

LEGALFLY

REPORT

The AI Governance Gap

Actions governments and industry can take to ensure responsible adoption of AI

CONTENTS

Part One 02

AI governance can no longer wait 03

AI is now embedded in operations 05

Leadership are enthusiastic about AI's potential 07

AI risk perceptions 09

AI governance maturity is low 12

Part Two 17

Better AI governance must be prioritised 19

More guidance is needed from external sources 20

Operationalising AI governance 23

Six recommendations to enhance AI security & operationalise governance 25

Summary 27

PART ONE

Why AI governance is falling short

1

AI governance can no longer wait

AI is now central to business operations, but most organisations lack the governance needed to manage its risks responsibly.

AI is now embedded in how organisations operate. Over 90% use it, and more than half rely on it several times a week or more. These tools influence decisions, shape outputs, and automate increasingly sensitive tasks, from triaging investigations and reviewing financial exposures to flagging bias risks in hiring and promotions.

Despite widespread use, governance has not kept pace. Only 18 percent of organisations have a fully developed AI policy. Most rely on high-level principles around ethics and privacy, but few define how AI should be selected, deployed or monitored. In some cases, there is no oversight or accountability.

This creates material risk. AI tools are now used to analyse customer data, draft communications and support strategic decisions. Unlike traditional software, there is a risk they may produce outputs that are inaccurate, biased or opaque. Inputting sensitive data into unsecured systems further increases the risk of regulatory, contractual and reputational harm.

Yet there is no stopping AI. Enterprises see its strategic importance to drive customer experience and competitive advantage -17% already see it as essential to operations. Therefore, overly restrictive policies will likely drive AI usage underground and further increase risk.

> If we look at the speed in which the adoption of AI has happened in consumer, it's never been seen before. ChatGPT took about three months to reach its one hundred million users. Think about past technologies, mobile phones, for instance. It took about sixteen years, and the Internet took about seven years to reach the same level of users.

An Tips | Go To Market Lead, Microsoft

LEGALFLY commissioned this research to understand how AI is being governed in practice. Based on responses from 154 general counsel across the UK, France and Germany, and informed by expert interviews, this report outlines where governance is falling short and what can be done to close the gap.

The goal is not to restrict innovation, but to make it safer and more scalable. Some argue that more governance can stifle progress, but in practice, its absence creates confusion, inconsistency, and risk. Existing frameworks are often too vague or detached from how teams actually work. The real need is governance that works in practice: policies that guide behaviour, safeguards built into workflows, and leadership that takes ownership.

From this research, we present a set of clear, practical recommendations to help close the AI governance gap.

ABOUT LEGALFLY

LEGALFLY is the leading agentic AI workspace for in-house legal and compliance teams. Some of the largest enterprises across Europe use LEGALFLY to streamline contract reviews, drafting, and compliance workflows, increasing operational efficiency.

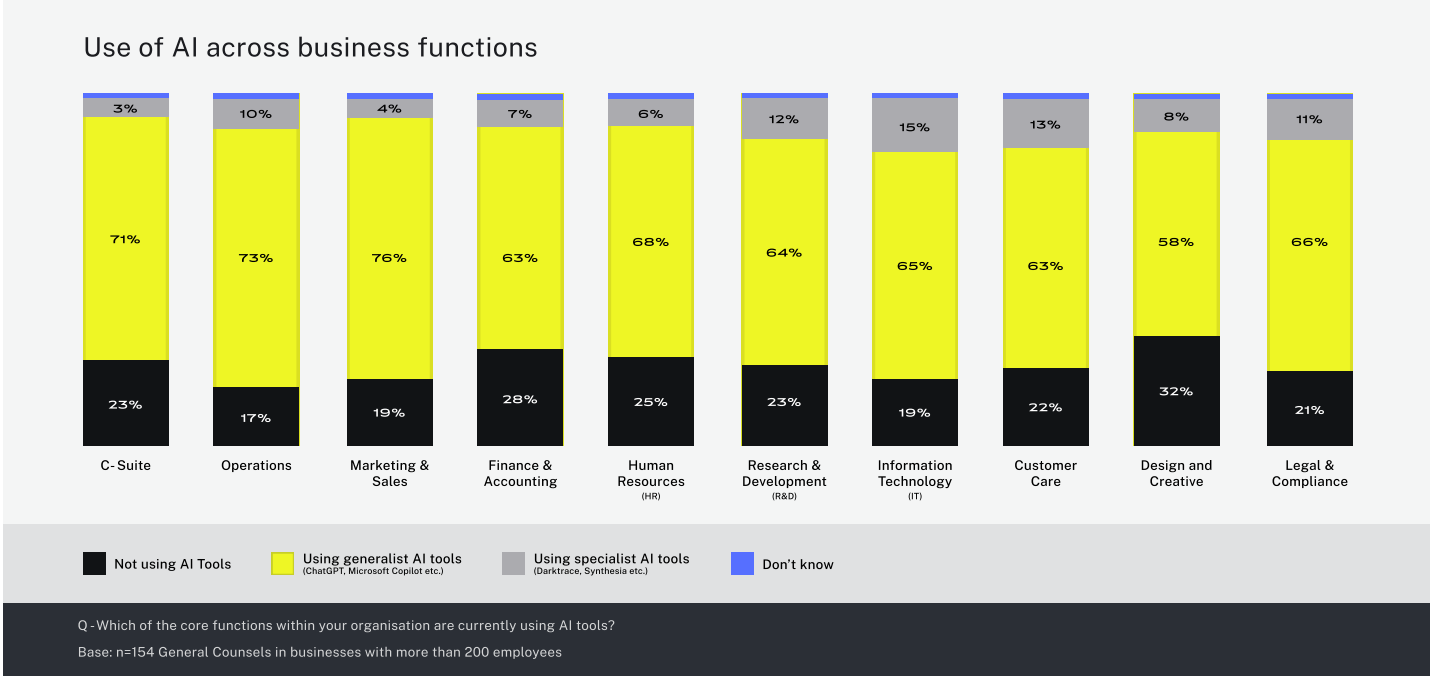
Founded in 2023, LEGALFLY's mission is to empower legal professionals to deliver more value with AI, safely and responsibly. LEGALFLY is the only Legal AI platform to anonymise all sensitive data before processing begins, ensuring that confidential information is protected at all times.

legalfly.com

2

AI is now embedded in operations.

AI is widely used across functions, with daily usage growing. Without consistent oversight, adoption often happens informally or invisibly.



AI is now deeply embedded across many core business functions. Ninety percent of organisations report active use of AI tools in day-to-day operations, spanning both generalist platforms such as Microsoft Copilot and ChatGPT, and specialist systems tailored for legal, compliance, finance and IT use cases.

AI adoption is widespread across functions

AI adoption is highest in operations (83 percent), marketing and sales (80 percent), and IT (80 percent), where automation and data analysis support efficiency and scale. Legal and compliance follow closely at 77 percent, driven by tools for document review, contract analysis and policy management.

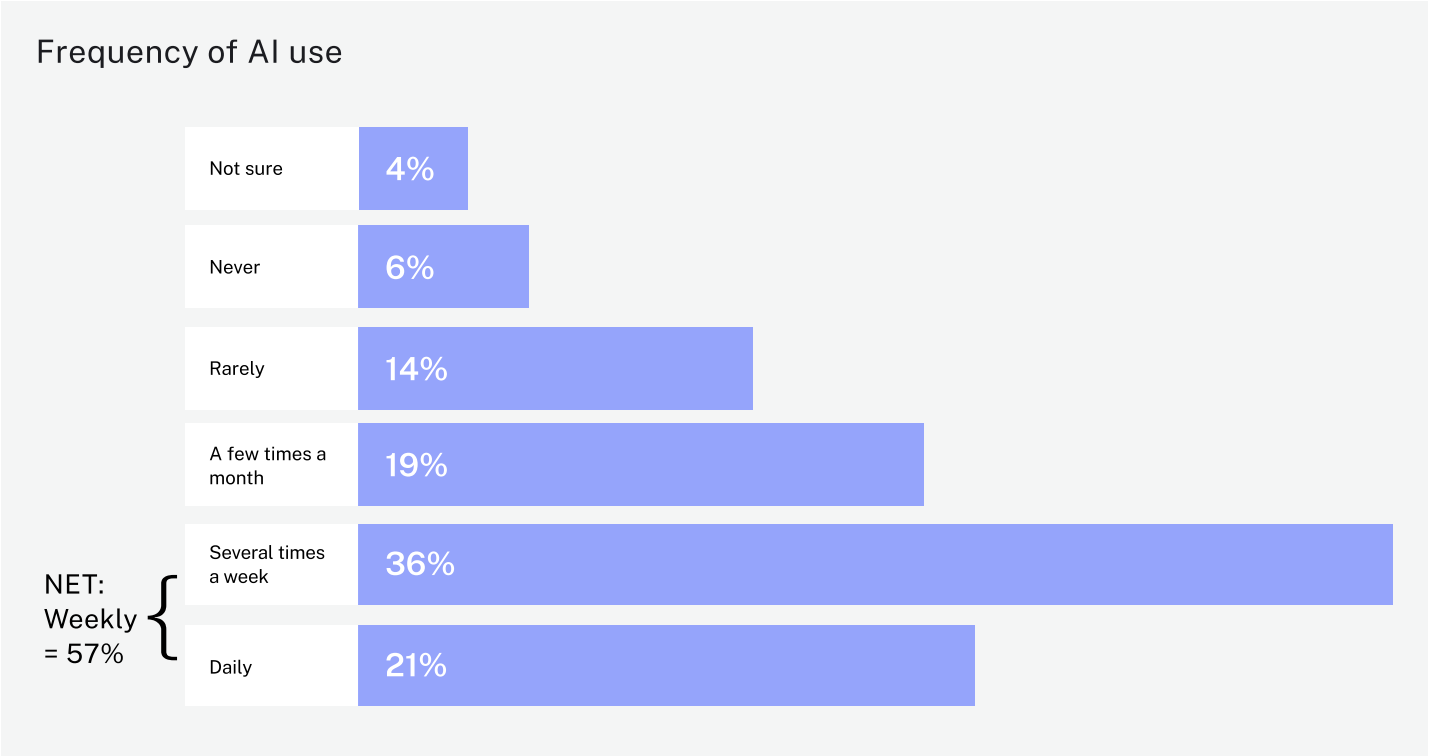
Use is also high in traditionally cautious functions. Seventy-four percent of HR teams and 70 percent of finance teams now use AI, often to summarise data, generate reports or triage requests. In practice, this might mean screening CVs in HR or automating variance analysis in finance.

At the leadership level, 74 percent of C-suite teams use AI to draft communications, explore strategic scenarios or review business summaries.

90%

of organisations use AI in day-to-day operations.

AI use is becoming routine



Q - How frequently do people in your organisation use AI tools as part of your work?

Base: n=154 General Counsels in businesses with more than 200 employees (UK, n=52; FR, n=51; DE, n=51)

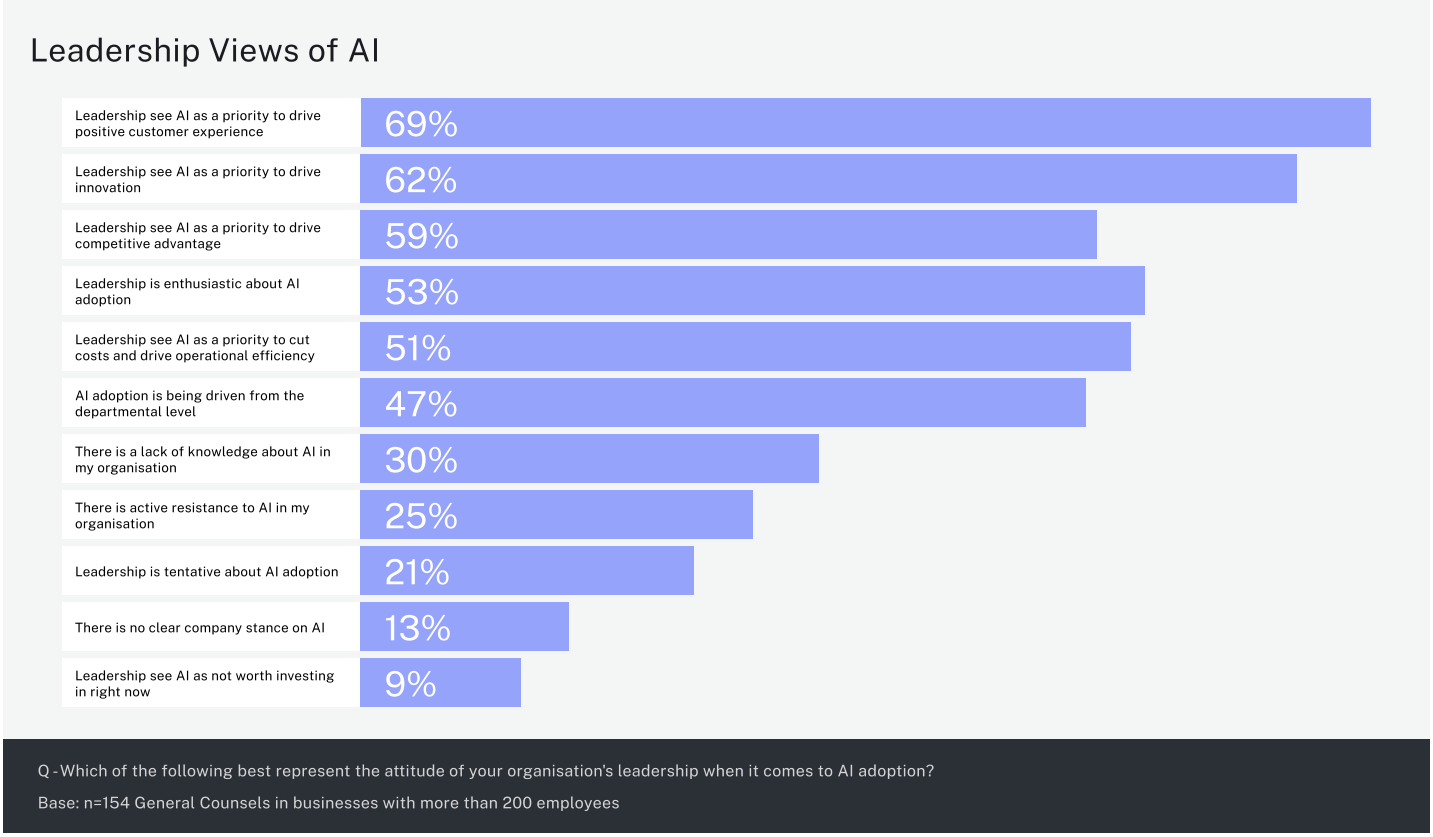
AI is becoming a standard part of how work gets done. AI is used by 57 percent of organisations several times a week. Among large organisations (2,500+ employees), daily use is at 34 percent. These firms are more likely to have formal governance structures in place, suggesting that scale and policy often go hand-in-hand.

The link between governance and adoption is clear. Organisations with formal AI policies are ten times more likely to use AI daily compared to those without. This suggests that where policies exist, teams feel more confident about when and how to use AI, and are less likely to delay adoption due to uncertainty or risk aversion.

By contrast, in organisations without clear oversight, adoption is slower or fragmented. A team may start using an AI writing assistant for client emails or case summaries, but without formal approval, training or tracking. This "shadow AI" can result in tools being used for purposes they were never evaluated for, increasing the risk of compliance breaches or reputational harm.

3

Leadership are enthusiastic about AI's potential



Q - Which of the following best represent the attitude of your organisation's leadership when it comes to AI adoption?

Base: n=154 General Counsels in businesses with more than 200 employees

Leaders support AI in principle, but lack of knowledge, coordination and ownership limits effective governance across the organisation.

Most leadership teams recognise the strategic value of AI, particularly in larger organisations. Over three quarters of organisations with 2500+ employees report enthusiasm for adoption at the top level.

However, this strategic intent is not always backed by structure. While 69 percent of leaders see AI as a priority to enhance customer experience and 62 percent to drive innovation, only 15 percent have good understanding of AI tools used and have governance frameworks in place. In nearly half of companies, AI adoption is led at the departmental level, without central oversight or shared governance.

This decentralised approach creates uneven standards. One business unit may evaluate tools for compliance and bias, while another adopts similar systems informally. Without a common framework, it becomes difficult to assess risk, track use, or ensure consistency.

Enthusiasm is high, but alignment is weak

AI is no longer just assisting decisions. In many teams, it is shaping them. Whether triaging customer issues, drafting contracts, or informing financial analysis, AI outputs increasingly influence core business operations. However, this shift has outpaced the readiness of many leadership teams.

A third of organisations cite a lack of AI knowledge at the leadership level. One in four report active resistance to adoption from senior stakeholders, often due to uncertainty about compliance, liability or public trust. In 13 percent of cases, there is no formal company stance on AI at all.

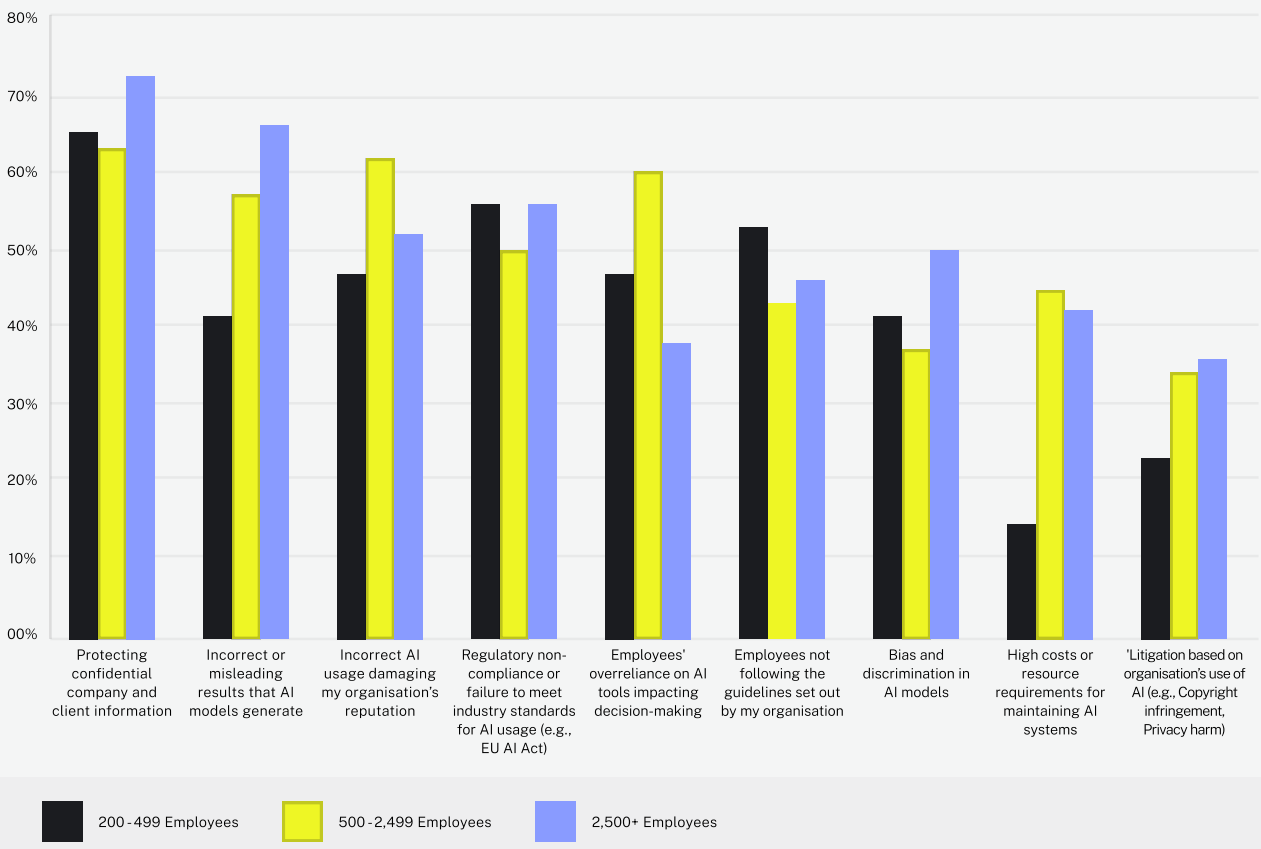
This misalignment suggests that policy gaps are not just operational – they reflect a deeper strategic ambivalence. Leaders may support AI in principle but are hesitant to commit resources, define boundaries or take accountability for outcomes.

25%

of General Counsels report active resistance to AI in their organisation.

4

AI risk perceptions



Q- When it comes to using AI in your organisation, which of the following do you feel poses the biggest risk?

Base: n=154 General Counsels in businesses with more than 200 employees (NET: 200-499, n=34; NET: 500-2,499, n=70; NET: 2,500+, n=50)

Organisations are most concerned about data misuse, inaccurate outputs and reputational harm – especially where AI is applied to sensitive or external data.

As AI becomes more deeply embedded in business workflows, risk awareness is rising. These concerns go beyond technical failure or productivity loss. They focus on legal exposure, reputational harm, and the risk of losing trust with clients, regulators and partners. These are not hypothetical risks. They reflect the practical consequences of AI systems being used without the right controls in place.

Risks are shifting from technical to operational

Most organisations use AI to process internal data, such as financials or performance metrics (56 percent). A significant share (34 percent) also use AI to analyse client and vendor data. For example, they might be reviewing contracts, summarising communications or drafting negotiation notes.

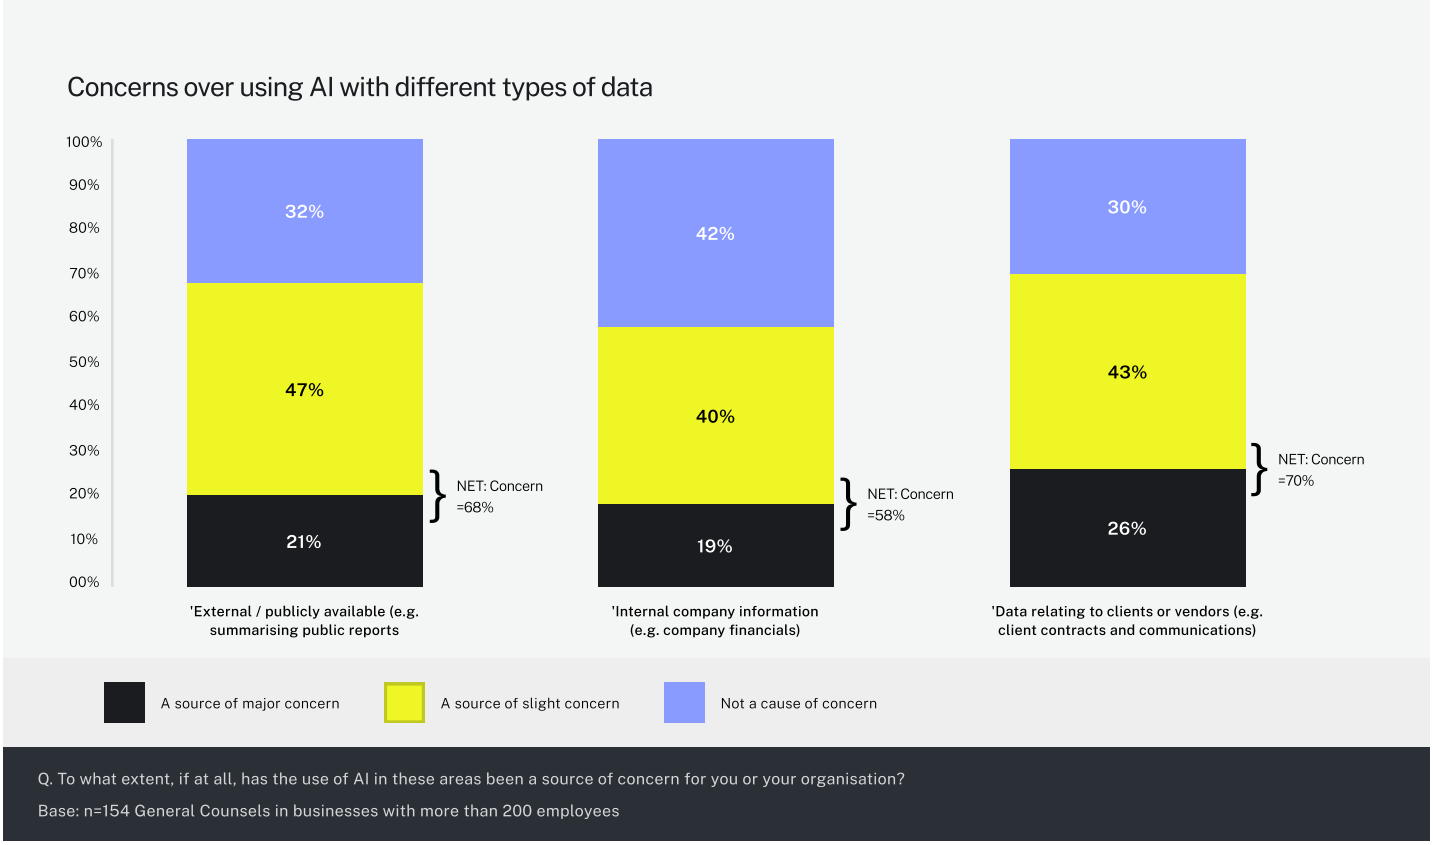
The most common concerns relate to protecting confidential data, preventing incorrect or misleading outputs, and avoiding reputational damage. These risks often overlap. For example, a financial AI tool might pull incorrect client data to generate investment advice. If that faulty output is shared, it could result in poor financial decisions, breach client confidentiality, and lead to a loss of trust in the firm.

AI is also being used to analyse and summarise performance data, customer interactions and internal reports. Without safeguards, these outputs can circulate beyond intended teams, surfacing information that is incomplete, outdated or commercially sensitive.

We’re now uploading and using a lot of sensitive data really frequently, daily, if not hourly, within this technology. So we need to make sure we’re doing that in a safe way. We're respecting our own policies, which we've created to tackle things like bias, hallucinations, and being good stewards of data but also recognising the fact that we're in a supply chain, and we need to respect our clients’ uses and policies on that.

Ravi Pau | Operations Director, Havas AI

Client and vendor data is the highest concern



Q. To what extent, if at all, has the use of AI in these areas been a source of concern for you or your organisation?
Base: n=154 General Counsels in businesses with more than 200 employees

Risk is high where AI systems are applied to client or vendor data. These datasets often contain contractual terms, personal information and commercially sensitive material. When processed by externally hosted or poorly governed tools, even a small oversight can lead to disproportionate consequences. An unintended disclosure might result in a contractual breach, fines under data protection law, or reputational damage that affects the organisation's standing with current and future clients.

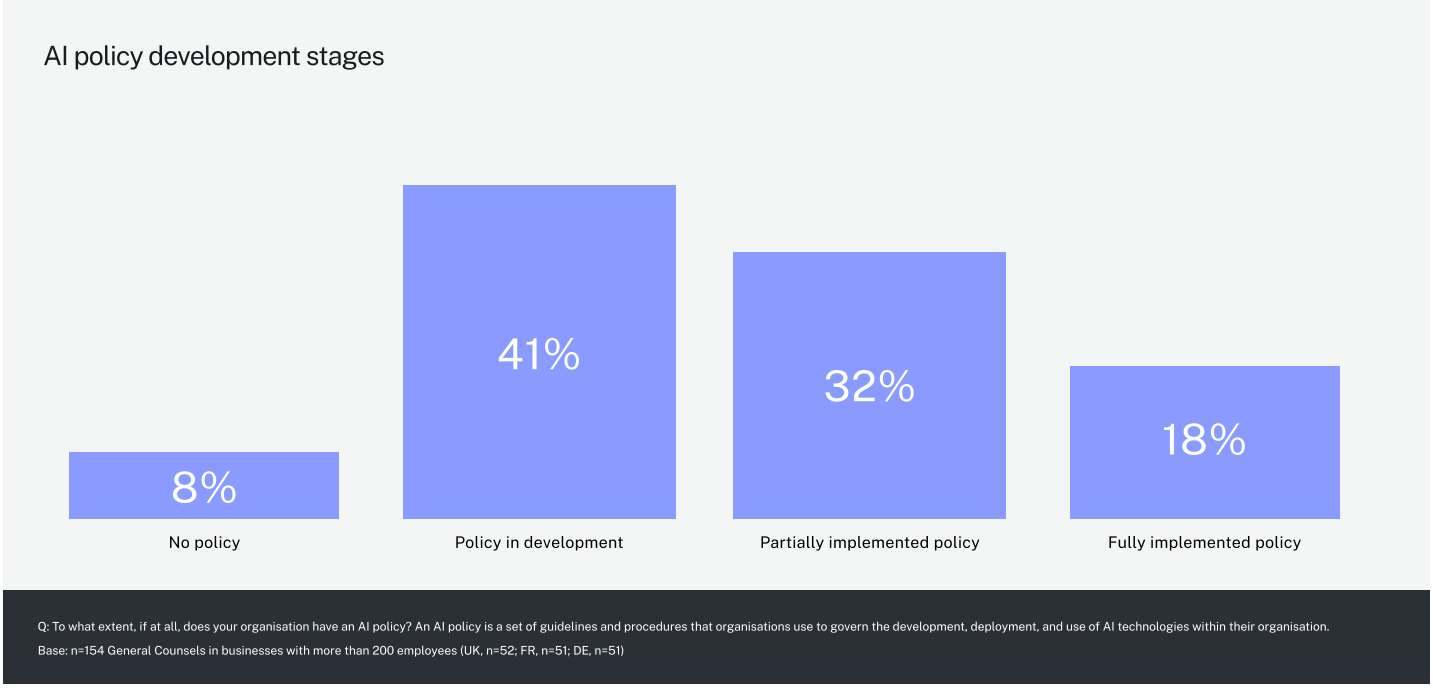
Internal data is also vulnerable. While less likely to result in external fallout, its exposure can undermine competitive positioning, reveal financial projections, or mislead internal decision-making.

Poor oversight amplifies risk

AI systems are not passive. They generate new content, often with the appearance of authority. If the underlying data is incomplete, inaccurate or mislabelled, the outputs may be wrong or misleading. These risks can scale quickly, particularly when AI tools are integrated into processes like decision support or document review. In regulated sectors, even minor inaccuracies can lead to compliance failures or regulatory penalties.

5

AI governance maturity is low



Most organisations understand their AI use, but few have formal governance frameworks, leading to fragmented oversight and hidden risks.

Ninety percent of organisations report active use of AI tools in day-to-day operations. Only 18 percent have a fully implemented AI governance framework. Many organisations face a mix of resource constraints, limited in-house expertise, and unclear ownership of AI governance. In some cases, frameworks are in development but stalled by uncertainty around evolving regulations or how to translate principles into daily operations. This means most organisations are using AI without the structures required to manage risk or ensure accountability.

The gap is more pronounced among frequent users. In many cases, frequent adoption occurs in the absence of clear oversight, increasing the risk of data misuse, compliance breaches or reputational harm.

Mid-sized firms are especially affected. Among organisations with 200 to 499 employees, 29 percent have no clear understanding of what tools are in use. By contrast, 32 percent of larger organisations (2,500 employees or more) report having a complete governance framework.

29% of smaller firms (200–499 employees) have no clear understanding of what AI tools are being used.

18%

of organisations have a fully implemented AI governance framework.

Fragmented oversight increases risk

Only a quarter of organisations track AI use centrally. In nearly half, oversight sits with individual departments, and in 14 percent, even individual employees are responsible for managing their own tools. Without a central registry or consistent reporting process, most organisations cannot say with certainty what AI tools are in use, who owns them, or how they are configured.

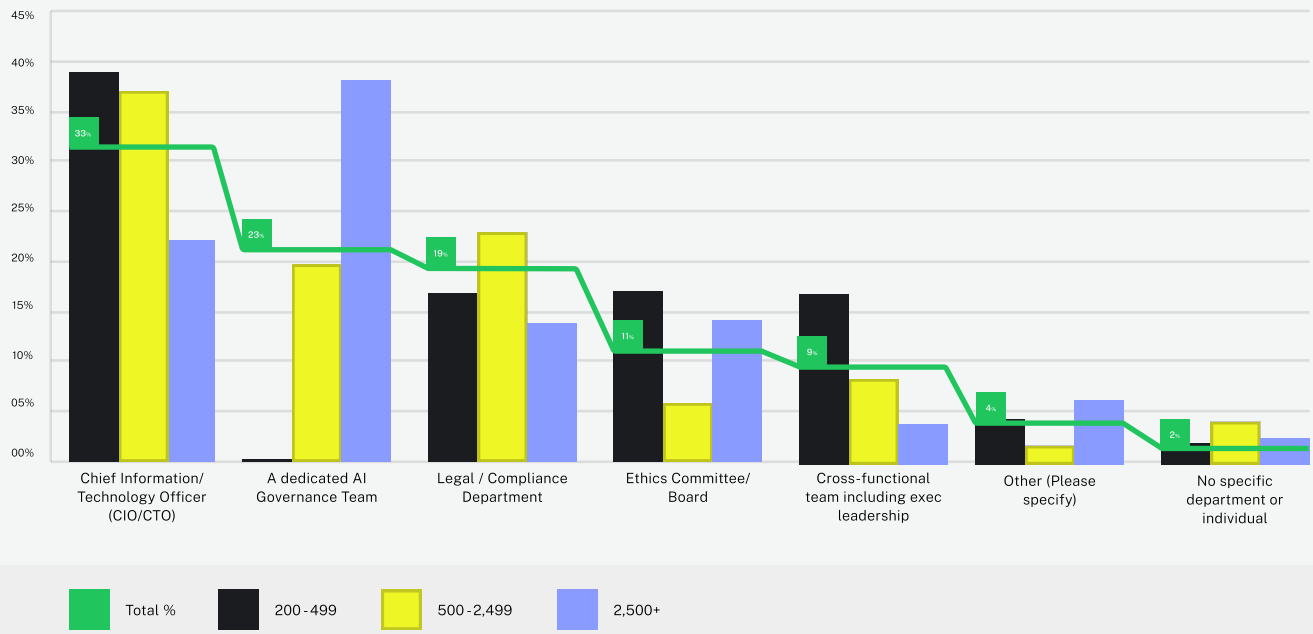
This lack of visibility creates major risks. 24% of firms admit they have limited or no understanding of what AI tools are in use, making it impossible to assess compliance or security exposure. Meanwhile, nearly 30 percent of companies with AI policies fail to explicitly protect client data, leaving a critical gap in privacy safeguards.

This lack of visibility allows “shadow AI” to thrive. For example, a customer service team might start using a free AI chatbot to summarise complaints. Without oversight, that tool could be processing sensitive customer data with unknown retention or sharing policies, potentially breaching GDPR or other data protection laws.

24%

have limited or no understanding of AI tools used across their organisation.

Who is responsible for AI governance?



Q23: Who in your organisation is most responsible for the implementation of your AI policy?

Base: n=141 General Counsels in businesses with more than 200 employees (NET: 200-499, n=23; NET: 500-2,499, n=69; NET: 2,500+, n=49)

Policy is not always enforced

Even where policies exist, implementation is inconsistent. Just 31 percent of organisations require formal approval for all AI use. Nearly half apply approval selectively, and 19 percent leave decisions to individual teams under general guidelines. This patchwork approach makes it hard to ensure that tools are used responsibly or that risks are identified early.

31% of organisations require formal approval for all AI use.

Policy adherence is also variable. Only 24 percent say usage rules are always followed. In most organisations, compliance depends on individual teams interpreting policy and voluntarily reporting use. This leaves large gaps in oversight, especially in fast-moving functions like marketing, sales or customer service.

Levels of approval needed and the extent to which policies are followed

Approval Policy		How often is it followed?			
All AI use requires prior approval	31%	Always 44%	Often 38%	Sometimes 7%	Rarely 11%
Only certain AI tools or use cases	48%	Always 17%	Often 51%	Sometimes 26%	Rarely 6%
Approval is not needed, but employees must follow AI usage guidelines	19%	Always 11%	Often 50%	Sometimes 25%	Rarely 14%

Q. You mentioned that your organisation's staff need to get approval before using AI solutions. Which of the below best describes your organisation's approach?
Q. And how often, if at all, is this approach followed by people working within your organisation?
Base: n=145 General Counsels in businesses with more than 200 employees, (selecting an approval process at Q10)

Principles are not enough. Without specific controls for model behaviour, vendor accountability and deployment oversight, organisations face exposure to:

- **Biased or inaccurate outputs**
- **Unapproved or high-risk use cases**
- **Data breaches linked to third-party tools**
- **Inability to meet regulatory expectations under the EU AI Act and DORA**

Policies must be more than compliance checklists. They need to guide real decisions and be embedded into how teams evaluate, procure and manage AI tools.

What AI policies cover, and what they miss



Q18 - You mentioned that your organisation has (or is) in the process of implementing an AI policy. What does it contain?
Base: n=141 General Counsels in businesses with more than 200 employees, with an AI policy

Policies often address high-level principles like ethics and privacy, but overlook operational controls such as vendor management, monitoring and use case evaluation.

AI governance is evolving, but many organisational policies are too general to be effective. While 82 percent of companies report having some form of AI policy, only a small share have moved beyond broad principles to operational frameworks that support safe and structured use.

Most policies stop at principles

Security (78 percent), data privacy (71 percent), and ethical commitments (65 percent) are the most commonly addressed areas. These are important foundations, but they do little to guide how AI is deployed, monitored or evaluated in practice.

Important safeguards are missing from many policies. Fewer than 41 percent include sections on vendor guidelines, system deployment, or performance monitoring. Only 35 percent address model training. Just 29 percent define how to assess or approve individual use cases. These omissions make it difficult to ensure that tools are behaving as intended, especially once deployed at scale.

Vendor oversight is underdeveloped

Many vendors across legal, marketing, HR, and finance now use AI to process sensitive data on behalf of clients. But only 41 percent of organisations include vendor governance in their AI policies.

With over 70 percent of organisations using AI and 40 percent applying it to client data, vendor activity represents a growing exposure. Even where internal controls are strong, third-party use of AI can create compliance, security, and reputational vulnerabilities.

Most AI systems are developed or hosted externally, limiting visibility into how they are trained, updated, or monitored. Without clear standards and contractual requirements, organisations may face data protection breaches, confidentiality lapses, or operational dependencies on unverified systems. In regulated sectors, these gaps can lead to enforcement action.

Smaller firms are most exposed

Among companies with 200 to 499 employees, only 6 percent report having fully implemented AI policies. By contrast, 30 percent of firms with over 2,500 employees report full implementation. This disparity reflects differences in resourcing, but also suggests organisations delay governance until AI is already embedded in day-to-day operations.

41%

of organisation include vendor governance in their AI policies.

PART TWO

THE SOLUTION

How to close the governance gap

Recommendations to enhance AI security and operationalise governance.

Some argue that stricter governance could slow innovation. But the answer to reducing risk around AI is not to limit usage. AI adoption is progressing at breakneck speed – 28% of enterprises with 2,500+ employees already see it as an 'essential tool' to their business operations. In reality, policies that restrict innovation are more likely to drive AI use underground. Appropriate governance enables safe, confident adoption at scale.

In order to accelerate the development of AI governance, and ensure responsible adoption of AI across European enterprises, three areas need to be tackled. They are:

- Better AI governance must be prioritised at board level.
- More guidance is needed from external sources
- AI governance must be operationalised

The following sections set out how to achieve this.

1

Better AI governance must be prioritised

Scaling AI safely requires more than top-level endorsement. It requires leaders to close the internal knowledge gap, set clear expectations for responsible use, and assign ownership for oversight and outcomes. Without this, AI governance will remain fragmented, with each department interpreting policy, or the absence of it, in its own way.

Organisations that succeed in embedding AI governance into enterprise strategy do three things well:

- They set a central policy direction that reflects legal, ethical and operational priorities.
- They give teams the guidance and tools needed to apply policy in practice.
- They ensure leadership remains involved in reviewing and evolving that governance as AI use grows.

Without clear ownership at the top, AI risk management becomes reactive and inconsistent. With it, governance becomes a strategic enabler, or a way to scale AI with control and confidence.

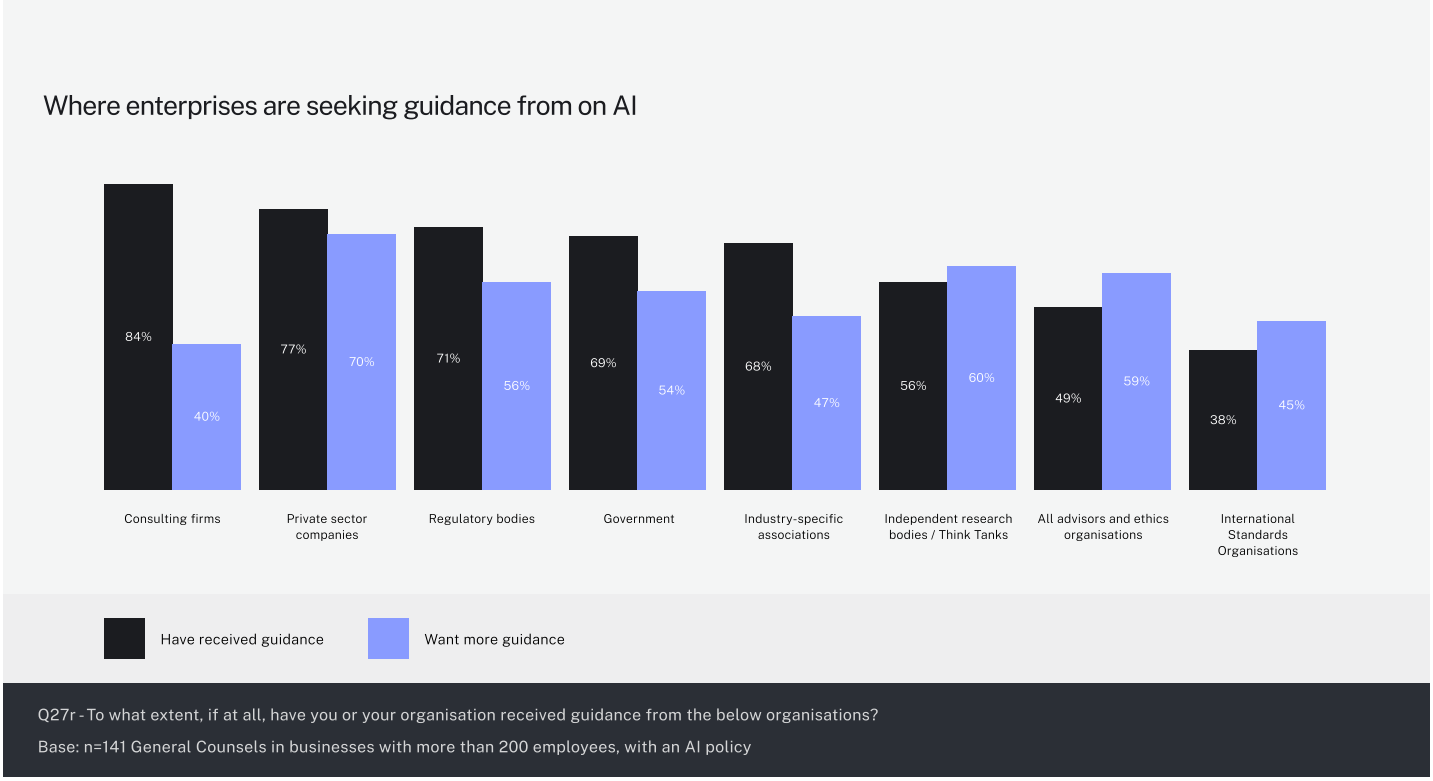


> The encouraging news is that governance and innovation go hand in hand – organisations that invest in proper AI frameworks see dramatically higher adoption rates and better outcomes. This research and its recommendations show that there's a clear path forward for companies ready to unlock AI's transformative potential responsibly.

Ruben Miessen | CEO & Co-Founder, LEGALFLY

2

More guidance is needed from external sources



Q27r - To what extent, if at all, have you or your organisation received guidance from the below organisations?
Base: n=141 General Counsels in businesses with more than 200 employees, with an AI policy

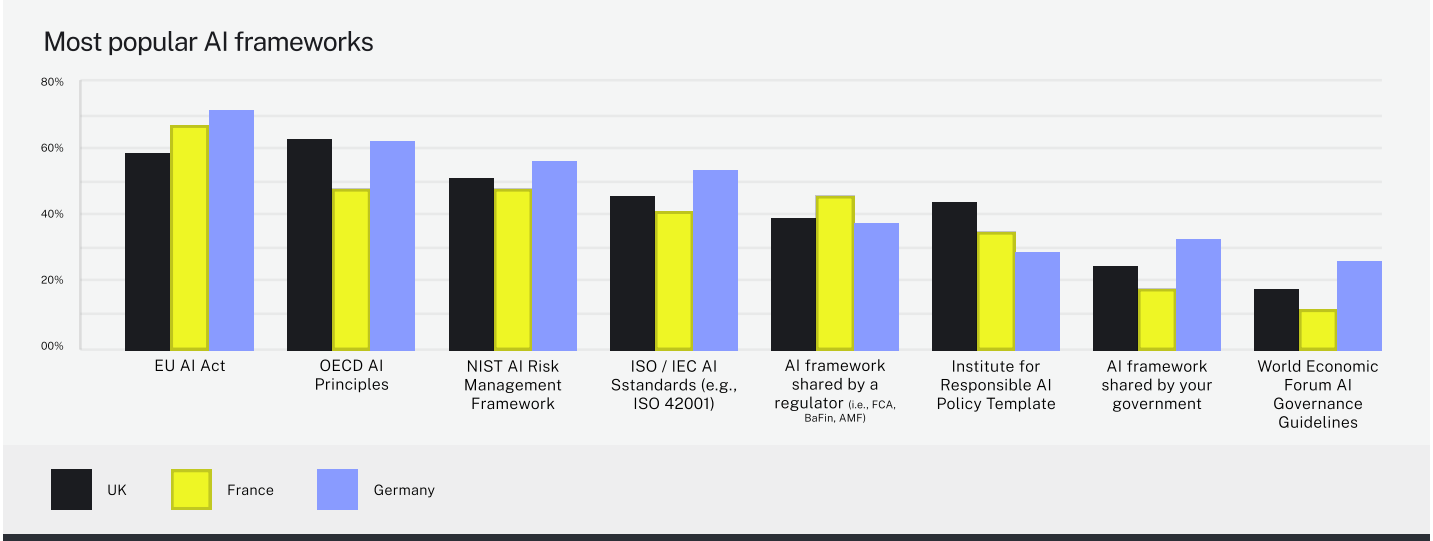
Generative AI is a complex new technology that is very rapidly becoming pervasive in European enterprises. General Counsels and organisational leadership cannot be expected to navigate the challenges of responsible governance alone. So far, organisations have received most guidance from consulting firms (84%) and private sector organisations (77%).

There is a significant appetite for more guidance from a range of sources. The majority of those surveyed would like more guidance from government and regulators, while 71% are interested in support from specialist private sector organisations, and 61% from AI advisors and ethics organisations.

56% of organisations say they would like more government guidance on AI.

How current AI frameworks are used

Many organisations are looking to multiple external frameworks to guide their AI policies. The EU AI Act is the most likely to influence AI policies and is seen as both robust and user-friendly. The OECD AI Principles and NIST's Risk Management Framework were also used by more than half of those surveyed to develop their AI policies.



Q20 - Is your AI policy based around any of the following AI frameworks?

Base: n=141 General Counsels in businesses with more than 200 employees, those with an AI policy (UK, n=48; FR, n=48; DE, n=45)

However, General Counsels are struggling to turn this guidance into action. While the principles are often clear, implementation is not.

For example, a company may be told to "ensure explainability" in AI outputs but not be given a clear definition of what that means in a commercial context, or the technical steps needed to achieve it. Another might be required to assess the risk of third-party AI systems without any benchmark for evaluation.

Matthew Kilkenny, Senior Policy Manager at the Open Data Institute, also points to a slowdown in regulation while industry bodies better understand the risks and how the technology will develop themselves.

"We're in what you'd call a 'sandbox period,' where we want to see what emerges before regulating. The slowdown in regulation across the US, UK, and EU reflects that. But none of these are unregulated industries. Every sector already has to comply with data protection laws, employment regulations, the Digital Markets Act, and more. It's not a legal Wild West. Businesses already face a huge regulatory burden, especially in the EU, where they must comply with a vast number of rules per country."

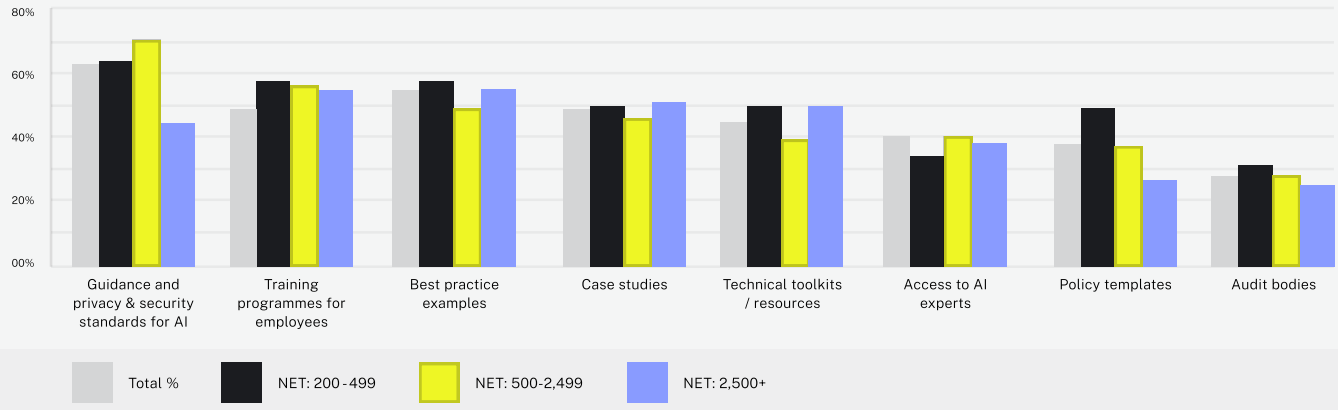
What guidance do enterprises seek?

The most consistent request from legal and compliance leaders is for actionable support. Most are not asking for new frameworks. They are asking for clarity, examples and operational tools.

This includes:

- Clear definitions of regulatory terms such as "high-risk system" or "meaningful human oversight"
- Case studies that demonstrate what good AI governance looks like in real settings
- Structured training programmes to build internal capability
- Templates and sample policies that show how others have translated rules into practice
- Guidance on the adoption of internal AI registries, risk-based approval processes, and designated ownership roles to embed governance into operations, not just documents

Guidance sought by General Counsels



Q28 - What additional guidance do you wish to see to help your organisation adopt AI?

Base: n=141 General Counsels in businesses with more than 200 employees (NET: 200-499, n=23; NET: 500-2,499, n=69; NET: 2,500+, n=49)

62% of those surveyed would like to see more guidance on privacy and security standards for AI.

3

Operationalising AI governance

In many organisations, AI governance is focused on written policies, but remains disconnected from daily work. Policies are often developed without input from the teams who are expected to apply them. Approval processes are added late, and oversight is distributed without clear accountability or support.

Effective governance must be operational. Organisations that treat it as a working system, rather than a static document, are better equipped to scale safely. This includes:

- Maintaining a central register of AI tools
- Linking approvals and reviews to risk levels
- Embedding checkpoints into procurement and onboarding
- Assigning ownership for AI use and outcomes

Without these elements, blind spots emerge. These can lead to data breaches, reputational harm or non-compliance. When in place, governance becomes a foundation for sustainable, risk-aware adoption.

Embedding governance into operations

Enforcement has to be part of how teams work. Success depends on:

- Leadership setting expectations and modelling responsible use
- Processes that reflect how teams actually work, not just policy templates
- Teams that are trained, resourced and empowered to apply governance in real time

Operational governance requires active monitoring, clear decision rights, defined roles for vendor oversight, and consistent enforcement. Organisations that approach policy as a dynamic framework – one built into workflows and maintained over time – are better placed to adopt AI responsibly and at scale.

Safeguards need to be built into systems, not added later

Policies alone are not sufficient. Organisations need technical safeguards that reduce risk at scale. Anonymisation is one of the most effective. When built into system design, it reduces the impact of unintended disclosure by masking or removing sensitive information before processing.

Additional measures such as access controls, data classification and audit logging help reduce reliance on manual oversight. They also provide the traceability needed to meet emerging regulatory requirements.

To manage this, organisations must move beyond policy documents and towards governance that is operationalised, technically supported and embedded into day-to-day practices. This requires close collaboration between legal, compliance, IT and data teams. If AI governance doesn't extend to how data is selected, tools are configured and outputs are validated, risk will continue to outpace control.

4

Six recommendations to enhance AI security and operationalise governance

1 | Recognise that GenAI is already in use

GenAI is already part of most organisations' workflows, often through personal accounts or unsanctioned tools. Employees adopt external apps when official channels lag, creating invisible risks. On average, companies use hundreds of AI tools, many of them untracked.

The first step is to acknowledge this reality and bring usage into view. Identify what is already in use and secure privacy-preserving alternatives that meet organisational and regulatory standards.

2 | Use a common risk framework

A shared risk framework helps legal, compliance and IT teams speak the same language and make coordinated decisions.

Adopting something like the NIST AI Risk Management Framework provides a structured way to identify, govern and reduce AI-related risks. It also supports consistency across departments and aligns internal controls with regulatory requirements.

3 | Build privacy and data safeguards into workflows

Governance is not just policy. It is system design. Technical safeguards should operate at scale, automatically and by default.

This includes data minimisation before information reaches an AI model, applying strict access controls to vector stores and retrieval caches, aligning retention with internal policies, and building in deletion or correction options to comply with data rights.

4 | Use anonymisation for sensitive data

Anonymisation is a practical, scalable way to reduce risk while preserving context. Tools like LEGALFLY automatically strip out personal, financial, or confidential business information before data is processed by AI systems. This applies not only to legal documents, but to any sensitive content.

5 | Treat prompt injection and model misuse as security issues

AI introduces new categories of application risk. Prompt injection, misuse of system instructions and ungoverned third-party integrations can widen the blast radius of a compromise.

Organisations should filter inputs and outputs, use layered prompt structures, isolate memory where possible, and run red-team testing to detect and contain failures. These are standard security practices that now need to apply to AI systems as well.

6 | Replace shadow AI with approved, monitored tools

Blocking tools outright drives adoption underground. A more effective approach is to offer sanctioned alternatives that are easy to access and tightly governed.

Build an allowlist of approved tools, route usage through single sign-on, and log all interactions. This allows for e-discovery, anomaly detection and auditability, without disrupting productivity.

This report set out to understand how AI is being governed in practice, where current approaches fall short, and what is needed to close the gap.

Based on responses from 154 general counsel and expert interviews across the UK, France and Germany, it is clear that AI adoption has outpaced governance, exposing organisations to operational, legal and reputational risk.

The pace of AI adoption is not slowing down. To keep Europe on the path to responsible AI, better governance needs to be prioritised at board level, and more guidance should be made available from government and specialised private companies. The most important actions for enterprise leadership to take are to move from abstract principles to operational safeguards for AI and the recommendations provided here provide some first steps to do so.

To continue the conversation about responsible AI adoption, email us at responsibleAI@legalfly.com

ABOUT LEGALFLY

LEGALFLY is the leading agentic AI workspace for in-house legal and compliance teams. Some of the largest enterprises across Europe use LEGALFLY to streamline contract reviews, drafting, and compliance workflows, increasing operational efficiency.

Founded in 2023, LEGALFLY's mission is to empower legal professionals to deliver more value with AI, safely and responsibly. LEGALFLY is the only Legal AI platform to anonymise all sensitive data before processing begins, ensuring that confidential information is protected at all times.

legalfly.com