

RIVERSAFE RESEARCH

AI GOVERNANCE IN UK RETAIL

Findings from 200+ cyber and security leaders across UK retail.

Research conducted by Censuswide, July and August 2026.

FOREWORD

A sector adopting fast, with governance catching up

AI has moved into retail operations quickly. Across the businesses we work with, tools are being adopted, connected and relied on at pace, reaching into supply chains, pricing, checkout and customer data. That appetite is real, and the commercial logic behind it is sound.

This research shows adoption running ahead of the governance around it. Tools are in production before the controls, ownership and assurance that make them safe to run have caught up. That reflects how fast the technology has arrived more than any failing on the part of security teams, and the work to close the distance is already underway.

We commissioned this research to give retail cyber and security leaders an honest benchmark against their peers. 200+ senior leaders took part and were candid about where they stand.

Alongside the survey, we regularly get retail security leaders in a room together under Chatham House rules, where people can be honest about what is actually working and what isn't. A lot of what gets said there puts flesh on the numbers, so where it helps, we've woven those conversations into this report, marked as coming from the room.

The findings are self-reported, and we present them that way. Where an answer is a matter of judgement, such as how serious an incident felt, we say so.

If any of it prompts a conversation you want to have, we are happy to talk.

Naveen Jalagadugu
Services Director, RiverSafe

EXECUTIVE SUMMARY

Adoption is outpacing governance

UK retail security leaders describe a sector where AI is everywhere and moving fast. It is in production, the pressure to keep adopting is near universal, and incidents have already happened. Underneath that activity, the governance foundations of documented review, clear ownership and control over what AI can reach are only partly in place. The distance between how quickly AI is being adopted and how completely it is governed is the thread that runs through this report.

Despite that gap, security leaders remain confident about what comes next, with 93% believing AI-related risk can be brought down to an acceptable level. The sector reads this as a solvable problem, not an unmanageable one.

94%

have experienced an AI-related security incident in the past 12 months

100%

are running AI tools with no formal security approval

99%

feel pressure to adopt AI faster than governance can keep up

88%

do not put every AI use case through a documented review

Four things retail leaders should take from this

1. Shadow AI is the baseline, and incidents follow it.

Unapproved AI is not an edge case waiting to be found. It is the working condition, at around ten tools per business, and the incident data sits right alongside it. Most retailers have already had an AI-related incident with real-world impact. Tools running outside an approval path are the ones least likely to be caught when something goes wrong, which is what makes the two findings inseparable.

2. The gap is in the scaffolding, not the will to adopt.

Almost no one is short of appetite. What lags is the structure around it. Most retailers review only some of their AI before it goes live, only a third use dedicated governance tooling, and just over half fully control what their agents can access. The pressure leaders describe is less about convincing anyone to move and more about keeping the controls level with how fast AI is arriving.

3. Ownership of AI risk is fragmented.

Asked who owns AI risk where execution is shared, no function holds a majority. Technology and IT is named most often, with security, legal, GRC, data and risk each taking a share. Leaders are not describing shared ownership. They are each pointing to a different team, which is how risk ends up sitting between functions rather than with any one of them.

4. Spending has followed regulation. The priorities point elsewhere.

The largest share of AI security investment is going to regulatory and legal review, with compliance tooling next. Detection and response sits lower. Looking ahead, that order flips, and detection and response becomes the single most-cited priority for the year ahead. Spend has tracked the clearest external obligation, while the stated priorities point at where leaders know they are thinnest.

ABOUT THIS RESEARCH

Methodology

This research was conducted by Censuswide between July and August 2026, among 200 senior cyber and information security leaders in UK retail. Respondents held roles at CISO, Director, Head of and BISO level.

All respondents worked at organisations with a turnover of £250 million or more and at least 500 employees. Responses were anonymous and are reported in aggregate. Censuswide is a member of the Market Research Society and the British Polling Council, and a signatory of the Global Data Quality Pledge.

Qualitative context is drawn from a RiverSafe retail security leaders roundtable held under Chatham House rules in July 2026. Those observations are anonymised by design and are used to illustrate the data.

A NOTE ON HOW TO READ THIS DATA

The findings are self-reported and describe how leaders assess their own position. Where an answer option is a matter of judgement, for example how serious an incident felt, we treat it as the respondent's assessment rather than an audited measure.
Percentages are rounded to whole numbers, so some figures may not sum to exactly 100.

SECTION 01

Adoption is outpacing governance

AI is in production across the sector, and the pressure to keep adopting is near universal

1.1

Shadow AI is everywhere

Every respondent estimated that their business runs at least some AI tools without formal security approval. The average estimate is around ten tools per organisation, and 9% put the figure above twenty. Not one leader reported zero.

How many AI tools do you estimate are being used across your business without formal security approval?



Shadow AI is not a fringe issue waiting to be discovered. Every respondent could estimate how much unapproved AI they are running, but an estimate is not an inventory.

Knowing roughly how many tools are in use is a long way from knowing which ones, what data they touch, and what they connect to. Continuous discovery, finding and mapping AI that sits outside any approval path, is the foundation the rest of governance depends on. Until that picture is accurate, nothing downstream, from risk assessment to incident monitoring, can be trusted to be complete.

1.2

Pressure to move faster than governance

Asked whether they feel pressure to approve or accelerate AI adoption faster than their governance and risk processes can keep up with, 99% said yes. Opinion then splits almost evenly on whether it is under control. 48% say the pressure is already creating real governance gaps, and 51% say they are managing to keep pace so far.

Do you feel under pressure to approve or accelerate AI adoption faster than your governance and risk processes can keep up with?



Pressure to adopt AI ahead of governance is prevalent in UK retail, not a problem confined to less mature organisations. Nearly half of leaders say the gap is already open, not a risk that might arrive later.

FROM THE ROOM: RIVERSAFE RETAIL SECURITY ROUNDTABLE

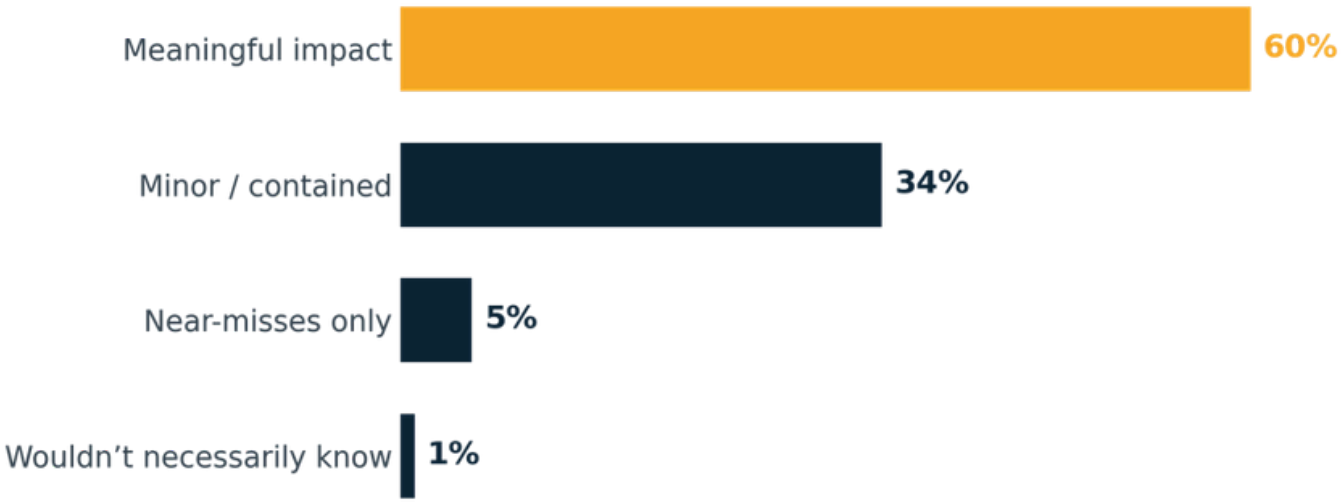
This matched what leaders told us in person. The recurring theme was boards were pushing for faster AI adoption while expecting risk appetite to fall at the same time, with security often learning of a decision after it had been made. Several described large AI rollouts being committed to before security had scoped them, then having to slow and strip back the deployment to make it safe.

1.3

Incidents have already happened

The majority of retailers have already dealt with an AI-related incident. Asked whether they had experienced one with real-world impact on customers, staff, data, operations, supply chain or compliance in the past 12 months, 94% said yes. Close to 60% reported an incident with meaningful impact, and a further 34% one that was minor or contained.

In the past 12 months, has your organisation experienced an AI-related security incident with real-world impact (on customers, staff, data, operations, supply chain, or compliance)?



READING THIS NUMBER

The split between meaningful and minor impact is the respondent's own judgement, so we present the 60% as how leaders characterise their experience rather than as a calibrated severity measure. For most retailers, an AI incident with real-world impact is something that has already happened, not a risk on the horizon.

SECTION 02

The foundations are only partial

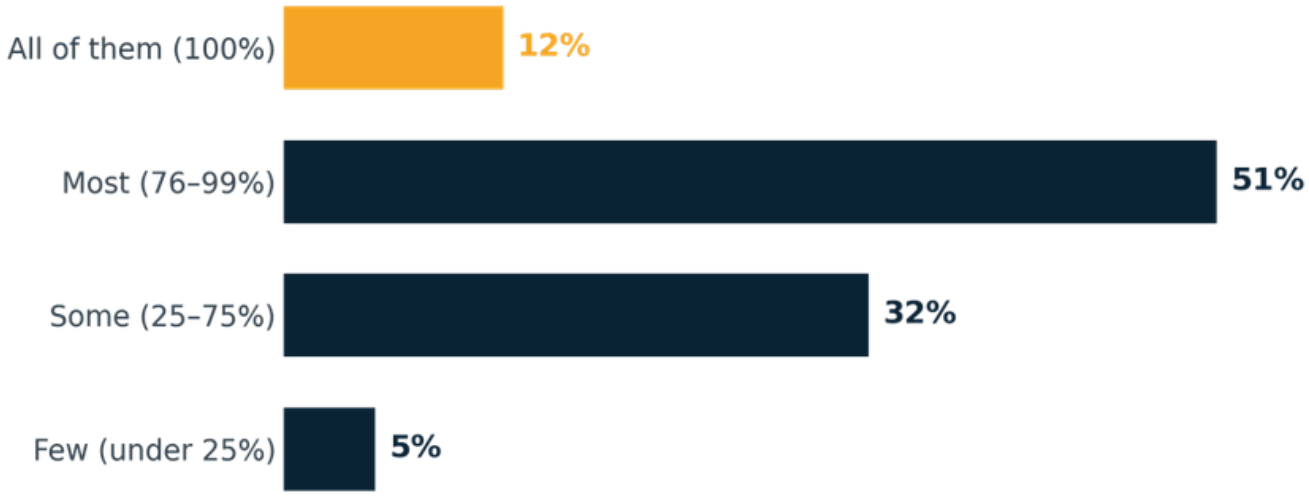
On the core controls, the fully in place answer rarely clears half

2.1

Documented review before deployment

88% of retailers do not put every AI use case through a documented security or risk review before it goes into production. Only around one in eight review all of them. Half review most, and roughly a third review three-quarters or fewer.

What proportion of AI use cases in production have been through a documented security or risk review before deployment?



The use cases most likely to skip a review are the quick, informal ones a team stands up for itself, and those often reach the data and systems that matter most. Coverage that is strong but rarely complete usually means the review process was built for a slower rate of change than AI now runs at. A lighter checkpoint that every use case has to clear will close that gap faster than a heavier gate that teams find ways around.

2.2

How AI governance is run day to day

No single operating model dominates. Just over a third have adapted existing GRC tooling to cover AI, a third use a dedicated AI governance platform, and 30% still run governance on spreadsheets and manual processes.

How does your organisation currently manage AI governance day to day?



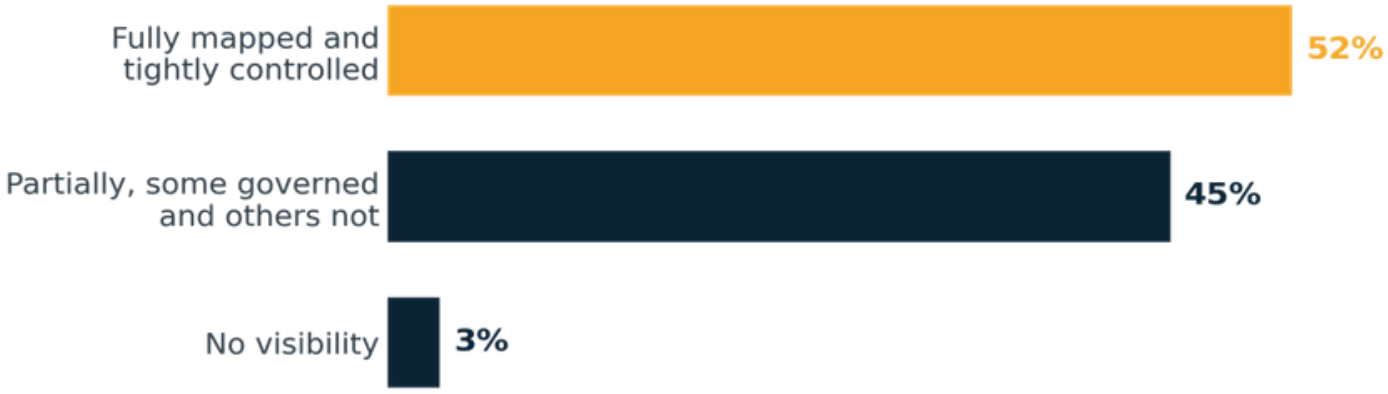
For organisations of this size and AI exposure, nearly a third managing AI risk by spreadsheet is a real capacity gap. It tends to hold up at low volumes and strain as the number of AI use cases grows. On the evidence of Section 1, that number is already growing.

2.3

Control over what AI can access

Just over half, 52%, say they have fully mapped and tightly controlled what their AI agents and tools can access without human sign-off. Another 45% govern some agents but not others, and 3% have no visibility at all.

Do you have clear visibility and control over what your AI agents and tools can access, systems, data, and actions, without human sign-off?



Agentic AI is only as safe as the boundaries around what it can reach and do. With just over half fully in control of those boundaries, the sector divides roughly evenly between organisations that have this locked down and those with partial coverage. That is a meaningful exposure as agents connect to more systems.

FROM THE ROOM: RIVERSAFE RETAIL SECURITY ROUNDTABLE

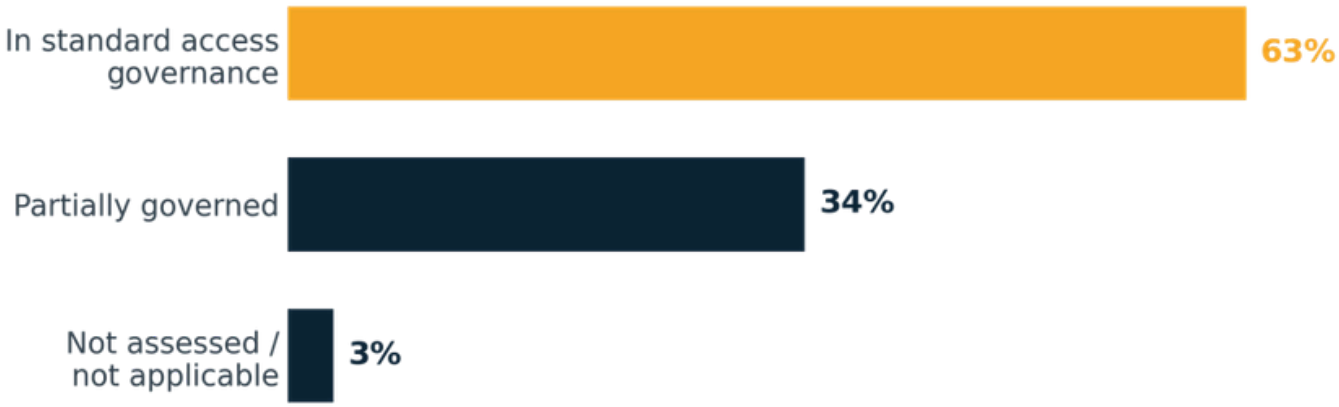
In practice, much of this AI is running as a service account. It sits inside existing systems, nobody asks it for permission, and nothing gates what it can do once it is asked. Leaders were candid that a front door for AI access is not yet built into how these tools are deployed, and several felt that gating should be native to the system rather than bolted on afterwards. A control that only filters inputs against a checklist is closer to data loss prevention than to governance of an autonomous agent.

2.4

Governing AI access into core systems

As AI tools connect to core retail systems such as inventory, pricing, checkout and customer data through APIs, 63% say that access sits within their standard access governance. The remaining 37% have gaps. 34% govern some integrations but not others, and a small group have not specifically assessed this at all.

As AI tools connect to core retail systems via APIs, is that AI-driven API access included in your standard access governance?



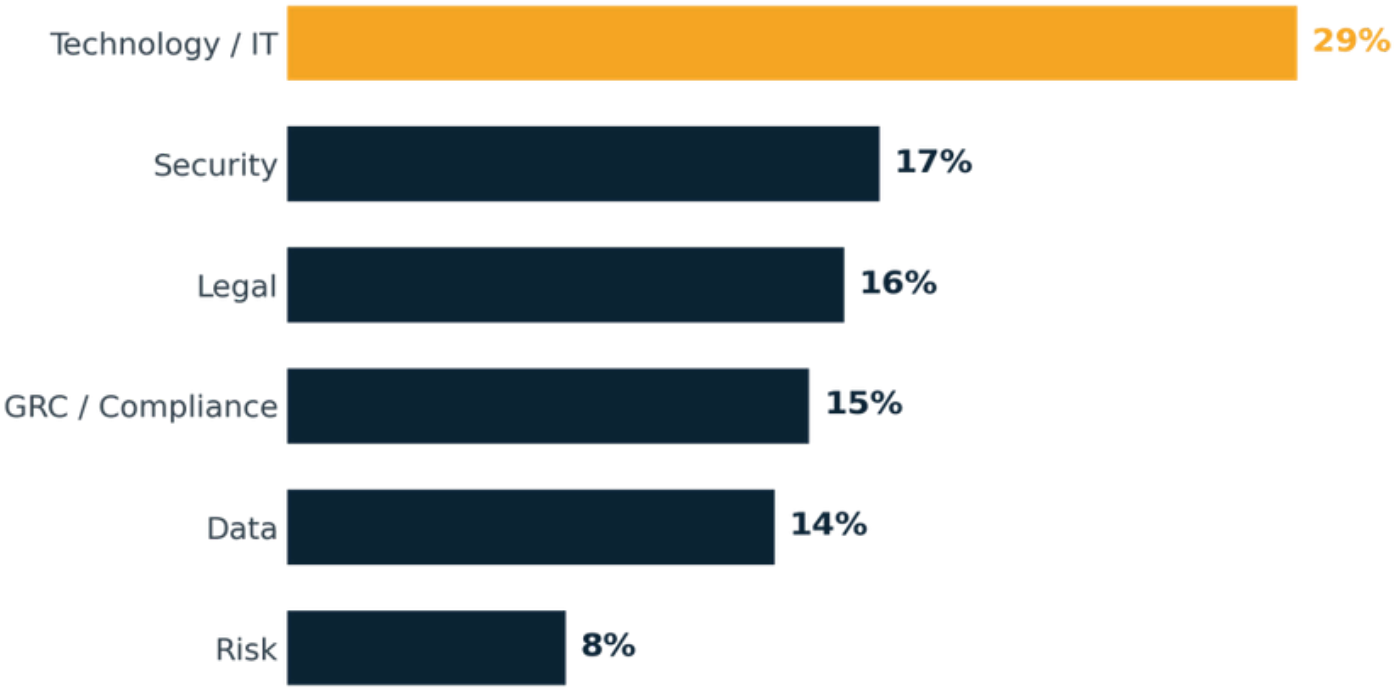
This is where AI risk becomes retail operational risk. An ungoverned integration into pricing or checkout is not an abstract policy gap, it is a direct route into the systems the business runs on.

2.5

Who owns AI risk

Technology and IT is the most common answer at 29%, but security, legal, GRC, data and risk each take a share, and none holds a majority.

Who owns AI risk in your organisation, even where execution is shared across teams?



Ownership of AI risk varies across the sector rather than settling on one function. Technology and IT is the most common answer at 29%, but no function holds a majority, and security, legal, GRC, data and risk each take a share.

In practice this means the answer to "who owns AI risk here" depends on the organisation you ask.

Technology leads, but there is no sector-wide default, so where accountability sits is something each retailer is still setting for itself rather than inheriting from a common model.

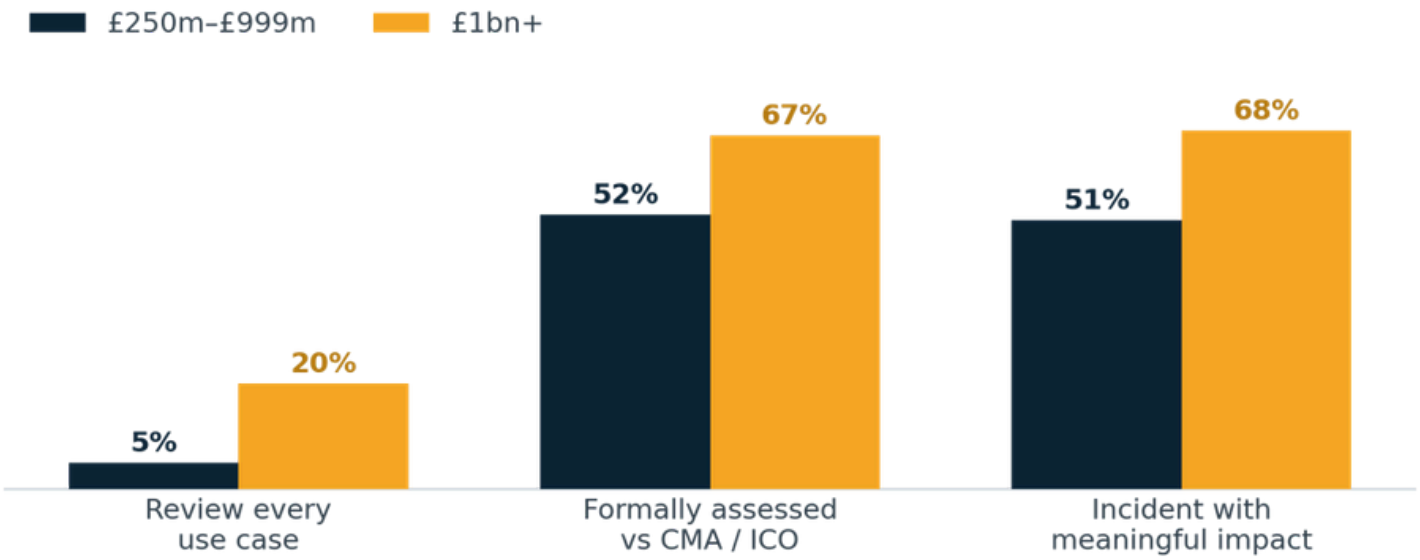
SECTION 03

The picture shifts with size

Larger retailers govern more, and absorb more risk

Splitting the sample by turnover gives two groups of 100, one at £250 million to £999 million and one at £1 billion or over.

Larger retailers are further along on documented governance. 20% of those at £1 billion or over review every AI use case before deployment, against 5% of the smaller group. 67% have formally assessed their AI against CMA and ICO requirements, against 52%.



Exposure scales alongside maturity. 68% of retailers at £1 billion or over report an AI incident with meaningful impact, against 51% of the smaller group.

The most-resourced players in the sector are not outrunning the risk. They are absorbing more of it, which is consistent with having more AI in production, a wider attack surface, and better means to notice when something goes wrong.

WHAT THIS MEANS IF YOU SIT IN THE SMALLER BAND

A lower incident figure is not automatically reassurance. It can reflect less AI in production, or less capacity to detect an incident when it happens. The gap on documented review, at 5% against 20%, is the more actionable signal. It points to review coverage as the first place a mid-market retailer can close ground on larger peers without a step change in budget.

SECTION 04

Investment today, priorities tomorrow

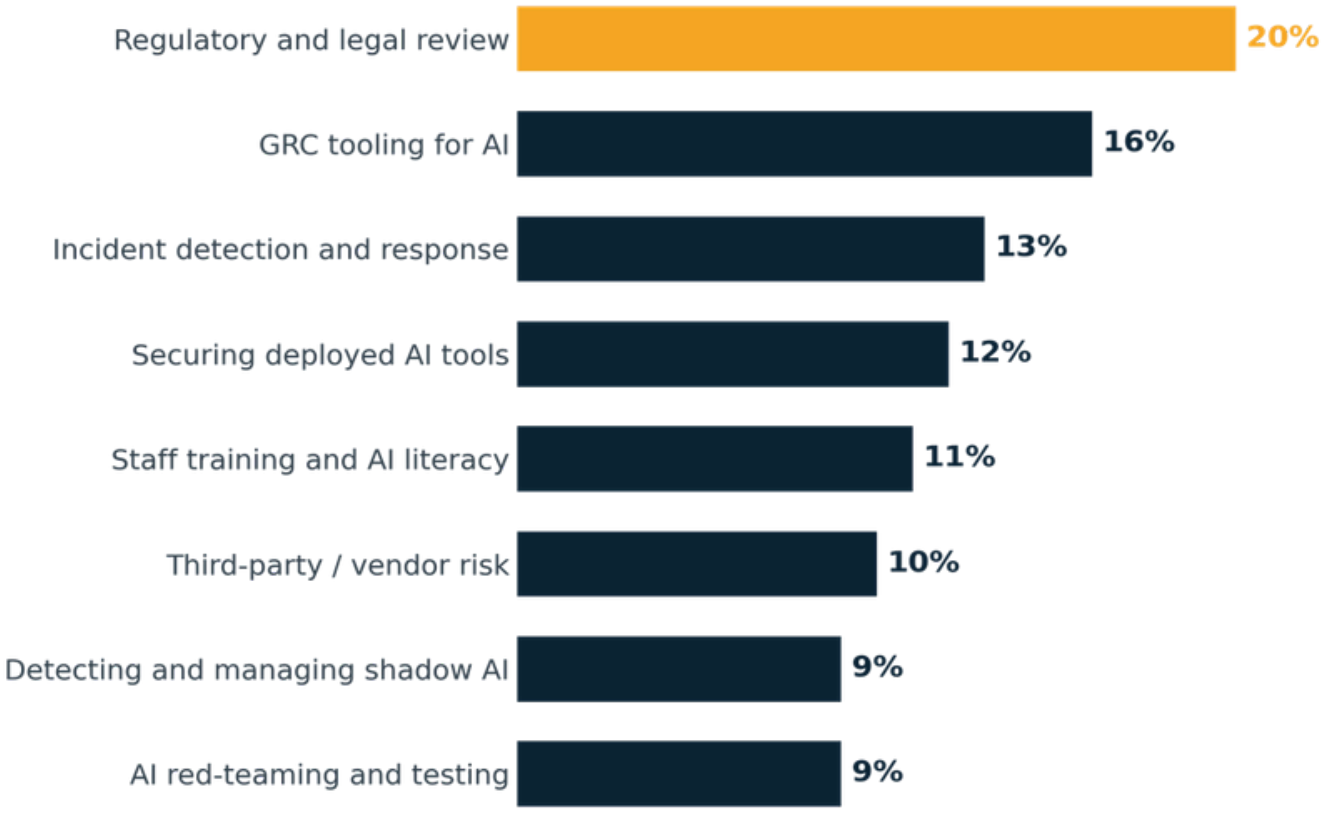
Spending has followed the clearest obligation. The priorities point somewhere else

4.1

Where the largest share of investment goes

Asked where the largest share of new AI-related security investment is heading over the next 12 months, regulatory and legal review leads at 20%, followed by GRC tooling at 16%. Incident detection and response sits third at 13%.

Where is the largest share of your new AI-related security investment going over the next 12 months?



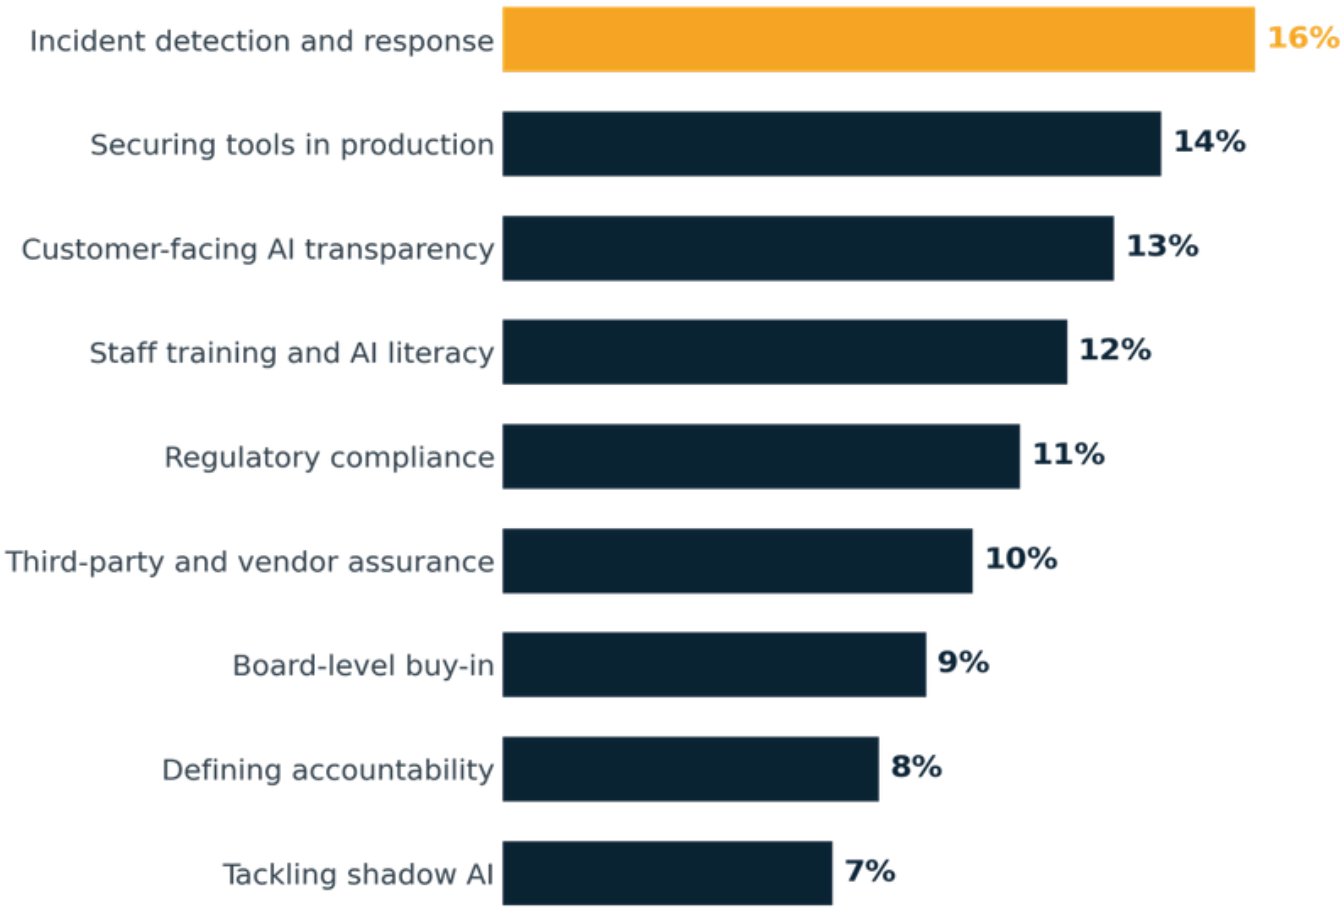
With CMA agentic AI guidance and ICO automated decision-making rules now in force, weighting investment toward regulatory readiness is a rational response to a real obligation. It also means the largest share of spend is currently going to compliance work rather than to operational detection.

4.2

The biggest priority for the year ahead

Priorities are spread. No single answer clears 16%. Incident detection and response tops the list as the biggest AI governance priority for the next 12 months, ahead of securing tools already in production and customer-facing AI transparency. Board-level buy-in sits near the bottom at 9%.

What is your single biggest AI governance priority for the next 12 months?



The stated priority, detection and response, is not where the largest share of money is going, with regulation and legal review taking the largest share.

4.3

Why funded and prioritised pull apart

The gap between what is funded and what is prioritised is not a contradiction. It reflects how each case gets made to a board.

Regulatory work is straightforward to fund. The obligation is external, dated and specific. CMA and ICO requirements exist, they apply, and the cost of ignoring them is legible to a board. Detection and prevention are harder to fund because the return is the absence of an incident. Justifying spend against something that has not happened, and may not, is a weaker pitch in a budget round than a named regulator with a live requirement.

That is why detection can be the honest top priority and still lose the funding race to compliance. The way to close that gap is to make the risk concrete, which is where the value of a tested scenario comes in rather than an abstract argument about resilience.

SECTION 05

Third party, regulation and outlook

The supplier route, regulatory readiness, and how the sector reads its own odds

5.1

AI in the third-party risk process

Recent attacks on UK retailers have largely originated through third-party suppliers, so respondents were asked how far their vendor and supplier risk process covers AI. 60% say they have a fully updated, AI-specific vendor risk process. Roughly 40% have only partially adapted it, or have not adapted it for AI at all.

Recent attacks on UK retailers have largely originated through third-party suppliers. How far has your vendor/supplier risk process been updated to cover AI-specific risk?

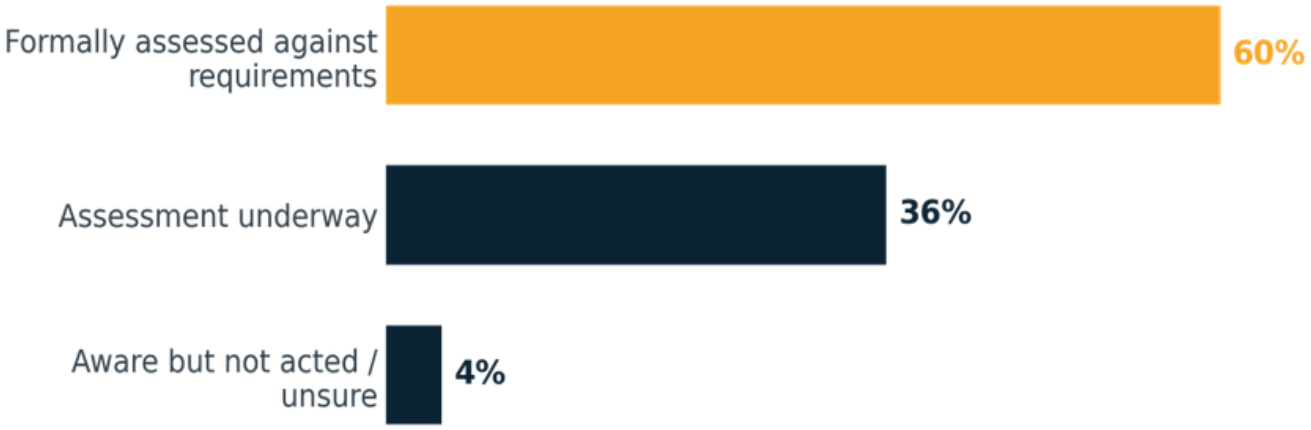


5.2

Response to AI-specific regulation

On current AI-specific regulatory requirements affecting retail, meaning CMA agentic AI guidance and ICO automated decision-making rules, 60% have formally assessed their AI systems against them. Another 36% have an assessment underway, and a small remainder are aware but have not yet acted.

Have you formally assessed your AI systems against current AI-specific regulatory requirements affecting retail (e.g. CMA agentic AI guidance, ICO automated decision-making rules)?



The direction is right and most organisations are engaged, but roughly 40% have not yet completed a formal assessment against requirements that already apply. Mapping which instruments apply to the AI tools currently in use is the practical starting point for closing that gap.

5.3

Retail security leaders back themselves to solve this

The picture in this report is one of adoption running ahead of governance, but it closes on a striking note of confidence.

Asked whether AI-related security risk in retail can be brought down to an acceptable level even with strong governance in place, 93% said yes. It means retail security leaders are not treating AI as an unmanageable threat to be contained, but as a problem they can get on top of with the right foundations.

The gaps this research identifies, in review coverage, ownership, agent control and detection, are real, but they are the practical, closeable kind and are being tackled. The sector knows what the work is and is working to get it done.

RIVERSAFE POINT OF VIEW

What this means for retail security leaders

The data describes a sector that has adopted AI quickly and is honest about the governance still being built around it. Five priorities stand out for the year ahead, drawn from the findings and from what leaders tell us in practice.

1. Bring known shadow AI under management.

Every respondent is running unapproved AI tools, so the task is not discovery but control. Build a defined route by which AI gets registered, risk classified, approved and monitored, and apply it retrospectively to what is already live. Treat this as a standing capability with a named owner, not a one-off audit.

2. Govern AI as an identity, not a feature.

The most useful mental shift we see is to stop treating an AI agent as part of an app and start treating it as its own identity with its own access profile. It gets provisioned, scoped to a purpose, and retired, the same way a human account does. That reframing is what turns the 37% with gaps in AI-to-system access into a solvable access governance problem rather than an open-ended one.

3. Put a real front door in front of AI access, and keep maintaining it.

A control that only checks inputs against a list is closer to data loss prevention than to governance. What holds up is a gate embedded in the system that decides what an agent can reach and do, scoped to the job it was brought in for. For frontline and store-scale rollouts, define a tighter permission profile for those roles rather than inheriting head office access. Permission boundaries are not set once. They need maintaining as tools, integrations and use cases change.

4. Fund detection against a scenario, not an abstraction.

Detection and prevention lose budget rounds to regulation because their return is the absence of an incident, which is hard to evidence. Make it concrete. A tabletop exercise that traces how an AI-linked incident would actually surface in your environment does two things at once. It shows where the gaps in tooling and process sit, and it gives the board a specific, credible thing to fund.

5. Treat AI apps and agents as their own attack surface.

The controls that matter most sit around the AI itself: what an app or agent can reach, what it is allowed to do, and what it connects to.

Run a security risk assessment on each AI application and agent the way you would any system with access to sensitive data, then set guardrails and access controls that match its purpose rather than inheriting broad default permissions.

Extend the same rigour to the integrations. Every API an AI tool opens into a core retail system is a route in that needs assessing in its own right, not assumed safe because the tool itself was approved. As agents connect to more systems and to each other, this is the layer that keeps AI risk from quietly becoming operational risk.

RIVERSAFE

ABOUT RIVERSAFE

A cyber security consultancy working across Retail.

RiverSafe works with retailers across the UK, giving security teams both the advisory perspective to know what to prioritise and the technical hands to deliver it.

That means securing the applications you build and ship, engineering detection so threats surface earlier and with less noise, hardening the wider environment, and keeping tight control over who has access to what.

The focus is on outcomes that hold up in practice, backed by ISO 27001, ISO 20000-1, ISO 9001, ISO 14001 and Cyber Essentials certification.

If any of the themes in this report match what you're working through, we're glad to talk.

www.riversafe.co.uk