TRUSTMARQUE

AI Governance Index

A report into the state of AI governance in enterprise

2025

trustmarque.com

IBM Gold Partner

TRUSTMARQUE

Contents

Foreword

AI no longer sits on the sidelines. It is a core pillar of business strategy, powering everything from customer engagement to process automation and strategic decision-making. What started as experimentation within individual teams has, in many cases, become embedded in day-to-day operations, with Generative AI driving much of this shift.

Implementation has been deliberate, often well-intentioned, and in some places highly effective. But while adoption has surged, governance has lagged. Many organisations are pushing ahead without the structure needed to manage risk, ensure accountability, or provide transparency. Formal processes are patchy. Ownership is fragmented. Technical safeguards are often missing altogether.

In many cases, AI initiatives are springing from the ground up, developed in pockets, outside of central oversight. These efforts are often innovative, but lack cohesion. Risk awareness varies widely. Cross-functional collaboration is rare. And executive leadership, if present, is often passive rather than proactive.

The result is a widening governance gap.

Organisations are deploying Generative AI faster than they can govern it, and the consequences are real. Privacy breaches, operational instability, ethical missteps and loss of stakeholder trust are all risks that grow in the absence of effective oversight. These aren't just compliance issues. They are long-term risks to business resilience and trust.

One of the clearest signs of this disconnect lies within the software development lifecycle. Most organisations haven't meaningfully adapted their SDLC or DevOps processes for AI. Testing for fairness, bias or explainability is inconsistent. Post-deployment monitoring is limited. AI risks are rarely embedded into development workflows in the way security or data quality risks might be.

There is progress, however. Many organisations are beginning to invest in better controls and rethink how AI is built and governed. What is needed now is structured action and executive direction, not to slow innovation, but to make it sustainable and secure.

This report, drawing on cross-sector research of the opinions of 500 IT decision makers, is both a temperature check and a practical guide to action, giving decision-makers the insight to move from reactive risk management to confident, coordinated AI governance. It highlights where organisations are today, where governance is falling short and steps to close the gap.

AI is here to stay. Governance must be too.

Seb Burrell
Head of the AI Centre of Excellence
Trustmarque

Key findings

1. AI IS BEING ADOPTED FASTER THAN IT IS BEING GOVERNED

- 93% of organisations are using AI in some form, over a third report widespread use throughout the organisation.
- Governance maturity lags significantly behind adoption - only 7% have fully embedded it, over half - 54% - have either no governance or it is very limited in scope.
- Generative AI tools are being used widely, but often without clear oversight or formal policy.

2. GOVERNANCE IS FRAGMENTED, INCONSISTENT AND OFTEN INFORMAL

- Responsibility for governance is split across IT, legal, compliance and data teams - with 19% reporting no clear ownership.
- Many organisations carry out security reviews or anomaly monitoring, but few apply bias detection or explainability testing.
- Generative AI tools are being used widely, but often without clear oversight or formal policy.

3. EXECUTIVE ENGAGEMENT IS PATCHY AND OFTEN PASSIVE

- Only 9% say AI strategy and governance are fully aligned with executive leadership.
- Boards and C-suites are driving governance in 20% of organisations - indicating a lack of top-down leadership.
- Most activity is being driven from the ground up, with teams managing risk as best they can.

4. INFRASTRUCTURE AND TOOLING ARE NOT READY TO SUPPORT SCALE

- Just 4% say their infrastructure and data environment is fully AI-ready.
- Model versioning, registries and audit trails are mostly manual or missing entirely.
- Skills shortages, budget limitations and unclear ownership are the top barriers to better governance.

5. TESTING AND LIFECYCLE MANAGEMENT ARE UNDERDEVELOPED

- There is testing for AI in the SDLC with 44% applying governance or testing before deployment, and 41% having post-deployment monitoring in place.
- Only 8% have fully integrated AI governance into their software development lifecycle.
- Most organisations describe their approach as either ad hoc or fragmented across teams.

6. CROSS-FUNCTIONAL COLLABORATION IS STILL IN EARLY STAGES

- 42% say they occasionally involve legal, ethics or HR in AI-related decisions.
- Only 20% have formal governance groups that span functions.
- A lack of joined-up oversight increases the risk of inconsistent or non-compliant use.

CHAPTER 1

Organisation maturity and AI adoption

If you don't align governance with adoption, your AI efforts will remain at risk. The faster AI is integrated into business functions, the more critical it becomes to manage its risks and accountability.

AI is embedded in the plans and operations of most enterprises. Nearly all organisations in our study reported active AI usage, with over a third saying it is already embedded in multiple functions. Adoption is no longer the bottleneck, or the question.

93% of organisations are using AI.

7% only have fully embedded governance practices.

54% report using AI at scale for two years or less, signalling early maturity.

Governance maturity is not keeping pace. More than half of respondents are operating with limited or no formal AI governance in place. Just 7% have a fully embedded framework. When asked about governance drivers, only 20% of organisations pointed to executive or board-level pressure, a sign senior leadership engagement remains limited. There is an opportunity and a need for stronger oversight and influence.

The disconnect exposes organisations to unchecked risk.

Top AI use cases

Data analysis and reporting	62%
Customer support/chatbots	50%
Document automation	43%
Content generation	39%
Code development	34%
Predictive modelling	33%

Most AI adoption is still happening at the departmental level, led by teams chasing operational improvements. It's fast, it's iterative, and often effective. It's also largely invisible to central governance teams. Without oversight, you can't assess risk, ensure fairness, or apply consistent standards across your AI use cases.

Only 6% of leaders said there was no awareness of AI governance within their organisation. Awareness exists. Execution doesn't. This suggests that the blockers are less about intention, and more about ownership, resources, and clarity on next steps.

54% Report their governance is either non-existent or minimal.

8% Only have fully embedded governance practices.

Top governance drivers

Security concerns	62%
Compliance mandates	55%
Reputational risk	44%
Executive/board pressure	20%

The consequences of this maturity gap are not abstract. Without governance, AI systems can drift in behaviour, replicate bias, or make decisions that are impossible to audit. All this erodes business resilience and public trust.

AI success depends on more than smart deployment. It demands deliberate design, accountable ownership, and cross-functional governance from the outset. Until adoption is matched by maturity, organisations will remain exposed.

CHAPTER 2

Governance frameworks, risk and compliance

Governance isn't just a set of policies. It's the infrastructure of trust that makes AI scalable, safe, and sustainable. Even with growing awareness, most organisations are still figuring out how to turn principles into practice.

Many organisations understand the risks AI presents, from security breaches to biased decisions. But understanding isn't the same as action. While awareness is high, our findings show that governance remains underdeveloped. Only 7% of organisations have adopted a formal, enforced governance framework. A further 54% are in the early stages, experimenting with tools and policies but without consistent application. The gap between intent and implementation is one of the biggest challenges facing enterprise AI leaders.

What kinds of risks are being prioritised?

The focus is heavily weighted toward traditional IT and compliance concerns: security threats and privacy violations top the list. Half are focused on inaccurate or misleading outputs. But fewer organisations are tackling the deeper, AI-specific risks, such as explainability, model bias, or model drift.

Risk priorities

Security threats	65%
Privacy violations	52%
Inaccurate/misleading outputs	49%
Lack of explainability	29%
Model bias	28%
Model drift	18%

While risk awareness is widespread, risk controls are patchy. Nearly half (46%) monitor for anomalies, and 43% conduct security reviews before deployment. But fewer apply bias detection (28%) or track data lineage (32%). Worryingly, 19% apply no risk controls at all.

Most organisations don't consistently measure the effectiveness of their governance. A quarter rely on ad hoc reviews, and only 18% have continuous monitoring with KPIs. Another 18% don't measure governance effectiveness at all. This leaves a critical blind spot in their AI risk strategy.

24% Use ad hoc reviews

22% Periodic internal audits

18% Use continuous monitoring against KPIs

18% Do not measure effectiveness at all

When it comes to external support, organisations are still largely self-reliant. 43% manage governance entirely in-house, and only 5% use external partners for ongoing oversight, despite the availability of expert services and tools.

Governance responsibility is also fragmented. In most organisations, accountability is shared between IT, compliance, legal, and data teams. But with no clear owner, implementation is harder, slower, and less effective.

The foundations for AI governance are being laid, but they are uneven and incomplete. Without clearer ownership, more consistent controls, and a shift from awareness to action, organisations risk leaving critical vulnerabilities unchecked.

CHAPTER 3

Development lifecycle and deployment

Organisations must embed AI governance directly into how they build, test, and deploy AI. Without this, even well-intentioned initiatives risk failing silently, or causing damage.

While many organisations have started to explore governance frameworks, far fewer have integrated those principles into their development lifecycle.

AI introduces unique risks, from data drift to explainability gaps, that traditional DevOps practices weren't designed to handle. Yet most organisations still rely on legacy processes that overlook these issues.

Our research shows that just 7% of organisations have fully integrated AI governance across the software development lifecycle. A further third say it's partially implemented in some teams, while 26% admit they have no integration at all.

Governance in the development lifecycle

- **7%** Fully integrated
- **33%** Integrated in some teams
- **33%** Informal/inconsistent
- **27%** No integration

Governance touchpoints are similarly inconsistent. While 46% of organisations apply testing or governance checks pre-deployment, and 41% have post-deployment monitoring in place, the research found nearly one in five report no governance activities at any point in the development lifecycle.

Governance touchpoints during the development lifecycle

- **44%** Pre-deployment
- **41%** Model design
- **41%** Post-deployment
- **20%** No governance

Testing approaches vary widely. Functional testing (61%) and security checks (59%) are most common. But testing for bias (33%), performance drift (46%), or interpretability (22%) is less consistent. And 14% of organisations conduct no testing on their AI models at all.

There's also a structural challenge around ownership. In most organisations, governance responsibility is distributed across multiple teams.

- **49%** Cite central IT leadership
- **30%** Rely on risk or compliance functions
- **28%** Name data science or analytics teams
- **26%** Point to product or application owners
- **19%** Say there's no clear governance owner in development, a major block to accountability

In practice, many organisations are attempting to adapt SDLC practices through tactical fixes, layering in AI-specific risk assessments or MLOps tooling.

But these efforts are often underfunded or siloed. This is reflected in how organisations assess their own maturity.

- **21%** Have no formal governance or testing process
- **25%** Reactive or ad hoc approach to governance
- **38%** Describe their governance as adequate but fragmented
- **16%** Governance is well-structured and consistently applied

The research found that 88% of IT decision makers treated AI models differently to traditional software applications in terms of testing. This demonstrates the challenges of integrating AI into a traditional SDLC.

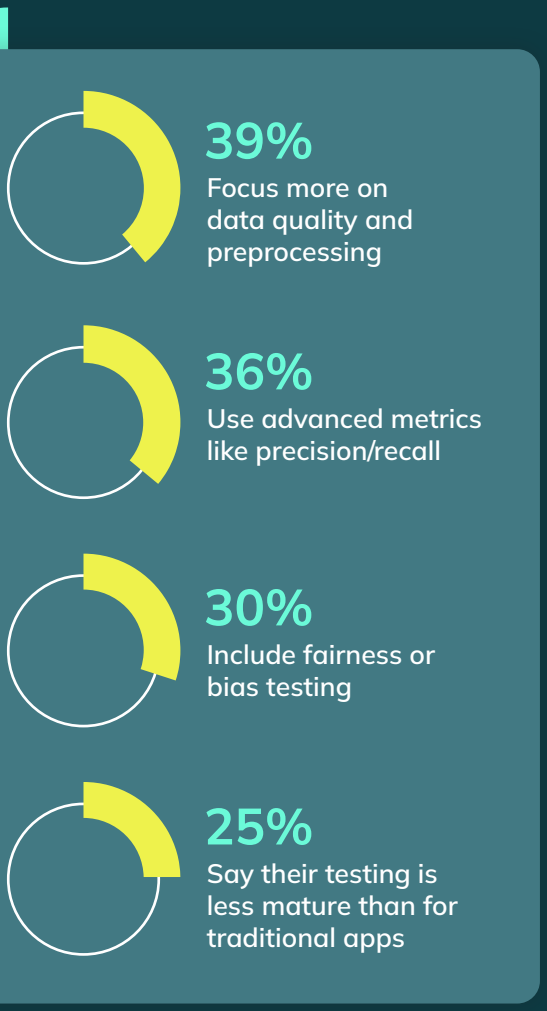

AI governance only works if it's built into the pipeline, not bolted on at the end. Organisations need to move beyond policy and make governance a living part of the development lifecycle. That means adapting testing protocols, clarifying ownership, and funding the operational changes that AI demands.

CHAPTER 4

Architecture, infrastructure and operations

Effective AI governance goes beyond policy and risk management. It depends on solid technical foundations. From data infrastructure to model registries, audit trails to monitoring tools, the operational layer plays a critical role in enabling governance at scale. Yet for many organisations, the underlying systems needed to support secure and sustainable AI deployment are incomplete, underfunded or still in development.

Just 4% of organisations say their infrastructure and data environments are fully AI-ready. Most fall into moderate (35%) or basic (31%) categories, and nearly one in five (18%) describe their infrastructure as very limited. Data environments mirror this trend, with only 3% claiming full readiness.

AI infrastructure readiness

- 4% Fully ready
- 12% Scalable
- 35% Moderate
- 31% Basic
- 18% Very limited

Model management is another gap. Only 11% of organisations maintain an automated and centralised model catalogue, while 19% rely on manual registries. A third (34%) have no registry at all, and 37% say theirs is still in development.

Model registry status

- 11% Automated and centralised
- 19% Manual
- 37% In development
- 34% None

Traceability is similarly unreliable. Just 16% have full audit trails across versions and deployments across AI model configurations while 19% report partial automation. Over a third rely on ad hoc or manual tracking of versioning, and 27% have no versioning practices in place.

The top barriers to infrastructure improvement reflect familiar enterprise challenges.

Barriers to infrastructure maturity

Lack of skilled personnel	54%
Budget constraints	54%
Competing priorities	44%
Unclear ownership	31%
Low executive prioritisation	27%
Fragmented tooling	22%

To close the gap, organisations are investing in people. Nearly half (46%) are upskilling internal teams to address shortfalls, while 20% are partnering with external consultants. But a third (33%) aren't yet addressing the skills gap at all. This is a worrying signal given the operational demands of AI.

Organisations are responding to the operational demands of AI in different ways.

While many are taking steps to build internal capability or bring in external expertise, a significant proportion have yet to take action.

How organisations are dealing with operational demands of AI

46%
Upskilling internal teams

20%
External partnerships

33%
Not addressing

What's clear is that even well-written governance frameworks can fail in practice without the technical plumbing to support. Without monitoring platforms, model registries, versioning and audit tools, it becomes difficult to enforce policy, ensure transparency, or respond to compliance risks in real time.

Governance maturity depends on infrastructure maturity. To move from principle to practice, organisations need to invest in the tooling and operational architecture that makes oversight possible, and reliable.

CHAPTER 5

Culture, leadership and strategic fit

Strong governance starts with organisational culture: on leadership that prioritises responsible AI use, on cross-functional collaboration, and on a shared understanding of AI's role in business strategy. Without these foundations, even the best technical controls will fall short.

Only 4% of organisations say AI is fully embedded in their business strategy, and just 15% report strong alignment. The majority fall somewhere in the middle: 40% describe alignment as partial or inconsistent, while a combined 41% say it is minimal or absent altogether.

AI and business strategy alignment

4%	Fully embedded
15%	Strong alignment
26%	Minimal
15%	No alignment

Executive involvement follows a similar pattern. Only 9% of organisations report full strategic alignment between AI governance and IT leadership. Another 25% collaborate regularly, and 33% say there is occasional consultation. But low engagement (22%) or no involvement at all (10%) is still common.

When support does exist, it's expressed in different ways. The most usual indicators of leadership support are inclusion in strategic roadmaps and budget allocation. Some organisations go a step further with executive committees and performance-linked KPIs. But in a significant proportion of organisations, leadership support is limited to simple verbal encouragement.

Leadership support indicators

Strategic roadmap	45%
Budget allocation	32%
Executive committees	30%
KPIs tied to governance	28%
Tooling investment	27%
Verbal support only	20%

Culture also plays a major role. 36% of organisations describe themselves as cautiously supportive of IT-led innovation, and 21% say they are innovation-focused. Yet 24% report a risk-averse culture, and just 8% describe their approach as proactive and experimental. Another 10% say their organisation is resistant to change.

Cultural posture chart

Cautiously supportive	36%
Risk-averse	24%
Innovation-focused	21%
Resistant to change	10%
Proactive and experimental	8%

Cross-functional collaboration is another weak spot. While 30% of organisations report regular collaboration with legal, HR, or ethics teams, 42% say it only happens occasionally. 19% report no collaboration at all, and just 10% have a formalised, multi-disciplinary governance group in place.

Without shared responsibility and cross-team alignment, governance becomes fragmented. Strategic decisions about AI are harder to coordinate, risk ownership is unclear, and frameworks are harder to embed.

Culture is the control layer that sits above governance. Without visible leadership, clear alignment to strategy, and true collaboration across teams, governance efforts will struggle to take hold and last.

Conclusion

AI adoption is accelerating, but governance is not keeping pace. Most organisations recognise the need for oversight, but few have the policies, processes, or infrastructure to manage AI risks consistently across teams.

Our research reveals a picture of ambition constrained by fragmentation. Adoption is high, but governance is uneven. Technical safeguards are emerging, but cultural foundations are weak. Leaders are interested, but often not engaged.

This gap between intent and execution is the central challenge facing enterprise AI today. Bridging it will require more than frameworks. It demands action:

- Aligning AI strategy with business priorities
- Embedding governance into development workflows
- Investing in infrastructure and skills
- Making accountability a shared, cross-functional commitment

The foundations are being laid, but maturity will only come with clear ownership, sustained leadership, and operational follow-through.

Governance is not a blocker to innovation. It's what makes innovation sustainable.

The organisations that act now will be the ones best positioned to scale AI with confidence, resilience, and trust.

TRUSTMARQUE

IBM Gold Partner

AI governance can't wait. Whether you're scaling pilots, navigating compliance, or embedding AI into critical business processes, now is the time to build the foundation for sustainable, trustworthy AI.

Trustmarque and Acutest can help you:

- Assess your current governance maturity
- Implement risk and testing frameworks
- Align strategy, culture, and infrastructure

About Trustmarque

Trustmarque delivers the value of technology to bring real-life impact. As a trusted partner to both customers and technology vendors, together, we turn your vision into reality. Trustmarque's knowledge, experience, and technical expertise helps organisations acquire and adopt the right technology to create an environment of innovation.

AI testing and governance in action

Acutest, Trustmarque's specialist quality assurance and governance practice, helps organisations operationalise responsible AI.

With the support of cutting-edge Governance tools like IBM's watsonx.governance combined with their deep experience in software testing, risk management, and compliance, Acutest enables AI programmes to move from concept to deployment with confidence.

From validating fairness and explainability to embedding traceability, Acutest supports the practical implementation of governance controls within AI development workflows. Their AI Testing and Governance service combines state-of-the-art Governance tools, technical assurance and strategic insight, helping organisations stay ahead of regulatory expectations while maintaining pace with innovation.

Methodology

This report is based on independent research commissioned by Trustmarque and conducted between **10 May and 1 June 2025.**

The findings are drawn from a structured survey of **507 IT** decision-makers based in the **United Kingdom,** representing a range of industries and organisation sizes. The research explored how enterprises are approaching AI governance - what's in place, what's missing, and how organisations are adapting to manage risk and scale AI use responsibly.

Respondents included a mix of seniority levels and technical roles to ensure a representative and well-rounded dataset.

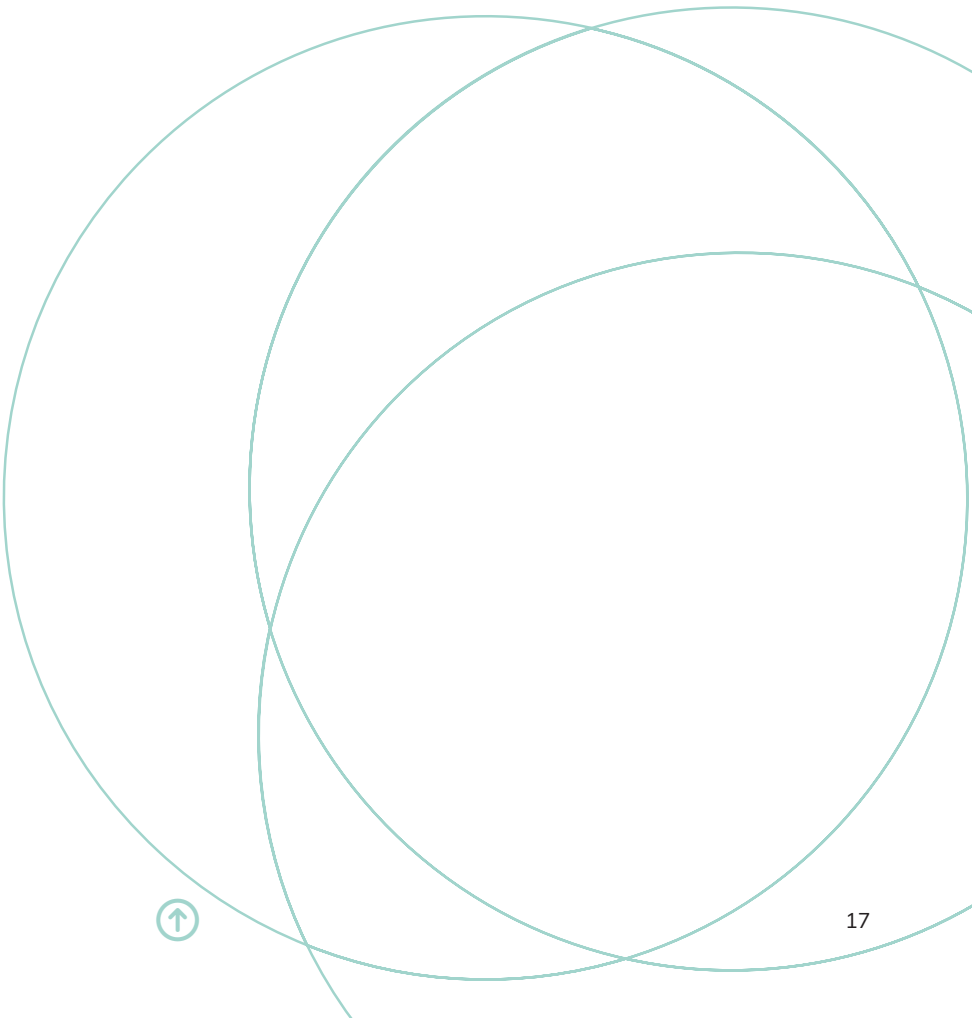

0845 2101 500 | info@trustmarque.com | trustmarque.com