CultureAI

Research

The State of Enterprise AI Usage:

The Illusion of Control

Contents:

Executive Summary

Generative AI is no longer an emerging experimental technology in enterprise environments. It is deeply embedded, decentralised, and shaping how core business functions now operate. However, governance, security, and compliance models have not evolved nearly as fast, and as a result, there is a growing gap between how AI is used in practice and how organisations believe it's being controlled.

To test whether perceived control reflected operational reality, CultureAI conducted a Q4 2025 survey of 300 senior technology, security, and risk leaders. The findings expose a widening gap between how AI is governed and how it is actually used.

What the Results Showed

AI usage is already widespread and accelerating:

of organisations report AI is widely used across teams.

expect AI usage to increase over the next 12 months.

In practice, however, that confidence is often misplaced:

of organisations have still identified unauthorised shadow AI usage.

AI risk is acknowledged, but rarely escalated:

of organisations rate AI risk as moderate or lower.

Despite this rapid, decentralised adoption, confidence in control remains high:

of organisations believe they have full visibility into AI usage.

These signals demonstrate a structural disconnect between adoption, perception, and operational control. The sections that follow throughout this document explore where this disconnect originates, why risk is difficult to detect and quantify, and how existing governance models break down in practice.

The Core Conclusion

Policies and oversight structures establish intent, but they do not consistently shape behaviour in day-to-day use. For security and governance leaders, this raises doubts about whether existing AI governance models can operate at the point where (and when) AI-based risks are actually created.

Introduction & Methodology

Between 2023 and 2026, generative AI moved rapidly from novelty to normality in both enterprise and mass-consumer environments, fundamentally changing how knowledge-based work is performed and forcing organisations to reconsider long-held assumptions about data governance and user behaviour. The speed of this shift was exceptional, with AI tools being simple to use, easy to trial, and increasingly embedded into systems employees already rely on, making adoption frictionless and largely user-driven.

The public launch of ChatGPT in late 2022, followed by its rapid growth to 100 million users within two months, is often cited as a turning point. It marked the moment when generative AI transitioned from a specialist capability to general-purpose consumer infrastructure. Enterprise adoption followed quickly, often without the deliberate planning, procurement, or any of the typical controls associated with traditional technology rollouts.

The security implications of this shift became visible just as fast, with widely reported incidents in 2023 demonstrating just how easily confidential information could be exposed through AI prompts and file uploads. For example, Samsung, restricted employee use of generative AI tools after internal data was leaked via prompts. Incidents like this helped accelerate regulatory scrutiny in parallel, including Italy's data protection authority, which ordered a temporary halt to all processing of personal data through ChatGPT due to concerns over transparency and lawful basis.

Since that initial surge, consumer and enterprise AI adoption has continued to expand across industries and functions without slowing down, reshaping expectations around productivity and speed. At the same time, it has altered how and where sensitive data is handled, now often being processed outside of the visibility of traditional security controls, and supercharged attacker playbooks, enabling new types of attacks that scale.

Against this backdrop, CultureAI conducted an industry survey in late 2025 to assess how prepared organisations are to adopt AI securely in practice, not just in principle.

The research was designed to answer three core questions:

- **How widely is AI already being used across organisations, and where?**
- **How do leaders perceive visibility, control, and risk today?**
- **How mature are governance, enforcement, and detection capabilities in reality?**

CultureAI commissioned an independent research study of 300 senior technology, security, and risk professionals across North America and Europe. Respondents included CISOs, CIOs, CTOs, Data Protection Officers, and senior IT and security leaders from sectors including finance, healthcare, technology, legal, and professional services.

The survey focused on organisational AI usage, governance structures, risk perception, enforcement capability, and regulatory readiness.

Survey Respondent Demographics

Age Groups

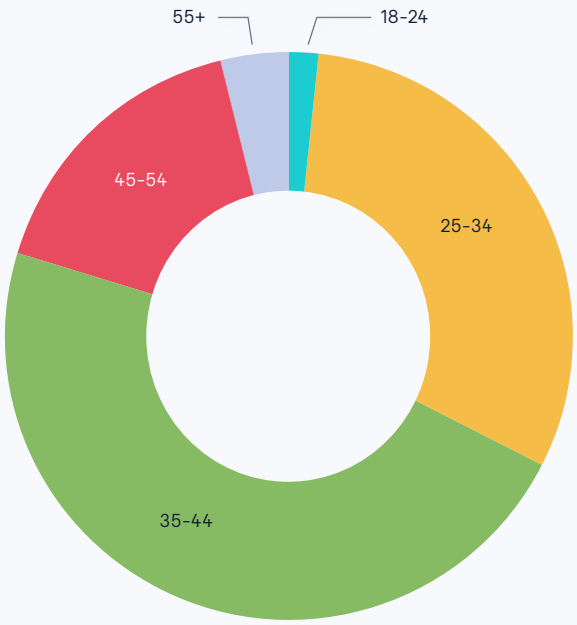

Company Size

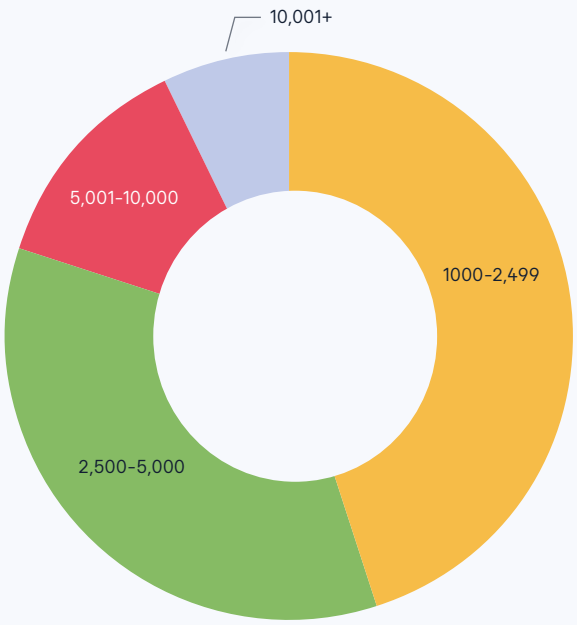

Role

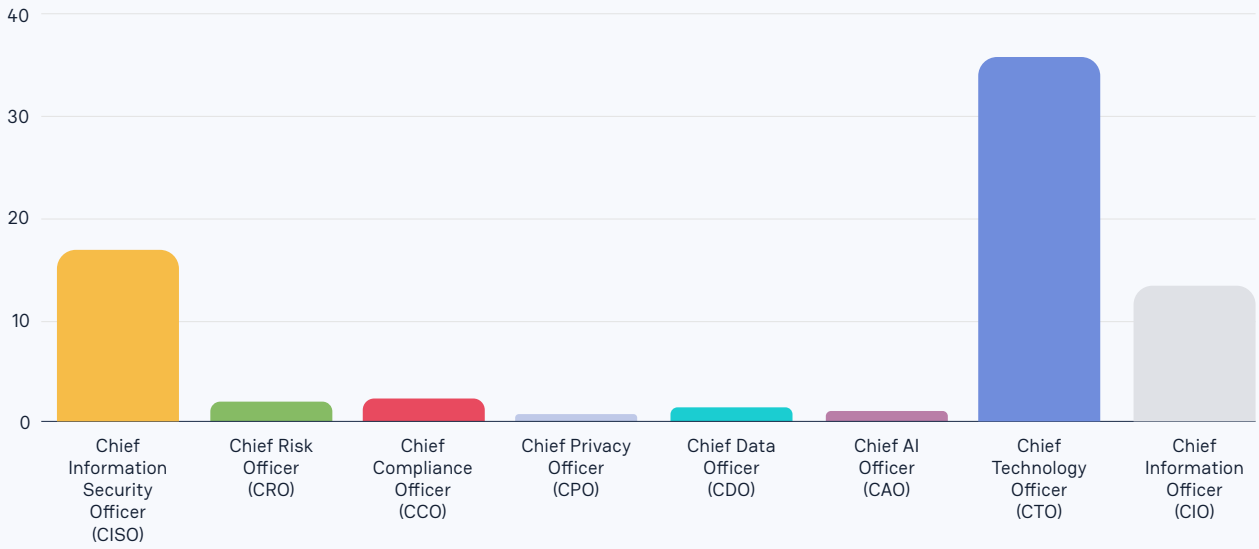

Country

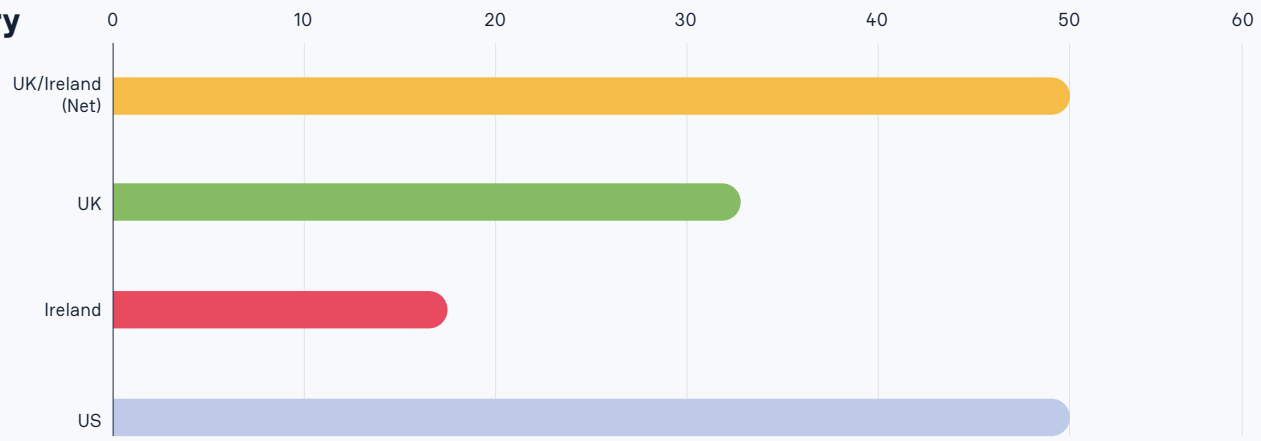

Key Findings at a Glance

This research identifies four core findings that explain why AI risk is growing faster than enterprise controls. Together, they describe how AI is adopted, why control breaks down, why AI risk is underestimated, and what must change operationally.

1 AI Adoption Is Already Widespread and Decentralised

AI is no longer confined to experimentation or specialist teams. It is embedded in everyday work and concentrated in business-critical functions such as engineering, revenue operations, and customer-facing teams.

Adoption is increasingly user-driven and SaaS-based, with AI capabilities embedded inside tools that are already approved. This fundamentally changes where and how risk is created.

2 Perceived Control Does Not Match Operational Reality

Most organisations express strong confidence in their visibility and governance posture, with formal frameworks, policies, and oversight committees now being common.

However, unauthorised AI usage, limited detection, and inconsistent enforcement capabilities remain widespread, creating an illusion of control: governance exists, but behaviour frequently escapes it.

3 AI Risk Is Recognised, but Rarely Escalated

Leaders consistently identify high-impact concerns such as data leakage, compliance exposure, credential compromise, and intellectual property loss.

Despite this, AI risk is still often rated as moderate, reflecting a measurement issue driven by the fact that many AI-related failures resemble ordinary work and do not trigger traditional security signals.

4 The Primary Gap Is Operational Control at the Point of Use

Most organisations have responded to AI adoption with policies, training, and governance structures. These establish intent, but they do not reliably influence behaviour where, or when, AI risk is actually created.

Closing this gap requires governance models that operate in real time and at the point of use, rather than relying solely on tool approvals or policy compliance.

How This Paper Is Structured

The sections that follow expand on these findings in sequence:

- How AI adoption has become widespread and decentralised
- Why traditional governance models create an illusion of control
- How AI risk manifests in practice, and why it is underestimated
- What organisations must do to move from intent to enforceable control

This structure is designed to move from evidence to failure modes, to actionable solutions for CISOs and governance leaders.

AI Adoption is Widespread and Decentralised

AI adoption is no longer experimental. The survey data showed that AI is already embedded in routine workflows across organisations and is expected to expand further, with the heaviest usage sitting within data-rich functions at the centre of core business operations. This combination creates a double-edged reality where significant productivity gains are obtained alongside an increased exposure to significant security and compliance risk.

Core Stat Snapshot

AI is used widely across teams:

67% of organisations report AI is widely used across teams.

27% report use in specific functions.

AI usage is expected to increase:

91% expect AI usage to grow, with

41% anticipating a significant increase.

AI use is focused on core functions:

72% Data Analysis and RevOps;

59% Software Development and Engineering;

43% Customer Support.

Which areas of your business, if any, are using AI tools most actively?

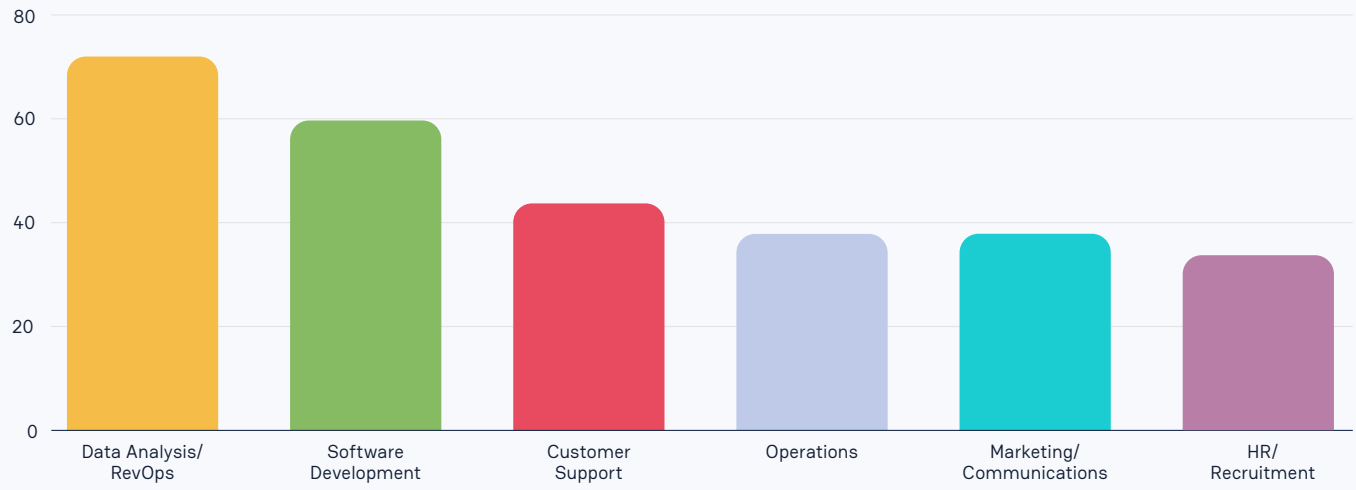

AI's Rapid Normalisation Across the Enterprise

The clearest signal of maturity is breadth. With **67%** of respondents confirming that AI is widely used across teams, adoption has moved beyond innovation programmes and pilots into day-to-day workflows. For security and compliance teams, this matters because data control models designed for "experimental tools used by a few teams" do not scale when adoption becomes ubiquitous.

Momentum compounds the challenge. **Over 90%** of respondents expect AI usage to increase, confirming that AI is not just a trend but an accelerating fundamental shift. Without controls that operate at the same speed as adoption, the gap between usage and governance will continue to widen.

A defining difference between AI and the emergence of earlier enterprise technologies is the ease of informal use. In practice, a user can move from curiosity to processing very real business data in minutes, without procurement, IT involvement, or a formal audit trail beyond browser history. For 2026 and beyond, AI usage must be treated as normal enterprise behaviour, not as exceptional or experimental activity.

Decentralised and User-Driven AI Adoption

AI adoption differs fundamentally from how we have traditionally performed enterprise rollouts. Most consumer AI tools are delivered as SaaS solutions and are accessible without central approval or even account creation. Users discover them through peers, social media, embedded assistants, vendor integrations, and third-party services rather than through managed IT programmes.

This decentralisation creates three predictable governance challenges:

1. **Tool proliferation outpaces inventory**

   Many AI tools bypass procurement and asset management entirely, making comprehensive inventories difficult to maintain.

2. **Fragmented accounts and access paths**

   Users interact with AI through corporate accounts, personal accounts, plugins, APIs, and integrated assistants, fragmenting control and auditability.

3. **Data handling decisions shift to individuals**

   Risk is often created by a single copy-paste or file upload, made by a user without full context on data sensitivity or regulatory obligations, and without mediation by established data protection systems.

Culture and incentives matter. When AI is positioned as a productivity multiplier, adoption is rewarded; however, when governance is perceived as friction, users will go far and wide to work around it. As such, effective security strategies must enable AI use while reducing risk, rather than attempting to suppress adoption altogether.

How AI is Used in **Day-to-Day Business Operations**

Enterprise AI usage is heavily concentrated in functions that routinely handle sensitive, regulated, or commercially significant data. Data analysis and RevOps teams work with pricing, pipelines, and forecasts, engineering teams handle source code, architecture, and security logic, and customer support teams process personal data, contracts, and customer communications.

Across these functions, AI is typically used in a small number of practical workflows:

- **Writing emails, proposals, policies, and customer response**
- **Summarising meetings, tickets, research outputs, and long documents**
- **Analysing spreadsheets, trends, segmentation, and anomalies**
- **Generating, refactoring, and troubleshooting code**
- **Supporting security testing, review, and development workflows**
- **Drafting templates, classifications, and knowledge base content**

The risks created here are not abstract, as the same actions that deliver productivity gains often involve transferring rich business context into third-party systems via prompts or file uploads. Without usage control, organisations often discover risky behaviour only after an incident, a complaint, or an audit request.

There is always risk in technology adoption; however, what distinguishes AI is the impressive speed at which it has been adopted relative to the creation of governance and security controls. In many cases, organisations and users alike have rushed adoption simply to keep pace with competition, assuming tools are “secure enough” or not considering security implications at all. This gap has already resulted in many data leaks, account compromises, and reputational damage with real financial impact.

The Illusion of Control over AI

Survey responses revealed that most organisations believe they have strong oversight of AI usage, yet most have also identified unapproved shadow AI use. At the same time, AI risk is commonly rated as moderate or minor, even though significant data leakage and compliance dominate concerns.

Together, these signals form what can be described as an "illusion of control". Governance structures exist, and confidence is high, but behaviour frequently has no oversight, and risk ratings do not reflect the true underlying exposure.

Core Stat Snapshot

Claim full visibility into AI usage:

72% report full visibility;

28% partial or no visibility

Top concerns:

56% compliance and privacy violations;

52% data leakage via prompts and uploads.

AI risk rated as moderate or minor:

76% rate AI risk as moderate or lower

Detected unauthorised AI use (shadow AI):

65% report detection, either occasionally or frequently

How much visibility, if any, do you have into which AI tools employees are using?

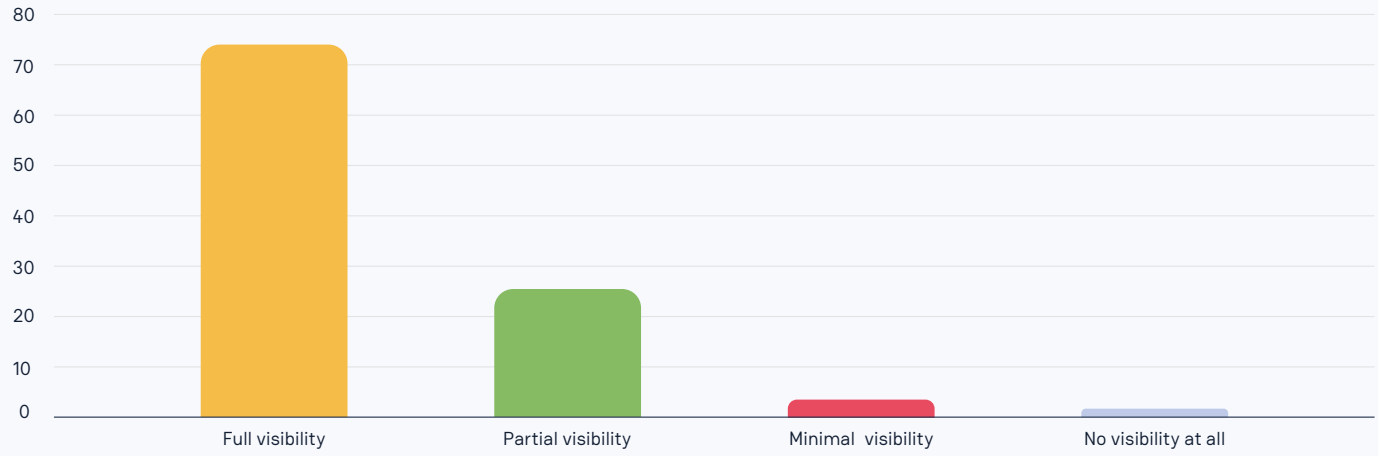

Have you identified any instances of "shadow AI" (use of unapproved AI tools)

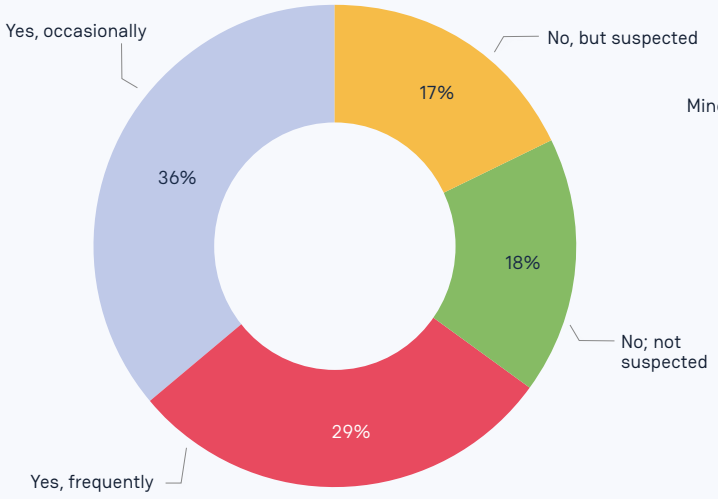

Overall, how would you rate AI as a security risk to your organisation today and why?

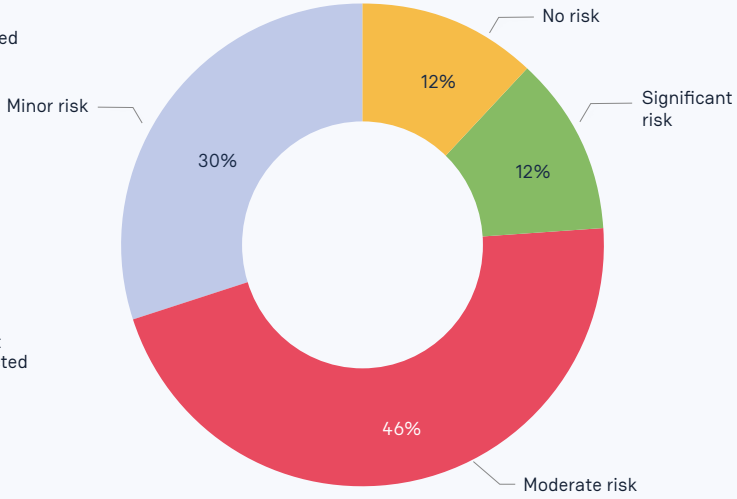

Perception vs Reality of AI Usage Control

If leaders believe they have strong visibility into AI usage, it is reasonable to ask why shadow AI appears so frequently in practice.

Perceived visibility is often built on familiar governance signals: Acceptable use policies exist; approved tools are identified; network traffic is monitored; logs are available for known platforms. Collectively, these controls create confidence that AI activity must be understood.

However, in practice, AI usage patterns consistently undermine these assumptions. Generative AI can be accessed directly in a browser, on unmanaged or personal devices, and via mobile applications. Furthermore, new AI features are increasingly embedded inside tools that are already approved, quietly expanding risk without triggering procurement or review.

The most sensitive actions also leave the smallest traces. A user can copy-paste confidential information or upload a document into an AI prompt without generating the file transfer events or access changes that security teams expect to monitor. The vast number of plugins and extensions also complicates detection, introducing indirect paths for AI interaction that are extremely difficult to inventory or audit.

The widespread presence of shadow AI is therefore not a failure of intent. It simply reflects the deeper mismatch between traditional governance models and the operational reality of AI use. Controls designed for managed applications and known data flows struggle when risk is created in seconds, through free-text interactions, at the discretion of individual users across thousands of tools, with new tools emerging every day.

Open survey responses reinforced this gap, with respondents consistently asking not whether AI is being used, but what data is being shared, by whom, and outside which approved channels:

"Exactly what data is being shared with unapproved generative AI tools and by whom across the organization."
(US, Q20)

"How can we gain full visibility into employee use of AI tools and sensitive data exposure?"
(US, Q20)

"What percentage of employee AI tool usage is currently happening outside of approved/visible channels, and what types of data are most frequently exposed?"
(US, Q20)

Recognised but **Underestimated AI Risks**

This visibility gap feeds directly into how AI risk is assessed. Nearly half of the respondents rated AI risk as moderate (**46%**), and almost a third rated it as minor (**30%**). At the same time, their stated concerns focused overwhelmingly on concrete outcomes:

Compliance / privacy violations:

Data leakage through prompts / uploads:

Account compromise / credential theft:

Confidential code / IP exposure:

Misinformation / brand impersonation:

Manipulated / poisoned models:

Which AI related risks, if any, most concern you today?

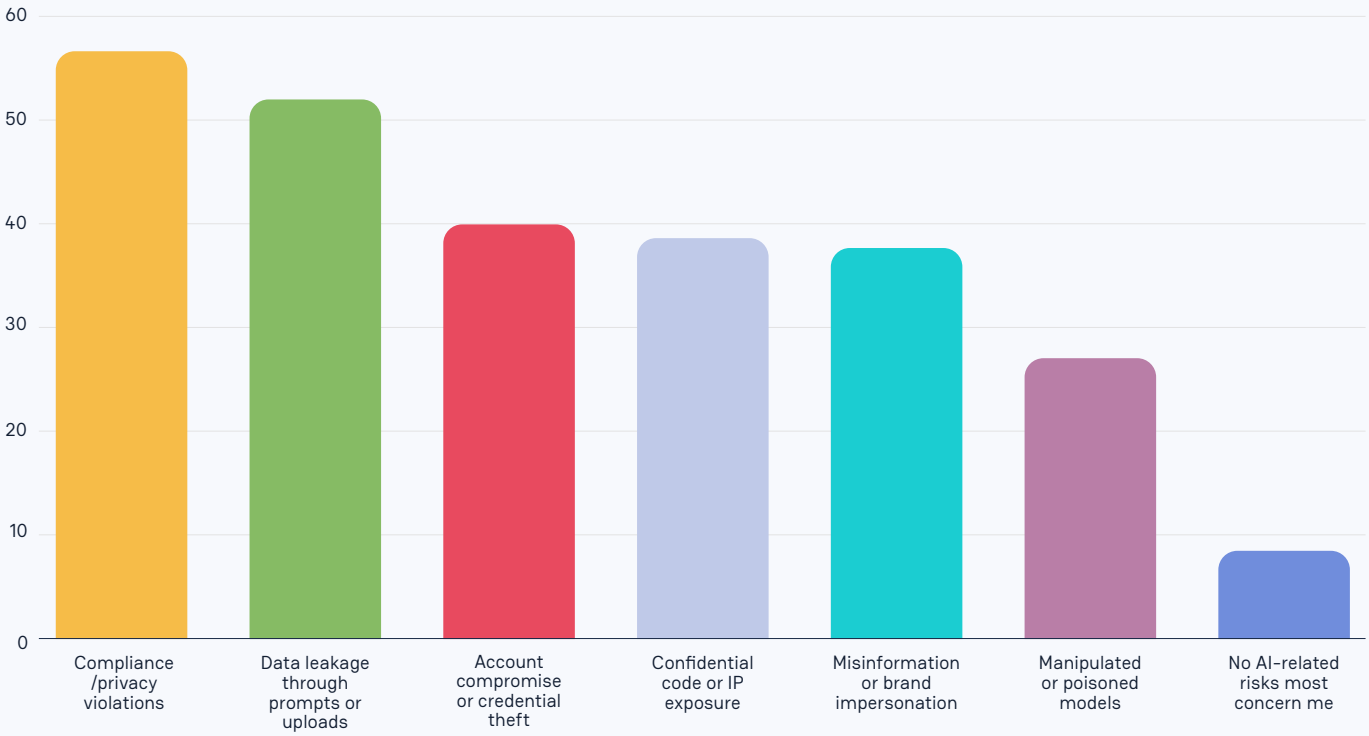

This apparent contradiction reveals that leaders are not dismissing AI risk, but they are struggling to accurately quantify it in an environment where damage often occurs without an obvious breach, alert, or outage.

A single paste of sensitive data into an external model might never trigger an incident response process; however, it may still create regulatory exposure, contractual breach, or a very real data compromise down the line. In the absence of detection and attribution capabilities, risk is difficult to measure, leading to conservative ratings that underestimate the reality.

Until organisations can consistently see, contextualise, and control AI usage in the moment, confidence will continue to outpace reality.

The Primary Security Failure in AI Adoption

The primary security failure in most organisations is not the decision to adopt AI (in many cases, adoption has been rational and beneficial). The failure is allowing AI usage to expand faster than the organisation's ability to define, detect, and enforce meaningful data boundaries around it.

Most organisations began sensibly with awareness programmes launched, acceptable use policies written, training rolled out, and governance committees formed. These steps created shared understanding and legitimised AI as an enterprise concern. However, problems arise when these early measures are assumed to be sufficient, even as evidence of shadow AI usage and data exposure continues to surface.

From a security design perspective, the gap is rarely about intent and almost always about a lack of capability:

- **No reliable inventory of AI tools or embedded AI features**
- **No consistent data classification or handling rules applied at the moment of use**
- **No enforcement layer to prevent or reduce unsafe prompts or uploads**
- **No measurement linking AI usage to data sensitivity or regulatory obligations**

Governance exists, but it is disconnected from day-to-day behaviour. Policies describe what should happen, but systems are not in place to ensure that it does.

AI Governance Exists, but is not Operational

Most organisations have implemented, or are developing, formal AI governance frameworks. Oversight committees are common and confidence in compliance is high. However, enforcement and detection capabilities remain inconsistent.

Core Stat Snapshot

Formal AI governance framework implemented:

Framework enforcement inactive:

Framework in development:

AI or risk committee exists:

Governance not actively enforced:

Lack of dedicated AI detection:

Training is the most common control:

Acceptable use policy is the most common control:

The Rise of Frameworks and Committees

For most organisations, the first response to the rapid spread of generative AI has been structural. Faced with a fast-moving technology that cuts across functions and regulatory domains, frameworks and committees offer an immediate way to impose order.

The survey data suggested that this structural response is now widespread, with nearly two-thirds of organisations reporting that they have already implemented a formal AI governance framework, while a further third are actively developing one. Similarly, two-thirds say they have established an AI or risk committee with explicit oversight responsibilities.

Do you currently have an AI committee?

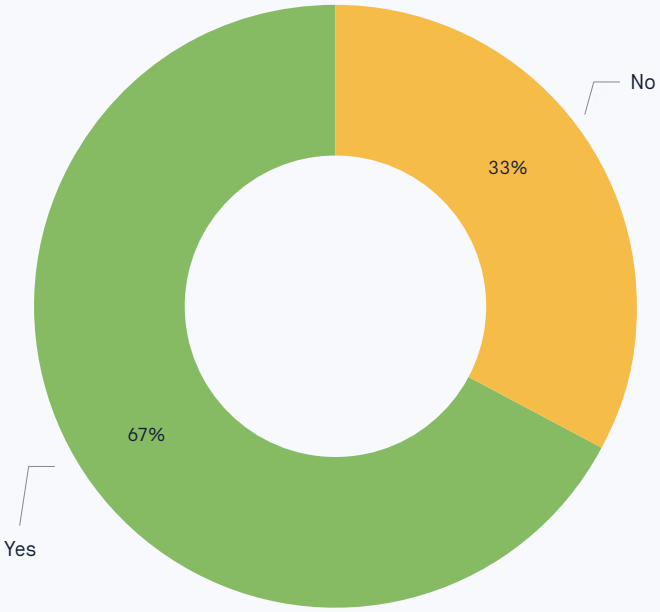

This progress matters, as it creates the conditions for consistent decision-making and cross-functional alignment. However, as previous sections have shown, the presence of frameworks and committees does not automatically translate into effective control. The next challenge for many organisations is ensuring that governance decisions are meaningfully enforced in day-to-day use.

Fragmented **Ownership and Accountability**

Even where AI governance committees exist, accountability is often fragmented. This is largely because AI spans multiple domains, each with legitimate but different priorities:

- Security wants risk reduction and visibility
- IT wants enablement and operational simplicity
- Legal and privacy want defensible compliance and minimised exposure
- Business leaders want speed, productivity, and competitive advantage

Without clear ownership, governance slows and enforcement becomes inconsistent. Policies may exist but are unevenly applied and tool approvals become a proxy for safety, while data handling rules remain unclear.

Effective governance does not require absolute centralisation, but it does need to be clear. In practice, this may mean assigning a single accountable owner for the AI usage policy, even if the broader AI strategy remains more widely distributed. It also requires metrics that reflect real behaviour rather than the mere existence of policies. Without these foundations, AI governance will remain well-intentioned but operationally fragile and inconsistent.

The Illusion of **Regulatory Readiness**

The survey revealed a second, closely related illusion around regulatory readiness. On the surface, confidence is high:

61% say they are very confident in their organisation's compliance posture.

62% report that they have implemented a formal AI governance framework.

33% are actively developing one.

Do you have a formal AI governance framework or policy?

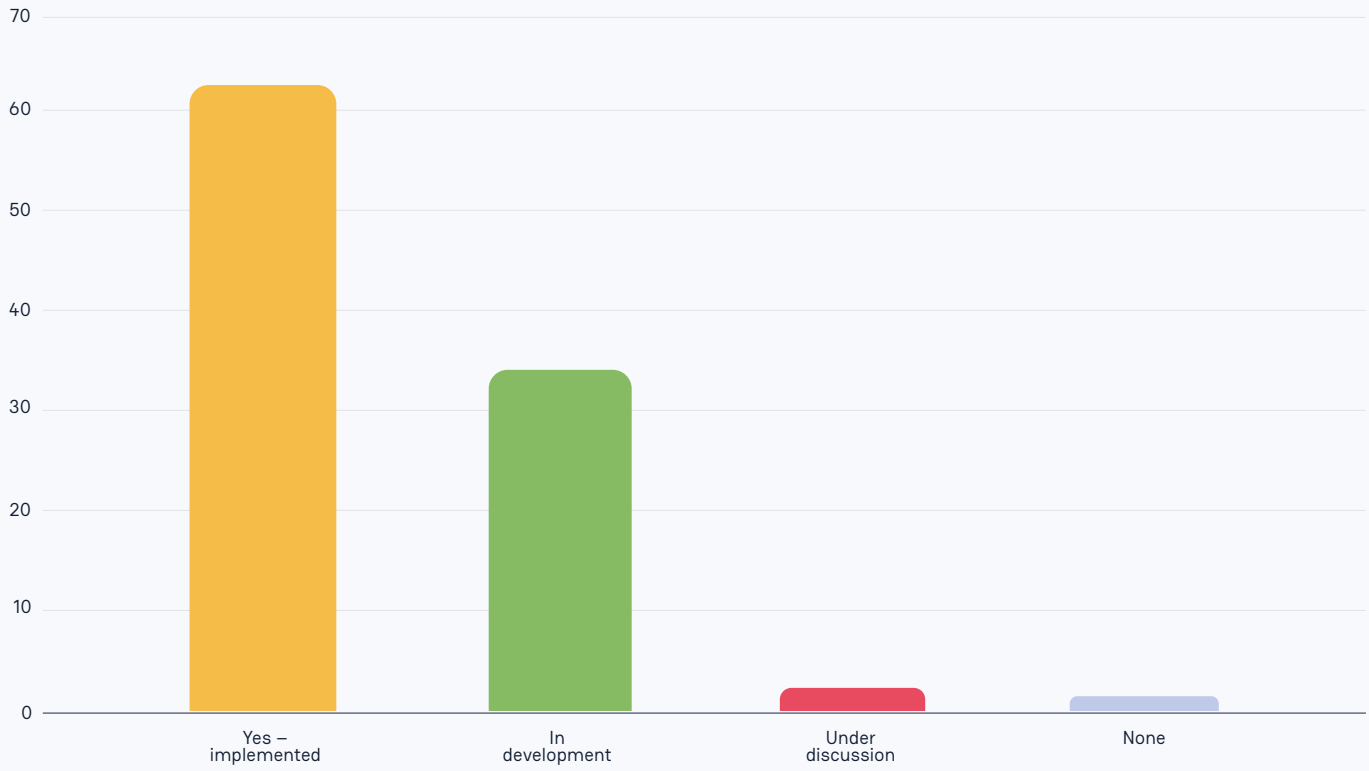

However, this confidence sits alongside clear operational gaps, with **20%** of respondents acknowledging that their policies are not actively enforced, and more than a third lacking dedicated AI detection capabilities altogether. This means that non-technical controls, such as policies and training, are far more common than mechanisms that detect or control AI usage in real time.

This will be a significant issue from 2026, and beyond, as regulatory pressure is no longer hypothetical and has clear milestones that organisations are actively tracking. For example, the EU AI Act, which entered into force in 2024, has a general date of application of August 2026, with full effectiveness expected by 2027. This shift moves compliance expectations away from intent and towards demonstrable, auditable control.

Without detection, enforcement, and measurement through AI usage control, perceived readiness risks collapse the moment it is tested.

What is the true risk from AI?

When AI risks are discussed, topics such as model poisoning or adversarial manipulation often dominate. These risks do matter, but they are not what most organisations will face in the immediate future.

The most significant enterprise risks are practical and immediate. Complications such as AI lowering the skill barrier for attackers and accelerating existing attack patterns, while simultaneously encouraging users to submit rich, sensitive context into third-party systems.

AI Is Rapidly Improving Attacker Playbooks

AI has not replaced common cybercrime tactics, but it has improved their volume, efficiency, and credibility.

Social engineering at scale

CrowdStrike reported a **442%** increase in voice phishing between the first and second halves of 2024, driven in part by AI-enabled impersonation.

Faster conversion from access to impact

CrowdStrike's "breakout time" reporting shows that in 2025, attackers were moving laterally in an average of 18 minutes, down from 48 minutes in 2024. AI-driven automation accelerates reconnaissance, scripting, and decision-making, reducing defenders' response windows.

AI Is Creating a New Attack Surface

The software supply chain problem has also worsened, with AI adding an entire new supply chain: models, datasets, and AI components.

Malicious packages are surging

The number of malicious open-source packages has increased 156% year-over-year as AI adoption accelerates dependency sprawl through wrappers, agent frameworks, and orchestration tooling pulled directly from public repositories.

Models are becoming a delivery mechanism

JFrog reported that more than 1 million new models were added to Hugging Face in 2024 alone, with many unvetted artefacts now accessible that may contain backdoors, unsafe code, or compromised dependencies.

The Biggest Enterprise Risk Is **Data Loss Through Normal Use**

AI tools extract value from detail, which incentivises users to paste proprietary, regulated, or commercially sensitive information into systems outside organisational control.

Sensitive data is already flowing

Menlo Security reported that **55%** of GenAI-related DLP events involved users entering personally identifiable information.

Real incidents show how quietly this happens

In 2023 Samsung engineers submitted confidential material to ChatGPT as part of routine work. Leakage was unintentional and productivity-driven, not malicious.

Breaches of AI services turn prompts into a data trove

In early 2025, Wiz reported an exposed DeepSeek database containing extensive logs, chat history and operational data. Separately, OmniGPT was reported to have suffered a breach involving the exposure of user data and 34 million chat logs. Even when sharing begins as an internal decision, once data is shared, protection depends entirely on third-party security posture.

AI Accounts and Chat Histories as High-Value Targets

AI tools are quickly becoming repositories of organisational memory: credentials pasted during debugging, customer details used in drafting, commercial plans in strategy work, and code shared for refactoring. That makes AI accounts valuable targets.

Group-IB reported that compromised ChatGPT credentials were being traded on underground markets. The risk is not only account takeover, but what attackers gain immediately: chat histories containing sensitive context, plus an interface that helps them summarise, search, and operationalise stolen information at speed.

Shadow AI Turns **Governance Gaps Into Exploits**

A CybSafe and National Cybersecurity Alliance study found that 38% of workers using AI for work admit to sharing sensitive information without their employer's knowledge.

Application security data reinforces this, with Harness's State of AI-Native Application Security 2025 reporting that 62% of organisations have no visibility into where LLMs are used, alongside real-world prompt injection and vulnerable LLM integrations.

Framing the True" Risk: **Probability, Impact, and Detectability**

The enterprise risk from AI is defined by three characteristics:

High probability

AI use is already widespread, and user-led adoption routinely bypasses central control.

High impact

Leaked context often includes customer data, credentials, IP, and strategic plans.

Low detectability

Many failures resemble ordinary usage until harm appears later.

This framing explains why AI risk is underestimated. Harm rarely presents as a dramatic breach. It emerges quietly, followed by delayed regulatory, contractual, or competitive consequences.

The Need for AI Usage Control

Although many organisations have started to implement formal AI oversight, and leadership attention is rising, adoption of controls that operate at the point of use is still lacking. That gap matters because AI risk is user-driven, data-centred, and rapidly evolving. Policies describe what should happen, but usage control determines what actually happens, with evidence.

Core Stat **Snapshot**

AI risk rating:

- 46% moderate
- 30% minor
- 12% no risk

Leadership stance on priority:

- 21% no addressed
- 10% low
- 2% not on the radar

Governance transition state:

- 62% implemented
- 33% developing

Stricter contols planned:

- 68% show demand for operational solutions.

Is AI security a priority for your organisation in 2026?

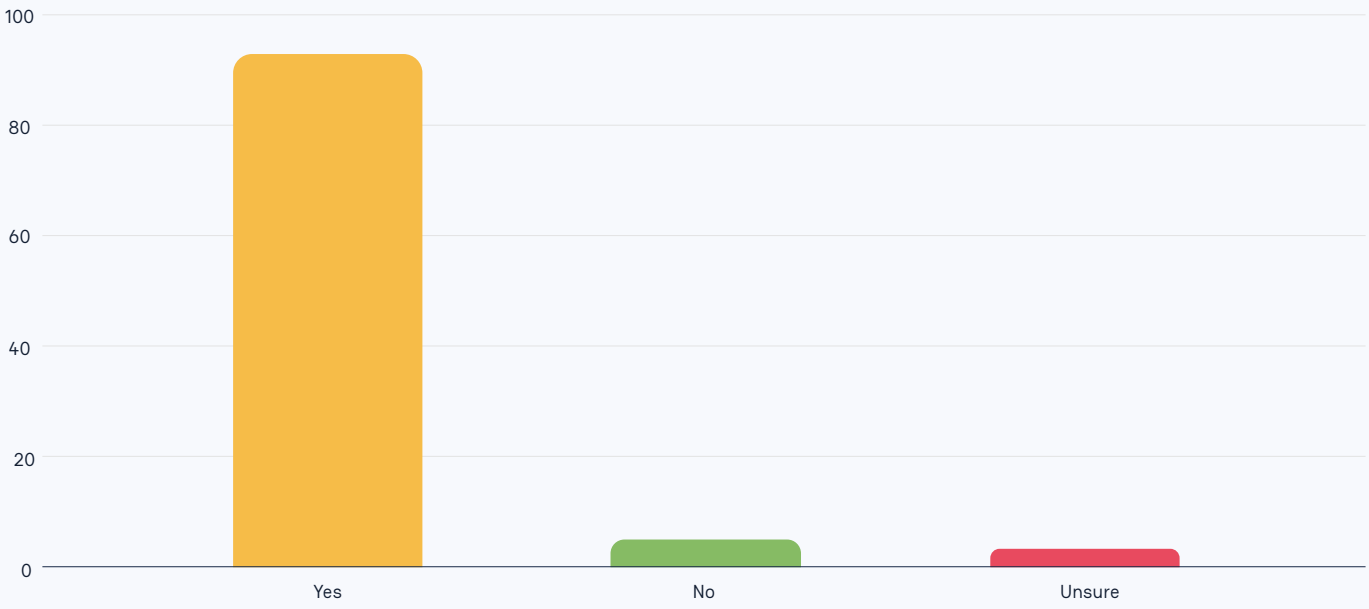

Security Maturity Spectrum

AI governance maturity is best understood as a spectrum rather than a binary state. Based on the survey responses and observed practices, most organisations currently sit between acknowledgement and governance, with growing momentum towards enforcement.

Maturity level	What it looks like	Typical controls	Typical failure mode
Unseen	AI use is informal and not measured	Ad hoc guidance	Unknown exposure
Acknowledged	Policy exists; training begins	Acceptable use policy, awareness	Shadow AI expands
Governed	Committee, framework, approval flow	Tool approval, risk reviews	Slow decisions, exceptions grow
Enforced	Usage is detected and controlled	Monitoring, data controls, response	Partial coverage, gaps in edge cases
Optimised	Governance supports productivity safely	Continuous measurement, coaching, automation	Risk accepted knowingly and managed

Foundation for Safe Enablement

For most organisations, the objective is not to restrict AI, but to enable its use at scale without introducing unnecessary exposure. This aligns with the prevailing “moderate risk” posture, where AI is seen as valuable but not yet fully controlled.

Safe enablement depends on the ability to answer four operational questions that governance programmes often assume are already resolved:

1. Which AI tools and features are in use, including embedded assistants?
2. What data types are being shared (customer data, source code, financials, HR information)?
3. Which users and teams are using AI, and for which workflows?
4. What guardrails apply at the moment of use (warn, redact, block, approve, log)?

Frameworks and committees can define these rules on paper; however, without usage control technologies, these rules cannot be applied consistently in real time. In practice, this leads to a dangerous misunderstanding that approval to use a tool can be treated as approval to use any data within it, as captured by some of the respondents clearly, when asked about existing misconceptions:

> *"Permission to use is a blanket clearance to use any data in that tool."*
>
> *(UK, Q18)*

> *"That AI tools are safe by default and don't pose data leakage or compliance risks."*
>
> *(US, Q18)*

Safe adoption requires the opposite assumption. Tools may be approved, but data use must still be governed dynamically, based on sensitivity, context, and workflow.

AI Usage Control as **Core Security Infrastructure**

AI usage control is the operational layer that turns governance intent into execution. It sits between policy and behaviour and ensures that decisions made by committees and leadership are enforced where risk is actually created.

A credible AI usage control capability should consist of (at least) of the following interconnected layers:

Discovery and inventory

1

Identify AI tools in use (approved and unapproved), detect AI features embedded in existing SaaS, and map usage and data residency by team, role, region, and device context. Visibility gaps are the engine of shadow AI.

Data-aware policy at the point of use

2

Move beyond “allowed” versus “blocked”. Apply rules based on sensitivity and context, distinguish low-risk prompts from high-risk uploads, and use user-supportive interventions such as warnings and just-in-time guidance.

Monitoring, alerting, and evidence

3

Leaders and governance bodies need proof that controls are operational: what was allowed, what was blocked, what exceptions occurred, and which risk patterns are emerging. Usage control should produce audit-ready logs and reporting that translate technical events into control coverage and risk language.

Response and continuous improvement

4

When incidents or near misses occur, answer quickly: what data was involved, which tool, which user or workflow, and which control failed or was missing. Feed those lessons back into training, policy refinement, and targeted coaching to reduce repeated unsafe behaviour.

This capability should not be confused with traditional DLP, as classic DLP solutions are optimised for known channels and file-based transfers. AI risk predominantly occurs through natural language interactions and routine uploads that look like ordinary work. As such, AI usage control focuses on the human-to-model interface where the majority of high-volume, low-friction exposure now occurs.

Usage control solutions allow the current assumptions to be replaced with real measurements and enable safe AI adoption through demonstrable safeguards rather than trust in defaults.

Strategic Implications for Security Leaders

For security leaders, the differentiator in 2026 will be demonstrable control. Organisations that respond effectively will not be those with the most comprehensive policy documents, but those that can answer simple questions with evidence.

1

Shift from tool governance to data governance in AI contexts

Approving an AI tool does not equate to approving all uses of that tool. Treat prompts, uploads, and pasted content as data movement and govern them accordingly.

2

Measure reality, not intention

Treat claims of “full visibility” as hypotheses to validate. Shadow AI detection is not a failure; it is insight into where productivity demand exceeds official enablement.

3

Build an AI control plane that matches how AI is actually used

Risk does not arrive through a single sanctioned platform. It emerges through browsers, plugins, embedded assistants, developer integrations, and personal accounts. Controls must cover these paths.

4

Turn committees into operating mechanisms

Governance groups should produce enforceable artefacts: policies that can be implemented, exception rules that can be executed, and metrics that reflect real behaviour.

5

Align AI controls to regulatory evidence requirements

Confidence is not evidence. Expect regulators, customers, and auditors to ask how AI usage is monitored and controlled in practice.

6

Optimise for safe productivity, not maximum restriction

If safe usage is harder than unsafe usage, shadow AI will persist. Proportional guardrails and just-in-time guidance reduce risk without undermining adoption.

Conclusion

Generative AI is no longer an emerging technology within the enterprise. It is embedded, decentralised, and heavily used across core business functions, often outside traditional IT and security controls. This shift has fundamentally changed how risk is created, and AI exposure now emerges through everyday actions such as prompts, uploads, integrations, and chat histories, driven by individual users and not centrally managed systems.

This research showed that most organisations now recognise AI as a governance concern in principle. Frameworks exist, committees have been formed, and confidence in compliance is high. However, these signals of readiness sit alongside widespread shadow AI usage, limited real-time visibility, and inconsistent enforcement (if any at all). The result is a gap between perceived control and operational reality, with AI risk being acknowledged but often rated as moderate because it is difficult to see, measure, and attribute until the damage is done.

The most significant AI risks in 2026 are not speculative or theoretical; they are practical, high-probability risks tied to data loss, account compromise, supply chain exposure, and regulatory failure, arising from normal and well-intentioned use. When sensitive information and context is shared quietly at scale, traditional security and governance models struggle to detect or prevent harm in time.

The core failure is not AI adoption itself but allowing usage to scale faster than the organisation's ability to detect and enforce security boundaries at the point of use. Policies and training set intent, but they do not control behaviour. As regulatory expectations shift towards demonstrable, auditable enforcement, confidence without evidence will no longer be sufficient.

The path forward is real-time operational AI usage control. Organisations must move beyond tool approval and paper governance towards real-time visibility, data-aware guardrails, consistent enforcement, and measurable outcomes. Once controls are implemented where risks are actually being created, AI can be adopted at scale without introducing avoidable or unknown exposures.

The question is no longer whether AI will be used, but whether **organisations can govern and secure its use in practice, not just in theory.**

Appendix 1 – AI Myth-busting

Many organisations are still operating on assumptions that no longer reflect how AI is actually being used across businesses. The following myths highlight the most common gaps between perception and reality, and why they matter for risk, control and compliance:

Myth 1:

"AI adoption is slowing down, so we have time to prepare."

Reality:

AI adoption is accelerating, not plateauing. Over **90%** of organisations expect usage to increase in the next 12 months, with **40%** expecting significant growth.

Why this matters:

Risk scales with usage. Waiting to "get governance right later" means exposure grows faster than controls.

Myth 2:

"AI is still restricted or experimental in most organisations."

Reality:

AI is already mainstream. **67%** of organisations report AI is widely used across teams, while only **7%** have it highly restricted or banned.

Why this matters:

The debate has shifted from whether to allow AI to how to enable it safely. Heavy restrictions often push usage underground, increasing risk.

Myth 3:

"AI is mostly used in low-risk areas."

Reality:

The highest AI usage is in Data Analysis, Software Development and Customer Support, exactly where sensitive data, IP and critical decisions live.

Why this matters:

AI risk is not theoretical. The blast radius includes customer data, proprietary code and strategic insight.

Myth 4:

“We already have good visibility and controls.”

Reality:

While **72%** claim full visibility into AI usage, **65%** have still identified shadow AI. Many tools, personal accounts and embedded AI features remain invisible to traditional controls.

Why this matters:

Approved tool lists do not equal real-world visibility. Blind spots are where incidents happen.

Myth 5:

“Policies, training and confidence are enough.”

Reality:

Compliance and privacy are the top concerns, yet almost everyone feels confident they comply. At the same time, many admit policies are not actively enforced, and detection of sensitive data leakage is only partial.

Why this matters:

Regulators and customers will ask for proof of enforcement, not policy documents or training slides.

Appendix 2 – Survey Stats

Percentages may total more than 100% because some questions allowed respondents to select multiple answers.

All questions - Top Two Responses

Q#	Question Topic	Top Response	Count	%	2nd Top Response	Count	%
Q1	AI Usage Policy	Widely used across teams	201	67%	Allowed for specific use cases	75	25%
Q2	Business Areas Using AI	Data Analysis/RevOps	216	72%	Software Development	176	59%
Q3	Visibility Level	Full visibility	216	72%	Partial visibility	76	25%
Q4	Shadow AI Detection	Yes, occasionally	107	36%	Yes, frequently	88	29%
Q5	Top Risk Concern	Compliance/privacy	167	56%	Data leakage	155	52%
Q6	Risk Level	Moderate risk	139	46%	Minor risk	89	30%
Q7	Leadership Stance	Clear priority	201	67%	Recognised not addressed	62	21%
Q8	Detection Capability	Yes-dedicated	190	63%	Partially-via DLP	91	30%
Q9	Top Control	Acceptable-use policy	183	61%	Employee training	181	60%
Q10	Ownership	IT/Engineering	133	51%	CISO/Security	91	35%
Q11	Governance Status	Yes-implemented	187	64%	In development	99	34%
Q13	Has Committee	Yes	202	67%	No	98	33%
Q14	Regulatory Confidence	Very confident	183	61%	Somewhat confident	115	38%
Q15	Growth Expectation	Increase moderately	150	50%	Increase significantly	122	41%
Q16	Top Planned Action	Implementing controls	203	68%	Updating governance	189	63%
Q17	Biggest Blocker	Lack expertise	77	38%	No single blocker	70	35%
Q19	2026 Priority	Yes	275	92%	No	16	5%

Analysis by Country

Country	Count	AI Widely Used	Full Visibility	Shadow AI	Has Governance	2026 Priority	Top Blocker
US	150	68%	75%	65%	69%	93%	24%
UK	100	63%	75%	63%	60%	88%	25%
Ireland	50	72%	58%	68%	48%	94%	32%

Analysis by Industry Sector

Sector	Count	AI Adoption	Shadow AI	Governance	Detection	Committee	2026 Priority
IT	106	65%	70%	60%	69%	77%	93%
Finance	67	61%	72%	69%	70%	73%	93%
Manufacturing	62	73%	53%	66%	63%	52%	90%
Retail	32	72%	47%	56%	56%	50%	84%
Banking	31	68%	77%	52%	35%	68%	94%
Legal	2	100%	50%	100%	100%	100%	100%

Analysis by Role/Job Position

Role	Count	AI Widely Used	Full Visibility	Clear Priority	Has Controls	Very Confident	Top Risk Concern
Chief Technology Officer (CTO)	108	76%	72%	72%	71%	68%	Data leakage
Chief Information Security Officer (CISO)	51	63%	78%	73%	69%	71%	Data leakage
Chief Information Officer (CIO)	40	68%	70%	73%	55%	60%	Data leakage
Director of Operational Risk Management	15	67%	80%	67%	80%	73%	Data leakage
Head of Information Security	13	62%	38%	54%	38%	46%	Data leakage
Head of Cybersecurity	10	80%	70%	50%	40%	40%	Data leakage
Chief Compliance Officer (CCO)	7	43%	86%	57%	43%	29%	Compliance
Information Security	7	57%	100%	57%	57%	29%	Data leakage
Chief Risk Officer (CRO)	6	67%	67%	67%	50%	67%	Data leakage
AI Governance Lead / AI Policy & Risk Lead	6	83%	83%	50%	67%	83%	Data leakage
Director / VP of Governance	6	50%	67%	50%	67%	17%	Data leakage
Risk & Compliance (GRC)	5	40%	40%	60%	40%	60%	Compliance
Chief Data Officer (CDO)	4	0%	75%	50%	25%	25%	Data leakage
Information Governance Manager	4	25%	50%	50%	50%	25%	Data leakage
Legal Compliance Manager - Data & AI	4	50%	75%	75%	75%	75%	Compliance

CultureAI

CultureAI is an AI Usage Control Platform that enables organisations to adopt AI safely, confidently, and at scale.

It gives security and compliance teams real visibility into how AI is actually used across 10,000+ AI-enabled applications, including sanctioned and unsanctioned tools, shadow AI, and embedded SaaS features.

CultureAI supports security leaders and AI governance teams across highlight regulated industries, who wish to enable AI adoption without compromising compliance or control.

Research by Oliver Simonnet
Lead Cybersecurity Researcher at CultureAI

www.culture.ai