concerning for organisations

(Public sector respondents)

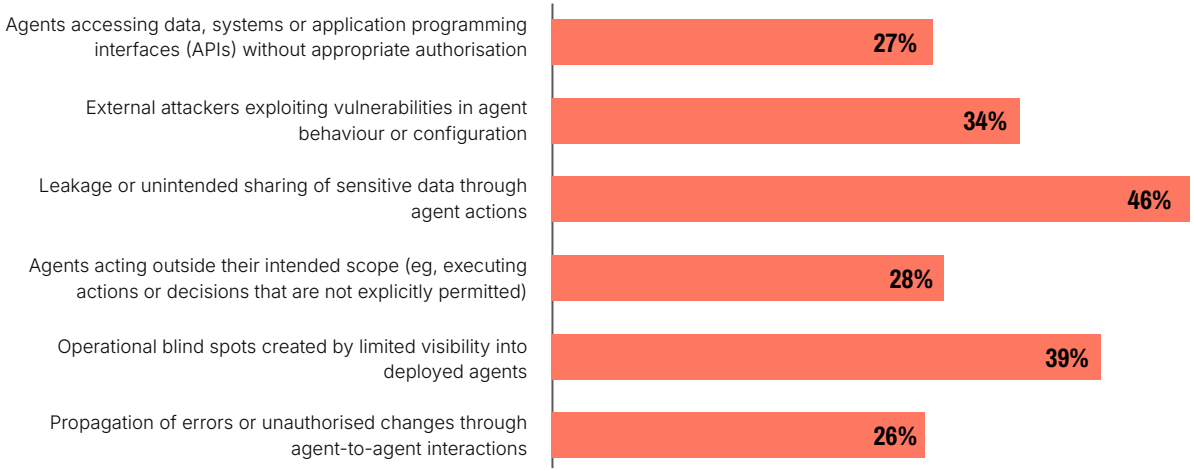

Question: Which of the following sources of cybersecurity risk are most concerning for your organisation? Please select the top two concerns.

ECONOMIST ENTERPRISE

Risks are wide-ranging, with no clear frontrunner

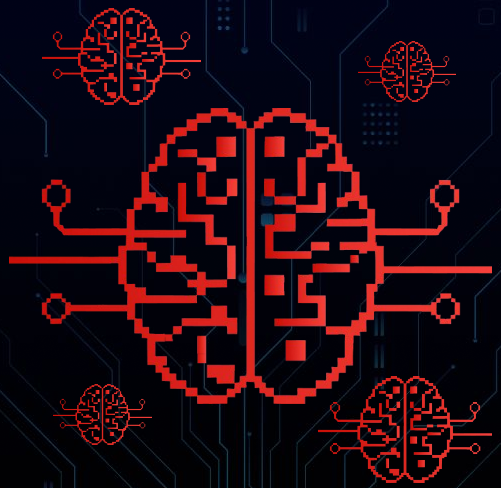

When asked about their leading cybersecurity concerns, leaders in the **retail / consumer goods** industry are relatively split.

This suggests concern is spread across the full range of agentic AI risks.

Top sources of cybersecurity risk most concerning for organisations

(Retail / consumer goods respondents)

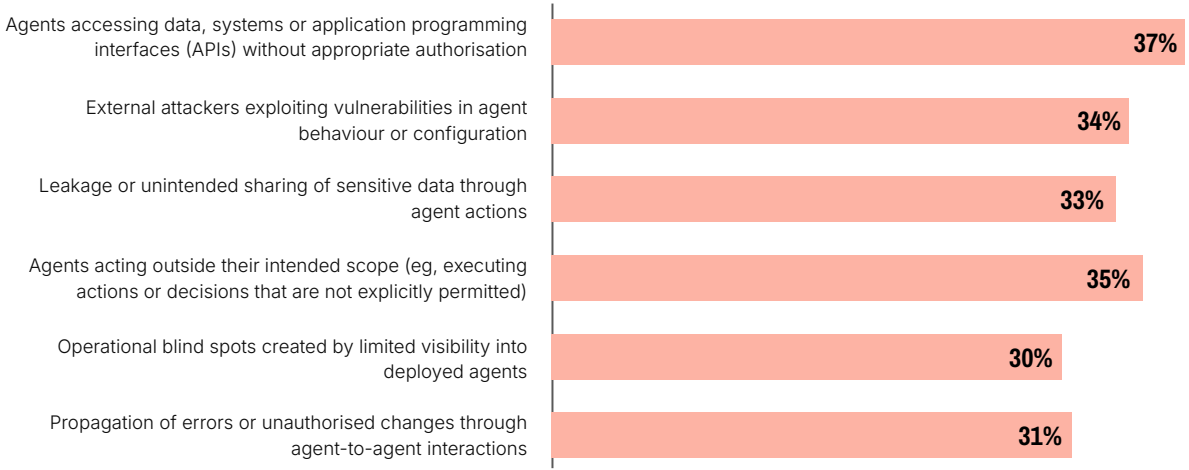

Question: Which of the following sources of cybersecurity risk are most concerning for your organisation? Please select the top two concerns.

ECONOMIST ENTERPRISE

Business leaders are aware of the inevitable risks associated with agentic

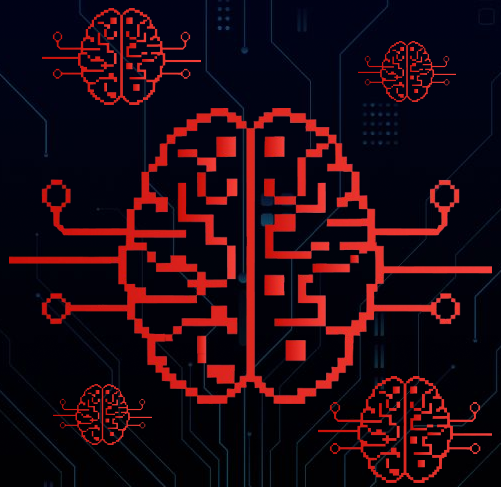

The consequences of agent-related incidents extend far beyond IT. They can disrupt operations, erode trust and create financial, legal and reputational costs.

The four business areas leaders view as most vulnerable to a major agent-related incident

- Regulatory fines
- Supply chain disruption
- Revenue impact
- Brand or reputation damage

Question: Which of the following business areas would be meaningfully affected by a major cybersecurity incident involving AI agents? Please select up to three.

ECONOMIST ENTERPRISE

3. How do we scale agentic with confidence?

If incidents are bound to happen, response matters as much as prevention

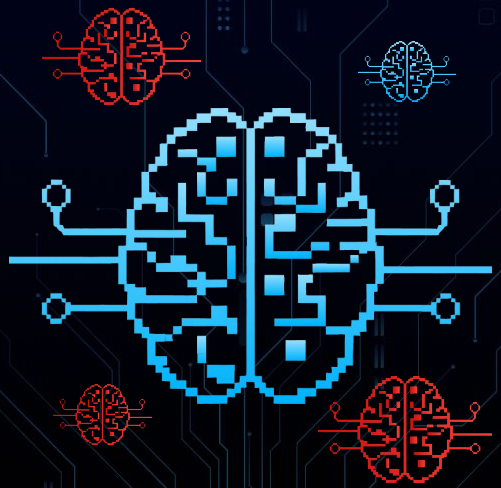

95% of organisations have set mean time to recover targets. But for many, those targets remain **informal** or **loosely defined**, creating a false sense of readiness when incidents occur.

Organisations that have defined target timeframe for recovering to full operational capacity after a cybersecurity incident involving AI agents

(All respondents)

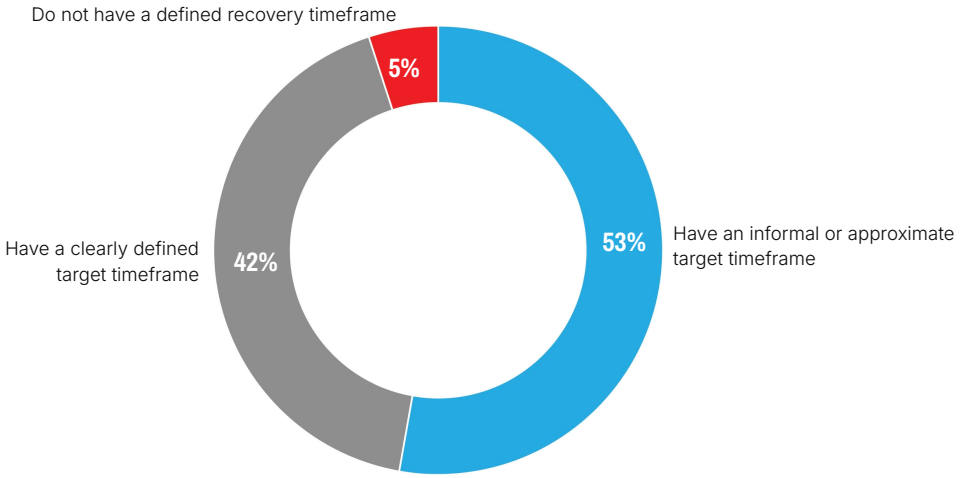

Question: Does your organisation have a defined target timeframe for recovering to full operational capacity after a cybersecurity incident involving AI agents? Please select one.

ECONOMIST ENTERPRISE

Cybersecurity investment still favours prevention over response and recovery, and leaders expect this imbalance to persist

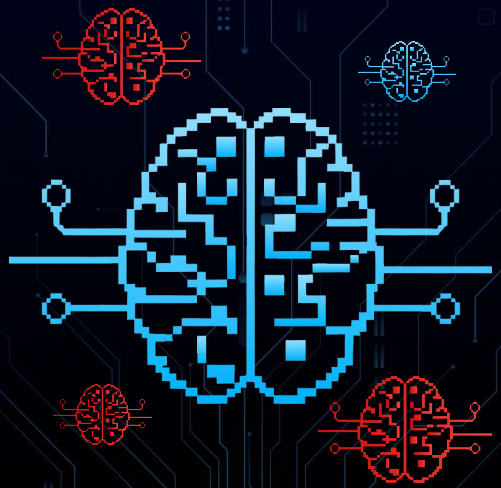

Yet as risk moves inside the organisation, fortifying the walls is no substitute for fixing the foundations.

Cybersecurity investment priorities broken down by prevention versus recovery or reversibility, today and as expected in 2030

(All respondents)

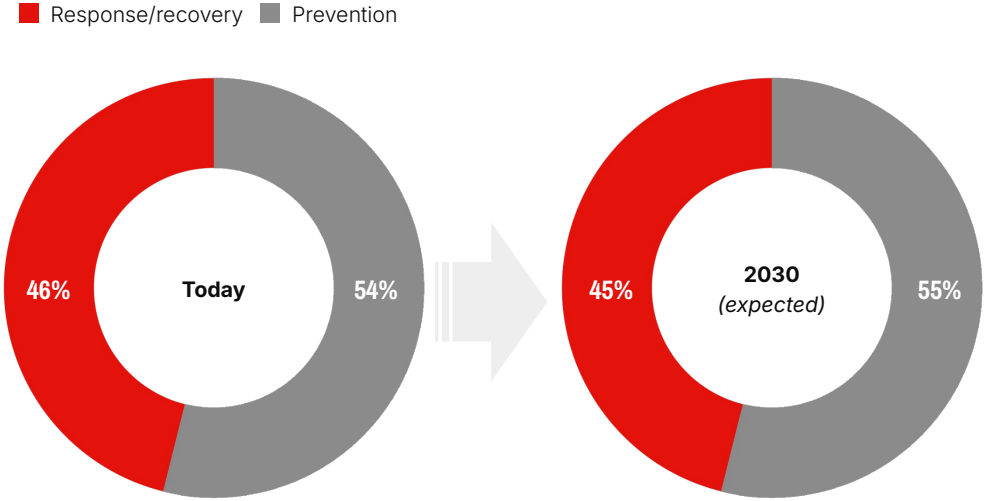

Question: Thinking of your organisation's cybersecurity investment priorities, approximately what share of spending goes to prevention versus recovery or reversibility, both today and as you expect in 2030? Please ensure each row totals 100%. If unsure, please provide your best estimate.

Cybersecurity investment still favours prevention over response and recovery, and leaders expect this imbalance to persist

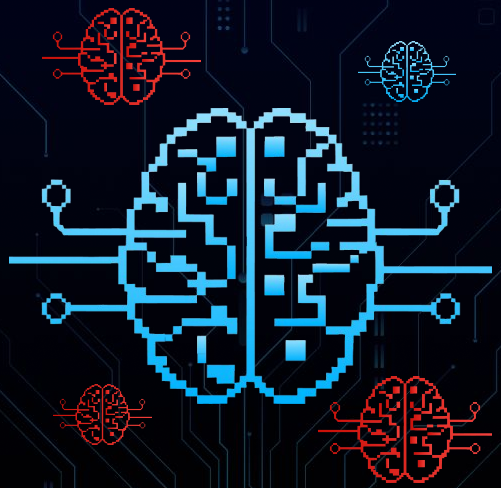

Boards are flying blind

Three-quarters of organisations **do not regularly report** cyber-recovery performance to senior leadership, leaving critical gaps in visibility and decision-making.

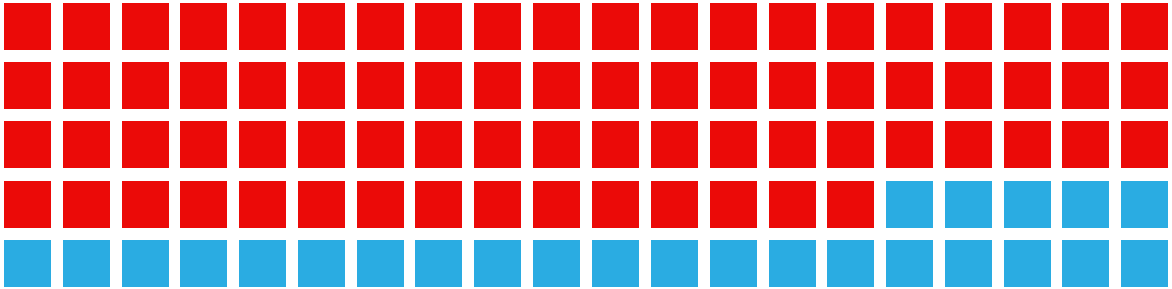

Only 25% actively tracked and reported to the board or C-suite regularly (e.g. quarterly or more).

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? Please ensure each row totals 100%. If unsure, please provide your best estimate.

ECONOMIST ENTERPRISE

To harness the power of agentic, businesses need to leverage three capabilities that define cyber resilience

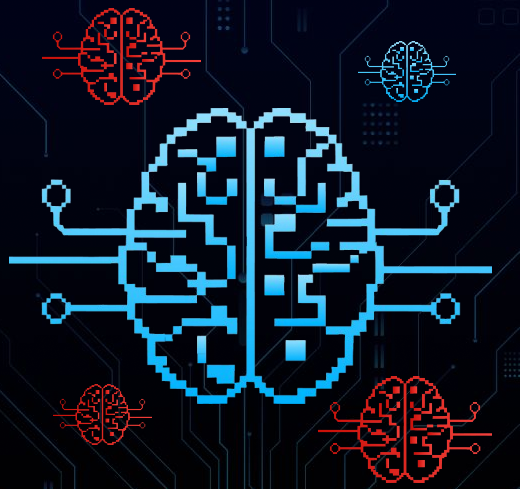

1. Full observability:

Full visibility into organisations' AI agents, ensuring ability to detect, contain and reverse failures when they occur.

2. Rapid response and recovery:

Capacity to quickly detect, contain, and recover from agent-driven incidents before they escalate.

3. Regular testing and validation

Ability to regularly test under real-world conditions the ability to safely reverse or undo an AI agent's actions.

ECONOMIST ENTERPRISE

1. Full observability

Businesses across sectors do not have complete visibility over their AI agents

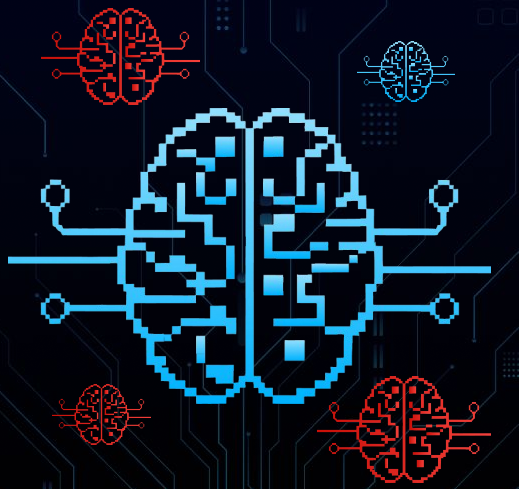

Visibility into AI agents operating within organisational systems is relatively low across sectors, with an average of **37%** saying they maintain a complete and up-to-date inventory of all AI agents, along with their authorised actions.

Another **37%** have a full inventory but lack observability into their actions, without a clear view of what could go wrong.

Visibility into AI agents operating within organisational systems

(All respondents)

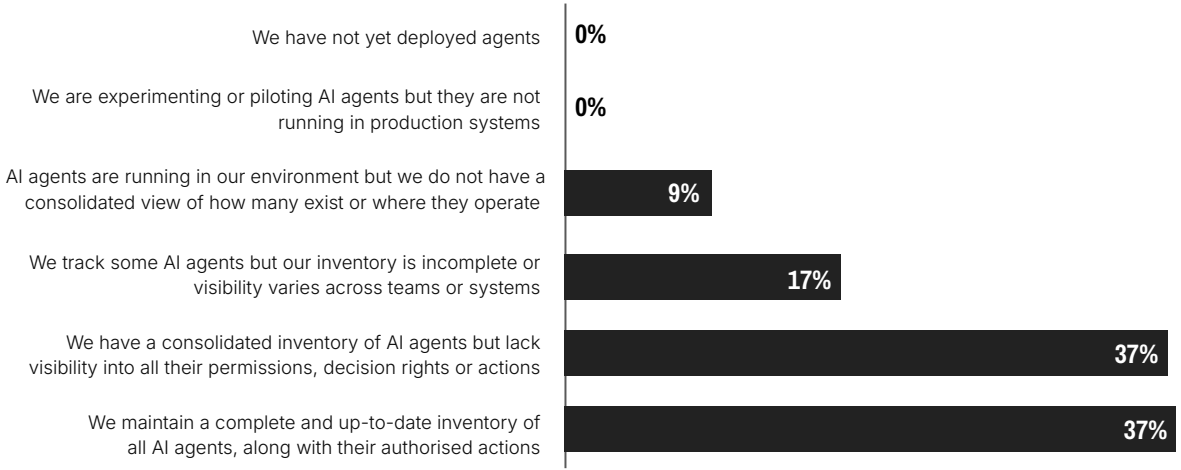

Question: Which of the following best describes your organisation's visibility into AI agents operating within your systems? Please select one.

ECONOMIST ENTERPRISE

1. Full observability

Businesses across sectors do not have complete visibility over their AI agents

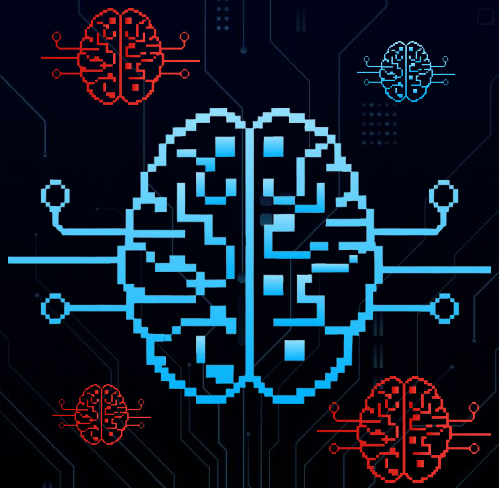

Visibility into AI agents operating within organisational systems is relatively low across sectors, with an average of **37%** saying they maintain a complete and up-to-date inventory of all AI agents, along with their authorised actions.

Healthcare / pharma / life sciences organisations are above average on complete agentic visibility.

Visibility into AI agents operating within organisational systems

(Healthcare / pharma / life sciences respondents)

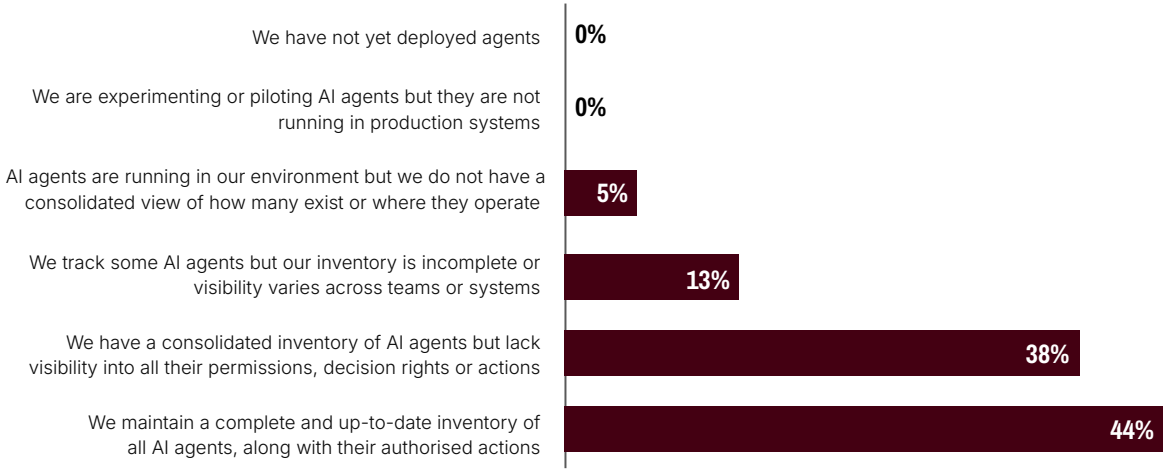

Question: Which of the following best describes your organisation's visibility into AI agents operating within your systems? Please select one.

1. Full observability

Businesses across sectors do not have complete visibility over their AI agents

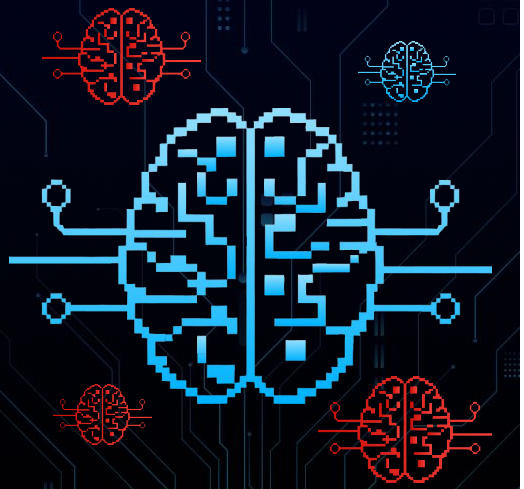

Visibility into AI agents operating within organisational systems is relatively low across sectors, with an average of **37%** saying they maintain a complete and up-to-date inventory of all AI agents, along with their authorised actions.

Infrastructure / construction organisations are above average on complete agentic visibility.

Visibility into AI agents operating within organisational systems

(Infrastructure / construction respondents)

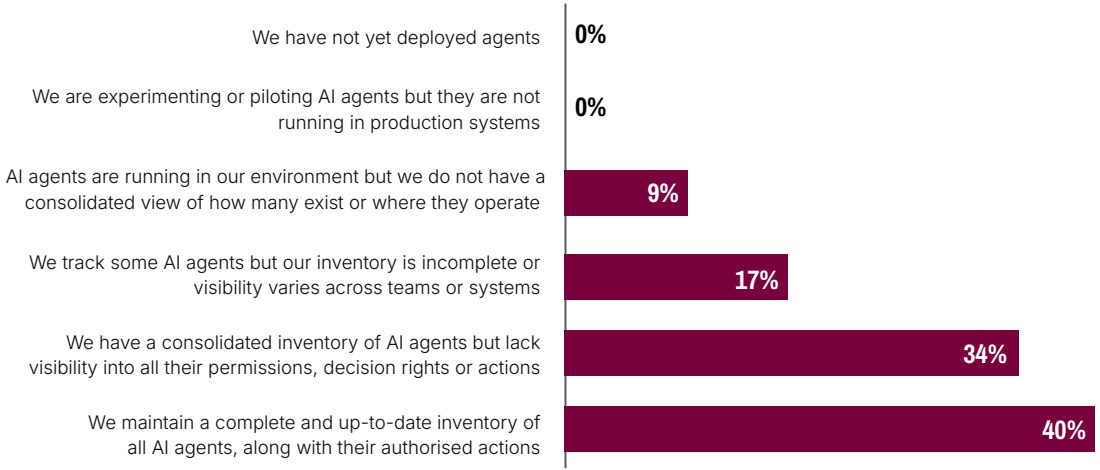

Question: Which of the following best describes your organisation's visibility into AI agents operating within your systems? Please select one.

ECONOMIST ENTERPRISE

1. Full observability

Businesses across sectors do not have complete visibility over their AI agents

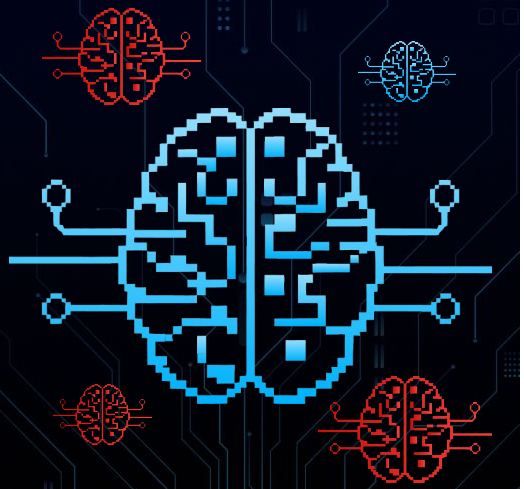

Visibility into AI agents operating within organisational systems is relatively low across sectors, with an average of **37%** saying they maintain a complete and up-to-date inventory of all AI agents, along with their authorised actions.

Technology organisations are below average on complete agentic visibility.

Visibility into AI agents operating within organisational systems

(Technology respondents)

Question: Which of the following best describes your organisation's visibility into AI agents operating within your systems? Please select one.

ECONOMIST ENTERPRISE

1. Full observability

Businesses across sectors do not have complete visibility over their AI agents

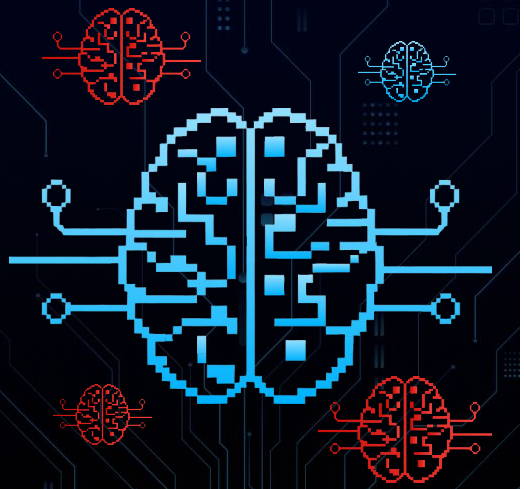

Visibility into AI agents operating within organisational systems is relatively low across sectors, with an average of **37%** saying they maintain a complete and up-to-date inventory of all AI agents, along with their authorised actions.

Financial services / banking / insurance organisations are above average on complete agentic visibility.

Visibility into AI agents operating within organisational systems

(Financial services / banking / insurance respondents)

Question: Which of the following best describes your organisation's visibility into AI agents operating within your systems? Please select one.

1. Full observability

Businesses across sectors do not have complete visibility over their AI agents

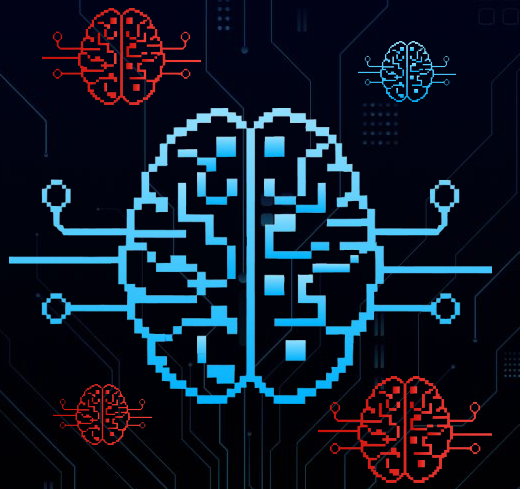

Visibility into AI agents operating within organisational systems is relatively low across sectors, with an average of **37%** saying they maintain a complete and up-to-date inventory of all AI agents, along with their authorised actions.

Public sector organisations are below average on complete agentic visibility.

Visibility into AI agents operating within organisational systems

(Public sector respondents)

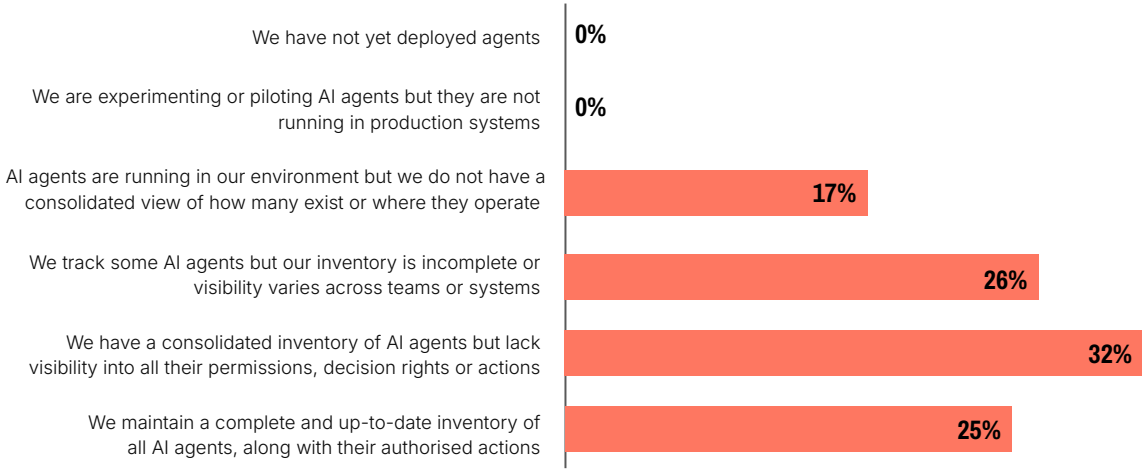

Question: Which of the following best describes your organisation's visibility into AI agents operating within your systems? Please select one.

ECONOMIST ENTERPRISE

1. Full observability

Businesses across sectors do not have complete visibility over their AI agents

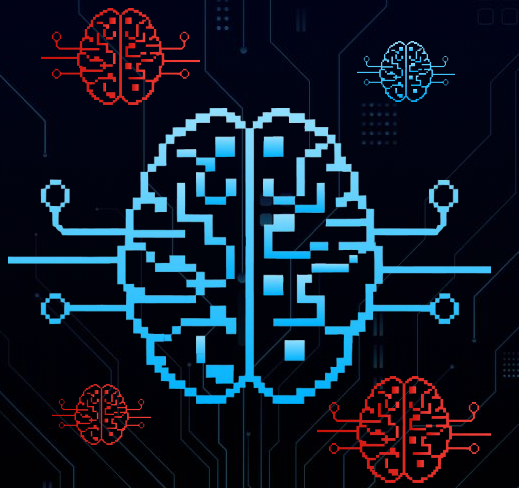

Visibility into AI agents operating within organisational systems is relatively low across sectors, with an average of **37%** saying they maintain a complete and up-to-date inventory of all AI agents, along with their authorised actions.

Retail / consumer goods organisations are below average on complete agentic visibility.

Visibility into AI agents operating within organisational systems

(Retail / consumer goods respondents)

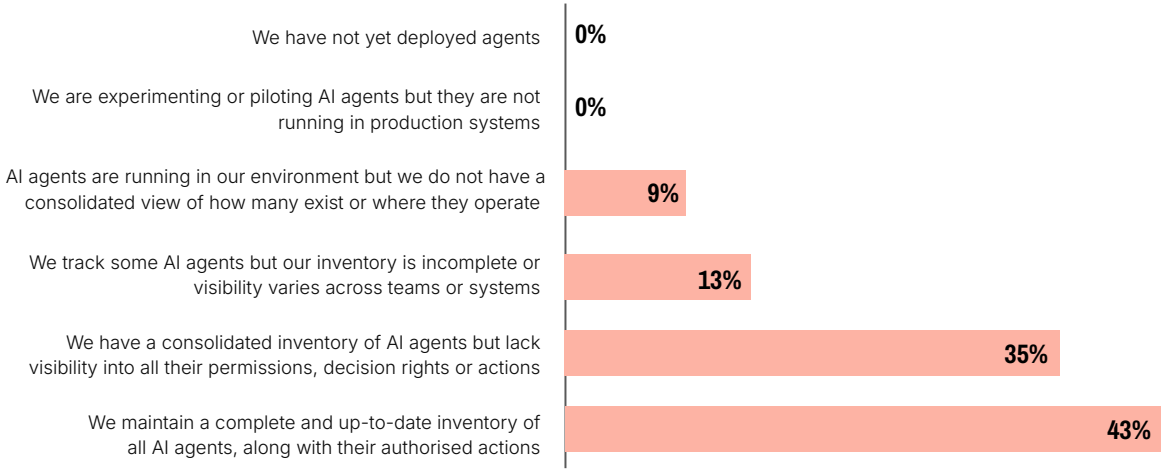

Question: Which of the following best describes your organisation's visibility into AI agents operating within your systems? Please select one.

2. Ability to respond

Rollback capabilities may exist but they remain insufficient

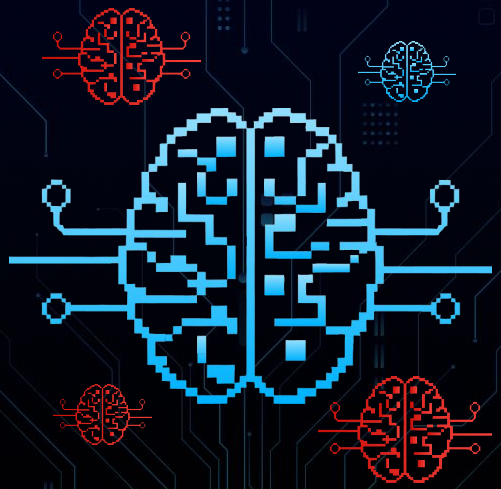

Almost a third (**30%**) of businesses overall have robust, fully-tested capabilities that allow them to reliably reverse or roll back harmful actions taken by AI agents and restore affected systems.

But the remainder have rollback mechanisms that are either **incomplete, untested** or **do not yet cover all agents or all incident types**.

Level of 'reversibility' for cybersecurity incidents involving AI agents

(All respondents)

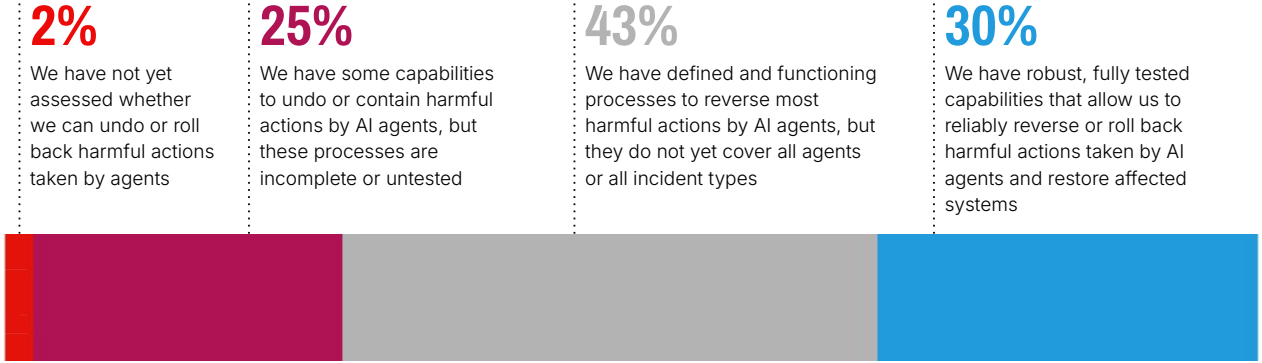

Question: Which statement best reflects your organisation's current level of 'reversibility' for cybersecurity incidents involving AI agents? Please select one.

ECONOMIST ENTERPRISE

3. Testing capacity

Most firms have defined their MVB but not all have tested

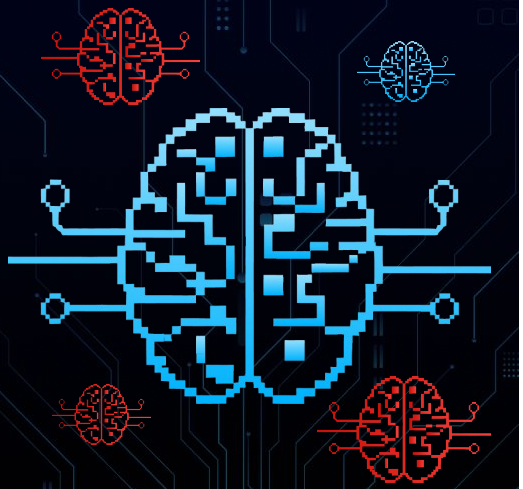

73% of businesses overall have defined their minimally viable business (MVB), highlighting recognition of the potential impact of a cybersecurity incident and the need to know the critical systems, data and operations required to maintain essential business functions in the aftermath.

But, of these only **40%** of businesses have also tested their MVB under real world conditions, with the rest unsure of whether their configurations will actually work in the case of an incident due to a lack of testing.

Progress towards establishing a minimally viable business (MVB) configuration

(All respondents)

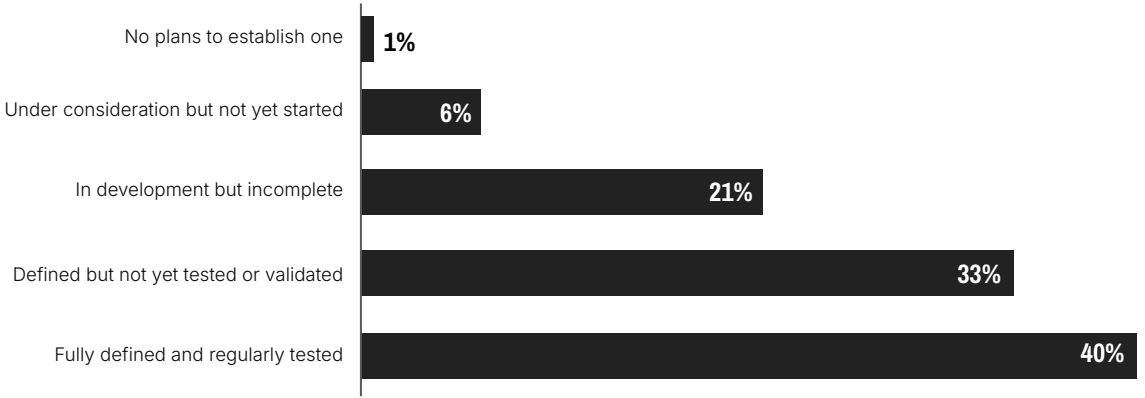

Question: Which of the following best describes your organisation's progress towards establishing an MVB configuration? Please select one.

ECONOMIST ENTERPRISE

4. Who is most prepared today?

Preparedness and confidence to rapidly detect, contain and roll back harmful agent actions vary across sectors

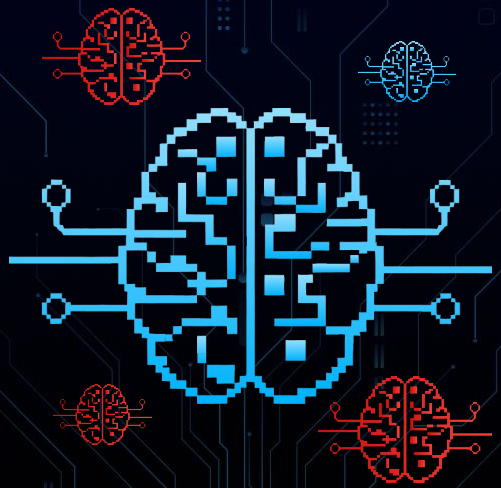

Regulated sectors tend to have stronger governance, while experimental sectors lag behind.

Healthcare
Most tested, least transparent

Infrastructure
Most exposed, least monitored

Technology
Most ambitious, least hardened

Financial services
Most visible, least automated

Public sector
Most threatened, least prepared

Retail
Most confident, least equipped

ECONOMIST ENTERPRISE

Healthcare organisations are the most tested, but least transparent

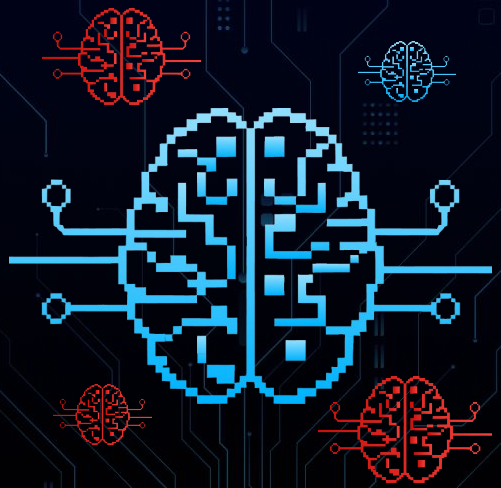

Healthcare

Most tested, least transparent

- More 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.

Infrastructure organisations are most exposed to risks and least monitored

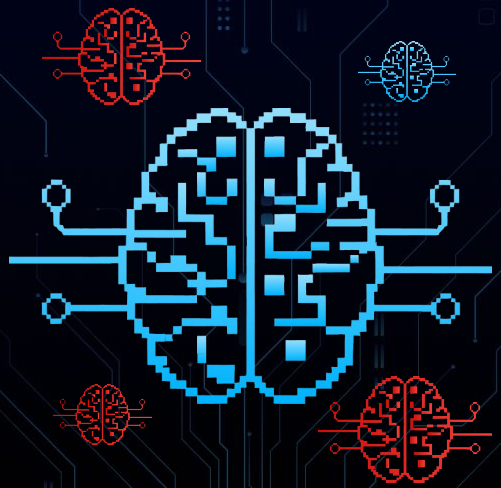

Infrastructure

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 CTO, 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.

Tech firms are naturally the most ambitious but least resilient

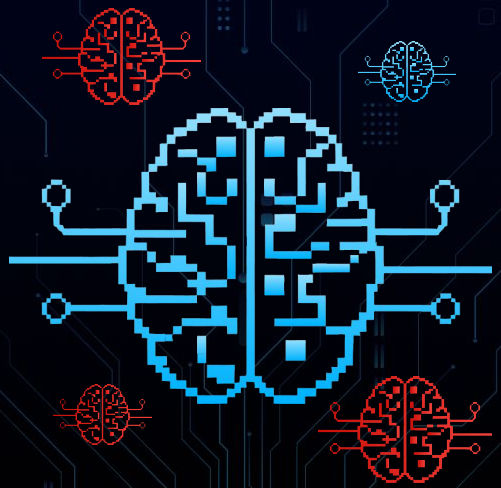

Technology

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, below every sector except public.
- 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.

Financial services lead in visibility but are the least automated

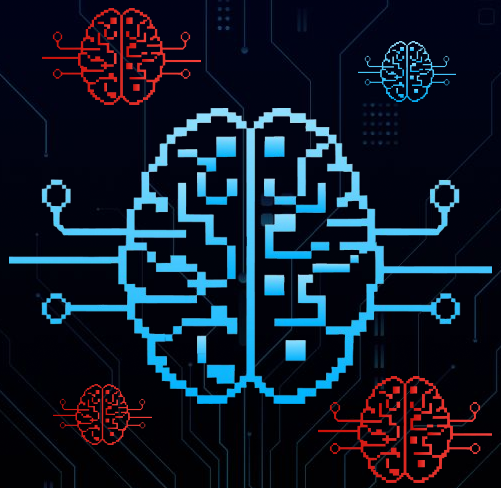

Financial services

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 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.

Public sector organisations are most threatened but least prepared

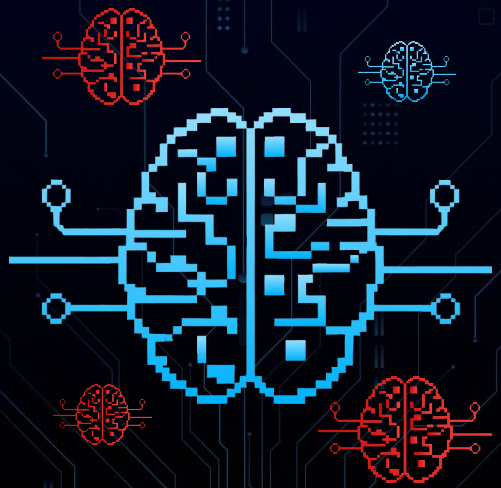

Public sector

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.

Retailers lead in confidence but are the least equipped, lacking expertise

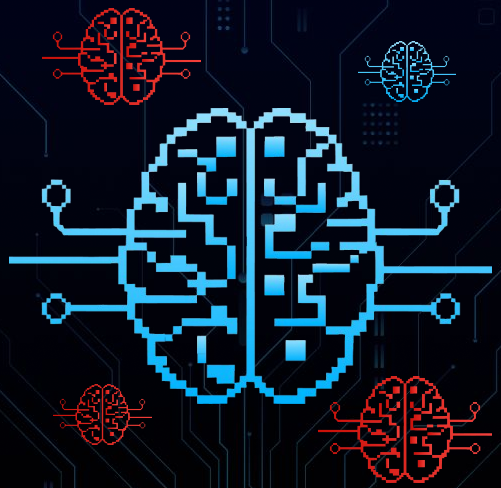

Retail

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 rates 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.

ECONOMIST ENTERPRISE

Thank you

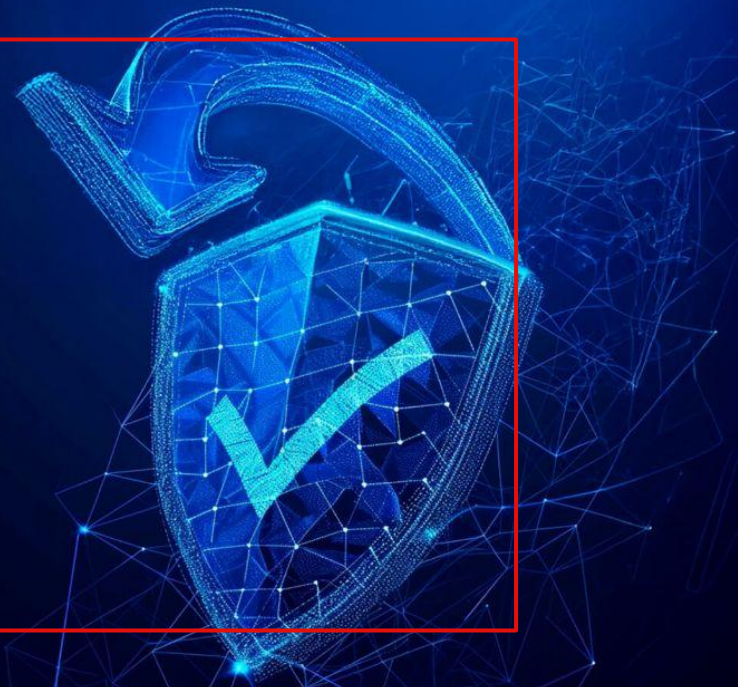

If you have any questions, please contact:

vaibhavsahgal@economist.com
ailiahaider@economist.com