




The New Realities of AI Governance

What the rise of AI agents means for governance and accountability in EMEA



Contents

Foreword	3
Introduction: The AI Governance Gap	4
AI Governance Reaches the Boardroom	5
Deploying AI Before Governance is Ready	6
Accountability and Liability Remain Unsettled	7
Governing AI Beyond the Enterprise	8
Regional Perspectives	9
Governance and Business Outcomes	10
Five Priorities for Governing Autonomous AI	11
The Next 18 Months: Barriers to AI Returns	13
Study Methodology	14

Foreword

Governing AI at the Pace of Change



Ana Paula Assis, Senior Vice President and Chair, EMEA and APAC, IBM

AI agents are opening up a new chapter for business. Agentic systems that can plan, decide and act on behalf of organisations promise to transform how work gets done, and across Europe and the Middle East, organisations are already moving to capture that potential.

Many are approaching this from a strong starting point. AI governance has moved firmly into the boardroom: nearly every organisation we surveyed has assigned responsibility for AI risk at board level, and when something goes wrong, it reaches senior leaders. Governance is also being refined through real-world experience.

But agents raise the bar. When agentic systems can act continuously, without a human approving each step, governance can't rely on sign-off alone. It must be built into how agents operate, with clear boundaries on what each one can do, clear ownership when things go wrong, and controls that work while agents act. That is what gives businesses the confidence to extend AI's role within their organisation.

So, how prepared are organisations for this shift? When AI incidents reach the boardroom, what happens next? Who is accountable when an agent gets something wrong? How far does a company's governance extend to the AI tools that its suppliers and partners are using? And are organisations ready for agents that act on their own?

We asked more than 3,200 senior executives across eight markets in Europe and the Middle East to find out. Our research, conducted with Censuswide, reveals where organisations are making progress, where the governance gap remains, and five priorities for closing it. The organisations that act now to get governance right will be the ones best placed to realise the full potential of AI agents.

Introduction

The AI Governance Gap

For most large organisations, the question is not whether AI governance is in place – nearly all organisations surveyed (98%) have formally assigned responsibility for AI risk at board level, while 26% have a dedicated board-level AI committee. As AI agents take on more autonomous work, the harder question to answer is whether governance can keep pace with how AI is being deployed.

Eight in ten executives in the study say an AI-related incident reached their board or C-suite in the past year. Three-quarters expect AI agents to have some level of autonomy towards decisions by 2027. But only 37% say their current governance is designed to scale with agentic AI.

This is the AI governance gap: organisations are deploying increasingly autonomous systems, but their ability to oversee, constrain and assign responsibility for those systems is still being developed.

When AI deployment outpaces governance readiness, 56% of organisations proceed within agreed risk limits while adapting

their frameworks; one-third (33%) wait until governance is ready. This highlights the challenge of moving forward without letting oversight fall behind.

The research findings point to four related challenges:

- Board-level attention does not always result in policy change
- Organisations must often navigate AI deployment outpacing governance readiness
- Accountability and liability remain unclear
- Controls often do not extend far enough across the supply chain

AI Governance Reaches the Boardroom but Does Escalation Change Policy?

AI incidents, such as an agent failure, data exposure, flawed AI-driven decision or security breach, are bringing risk to the attention of senior leaders. Of the eight in ten executives who say an incident reached their board or C-suite in the past year: 43% say it was formally reviewed at that level, while 37% say it was communicated informally.

Escalation does not always result in a change to governance policy. Among organisations

where an incident reached senior leaders, 40% cite a change to AI governance policy as a primary outcome. Similar proportions report changes to investment or deployment plans (38%), engagement of external advisers (37%) or notification of regulators (35%).

The findings suggest that board-level ownership is widespread, but organisations still need to translate attention into consistent policies, processes and follow-up.

Case Study

Banco Santander is establishing a comprehensive AI Governance framework to enable the safe, compliant, and scalable adoption of AI across the organisation. Covering traditional AI, Generative AI, and Agentic AI, the framework provides governance, risk management, compliance oversight, and end-to-end traceability across the AI lifecycle.

Leveraging IBM's technology and expertise, Santander has deployed a global AI Governance Platform built on [watsonx.governance](#) and [watsonx Orchestrate](#) across all operating countries. The platform enables the governance of AI models, agents, and LLMs while addressing country-specific regulatory requirements, providing end-to-end visibility and control of the bank's AI ecosystem.

Deploying AI Before Governance is Ready

Deployment of agentic AI is progressing beyond experimentation: just over a third of organisations surveyed are already using it in production and 75% expect agents to act without human approval of every action by 2027.

Maturity varies widely. Around one in five (21%) report no current agentic AI deployments and 29% are evaluating agentic AI but have not launched a pilot, while 16% are at the early pilot or proof-of-concept stage. A further 9% have AI agents in production in one or two functions, 13% across multiple functions and 12% embedded across the enterprise.

However, governance is often still catching up. Just over a third (37%) say their governance is designed to scale with agentic AI, while 45% say it could cope but needs strengthening. Although 54% govern agentic AI as a distinct risk category, only 37% have formally quantified the financial exposure created by governance gaps, with board review.

As agents take on more work, oversight will need to focus less on approving every action and more on setting clear boundaries, supervising outcomes and managing exceptions. Organisations can make progress while that shift takes place by letting autonomy grow in step with governance: starting with agents in well-defined, lower-risk processes; being clear about what each agent can do and when it must involve a person; and putting runtime controls in place, matched to the risk of each process, so its activity can be monitored or stopped.

As those controls prove themselves in practice, agents' authority can be extended, so deployment and governance advance together. Most organisations are already taking this route, with 56% continuing to deploy within agreed risk limits while their governance adapts.

75%

expect agents to act **without human approval of every action by 2027**

Accountability and Liability Remain Unsettled

Assigning AI risk to the board does not resolve who is accountable when an agent causes an adverse outcome. No single role is named by more than 17% of respondents.

The most frequently cited answer is the CIO or CTO (17%), followed by shared C-suite accountability (14.5%), the CRO or CCO (14%) and the Chief AI Officer (13%). Others identify the deploying team (11%), AI vendor (10%), business-line leader (10%) or CEO (9%).



Views on legal liability are similarly divided. One in five say it should be shared between the deployer and vendor

Views on legal liability are similarly divided. One in five (20%) say it should be shared between the deployer and vendor. Others say the issue is too complex for a general answer (19%) or that the necessary regulatory framework does not yet adequately exist (19%). A further 15% identify an individual accountable executive, while 14% assign liability to the AI vendor.

Executives expect several capabilities to be important as agent autonomy grows: formal vendor-liability frameworks (33%), defined thresholds for human escalation (32%), real-time monitoring and audit trails (30%) and the ability to pause agent activity across the enterprise (29%).

Governing AI Beyond the Enterprise

AI governance does not stop at the organisation's boundaries. **More than half of executives (54%) acknowledge gaps or limitations in how their governance covers AI used by vendors, partners and supply-chain participants.**

Where controls exist, they more commonly rely on contractual provisions (43%) and audit rights (39%). Continuous monitoring of third-party AI is reported by 21%, while 13% require independent assurance or certification.

As organisations adopt AI tools and agents provided by others, governance needs to account for what those systems can access and do, not only who supplied them. Contracts and audits can form part of oversight, but they may need to be complemented by runtime controls at the execution level and ongoing visibility into system behaviour and access.

The overall agentic system view is essential for governance. Risks are introduced not only from the behaviour of individual agents, but also from the interactions and data exchange between agents.



Case Study

Bertolotti S.p.A., a leader in advanced logistics and industrial automation, wanted to accelerate AI adoption while maintaining full control over sensitive data, proprietary know-how, and intellectual property.

To meet these requirements, the company partnered with IBM and Uno Informatica to develop *Gina On-Air*, an on-premises AI platform built on IBM watsonx.ai and watsonx Orchestrate, powered by IBM Fusion and Red Hat OpenShift. The solution enables secure, scalable AI integrated with existing

systems, with a first agent that automates the analysis of complex Request for Information (RFI) and Request for Quotation (RFQ) documents and supports commercial decision-making.

Gina On-Air lays the foundation for a broader AI orchestration platform that will progressively expand across engineering, procurement, administration, and HR, helping reduce manual tasks, increase productivity, enhance specialized expertise, and strengthen Bertolotti's competitive position through secure and scalable innovation.

Regional Perspectives

The pressures shaping AI governance vary across markets.

In the European Union, 41% of executives surveyed said their organisations are accelerating governance investment ahead of the EU AI Act's high-risk obligations, expected to come into force in December 2027.

In the UK, sector regulators are the leading reported driver of governance investment, cited by 27%. The UK is also the most advanced market for agentic AI in the study: 57% are running it in production, compared with 35% overall, and 24% have it embedded across the enterprise, twice the 12% average. 95% of UK respondents say their organisation is prepared from a governance perspective.

In Saudi Arabia and the UAE, national AI strategy requirements are the leading reported driver of governance investment, cited by 61%.

While the drivers differ, the underlying challenge is shared. As AI agents operate across borders and systems, organisations will need governance that applies the same controls wherever agents act, while meeting the specific requirements of each market.



Case Study

etisalat

As an AI-powered enabling telco, etisalat is expanding the use of agentic AI across its operations to improve efficiency while maintaining high standards of governance and accountability. In collaboration with IBM and Gulf Business Machines, etisalat developed an agentic AI solution based on IBM watsonx Orchestrate, integrated with IBM OpenPages and the broader watsonx portfolio. The solution enables employees and auditors access and interpret legal, regulatory and compliance information through clear,

traceable responses aligned with the company's governance requirements.

An eight-week proof of concept demonstrated how agentic AI could streamline compliance tasks, support faster and more consistent policy interpretation, and provide 24/7 self-service access. The initiative reinforces etisalat's focus on putting agentic AI into practice, combining operational efficiency with enterprise controls, explainability and human oversight to support responsible AI adoption at scale.

Governance and Business Outcomes

Most organisations have a formal framework for measuring AI return on investment (ROI): 87% say they have one. However, only 40% apply it across the enterprise. Among organisations with an AI ROI framework, 68% include governance costs in their measurement.

This is significant because governance is part of the conditions for responsibly deploying AI, rather than a separate consideration from business value.

Organisations need to understand both the benefits they expect from AI and the resources required to manage its risks.

Executives also identify several governance-related barriers to realising AI's returns, with no single issue dominating, as explored in the final section of this report. Addressing them will require organisations to connect governance with investment decisions and operating models.

68%

include governance costs in their measurement of AI ROI

Case Study

A1 Telekom Austria AG

As organisations increasingly adopt Agentic AI, governance, transparency and data sovereignty are becoming essential requirements. A1 Austria addresses these challenges through its European cloud platform Exoscale combined with IBM watsonx. The solution enables organisations to develop, deploy and govern AI applications while maintaining control

over their data, models and infrastructure. Integrated capabilities for AI governance, compliance and lifecycle management support transparency and help enterprises address evolving regulatory requirements. By combining sovereign European cloud infrastructure with enterprise-grade AI capabilities, A1 Austria supports organisations to scale agentic AI responsibly, securely and with confidence.

Five Priorities for Governing Autonomous AI

Governance is one of the keys to unlocking business value from AI. Success in the next phase of agentic AI adoption will depend not only on what agents can do, but on whether organisations can understand and govern their actions.

There are several practical questions that can help guide that work: Who are the agents in the system? What are they allowed to do? What did they do? Can their limits be enforced? Which runtime controls are put in place in the agentic AI system?

To establish effective governance for agentic AI, enterprise leaders should focus on five key priorities:

1. Know every agent you run

Governance starts with visibility. Organisations need a clear picture of the agents operating across the enterprise, including those introduced by individual teams or embedded in suppliers' systems without central oversight. With more than half of respondents acknowledging gaps in governing third-party AI, this has to extend beyond the agents an organisation builds itself.

2. Give each agent a clear identity and only the authority it needs

Every agent should have a defined identity, owner and scope: which systems it can access, which decisions it can make and when it must hand back to a person. Autonomy should be extended deliberately, based on the value it creates and the risk it introduces, and judged on what an agent can do rather than who supplied it.



3. Extend governance across end-to-end processes


A single request can pass through several agents, built on different applications, before it is complete. Governance applied within each application in isolation can leave gaps at the points where agents hand work to one another. Organisations should apply consistent policies across the whole process and the full agent lifecycle, define thresholds for human escalation, and ensure controls can limit or pause activity at any step.

4. Invest in governance processes at the level of individual AI agents and systems

Board ownership, committees and policies give AI governance its foundation, but they need to translate into decisions about each individual agent and how they interact with other agents and tools. The right controls depend on the business context of each AI system, so this starts with identifying the risks and legal obligations attached to a given system. From there, organisations can define the controls they need at design time and at runtime, verify through monitoring and audit trails that those controls have been applied, and ensure the agent does not expand beyond the scope it was assessed and approved for.

5. Assign accountability and measure the risk at board level

Board ownership of AI risk is near universal, but accountability for individual agents that are operational is less clear. When deployed in an enterprise, agents should be overseen by a professional who understands each agent's scope and can intervene if anything goes astray. Ultimate accountability may sit higher in the organisation, so lines of responsibility need to be clear, rolling up from that first-line oversight to executive leadership, with governance risk quantified for regular board review.



The Next 18 Months: Barriers to AI Returns

AI governance has reached the boardroom, but board-level ownership alone is not enough. As agents take on more autonomous work, organisations will need to make accountability clearer, extend oversight to third parties and ensure controls can operate while systems are active.

The research points to a gap between the autonomy organisations expect and the governance they say is ready to support it. Closing that gap is an immediate operational priority. With clear ownership, defined authority, enforceable boundaries and traceable actions, organisations can work towards extending AI's role with greater confidence.

According to the research, executives are divided on the most significant governance-related barrier to achieving AI return on

investment over the next 18 months.

Respondents most often cite frameworks designed for an earlier era of AI (16%) and data-sovereignty constraints (16%) as the largest barriers to AI ROI. Regulatory uncertainty came a close third (15%), followed by inadequate liability frameworks (14%). Measuring and attributing AI's return on investment is another concern, cited by around 13% of respondents.

Governance is about more than managing risks, it is crucial for enterprise AI value creation. To harness the power of agentic AI, organisations must invest in the right technology strategy, with clear accountability, operational control and direction from the boardroom.

Case Study

Celerity

IBM watsonx.governance provides lifecycle governance, risk and compliance workflows, asset documentation, evaluation and monitoring, while the agentic control plane in watsonx Orchestrate adds visibility, policy enforcement, guardrails, auditability and identity controls for enterprise agents.

Celerity's WatchPoint extends these capabilities by correlating activity across IBM and non-IBM AI services, identifying unauthorised use, and linking operational evidence to accountable owners, policies

and assurance processes. Using IBM Bob, Celerity accelerated WatchPoint from prototype to production readiness. In its first 30 days, WatchPoint brought more than 712,000 AI interactions into a governed organisational view and helped consolidate activity from 42 unauthorised services onto five approved enterprise platforms.

The resulting cross-platform evidence supported Celerity's ISO/IEC 42001 certification programme, while watsonx.governance and watsonx Orchestrate continued to provide lifecycle, risk and agent-control governance capabilities.

Study Methodology

The research was conducted by Censuswide on behalf of IBM among 3,221 executive management respondents, including C-level, at organisations with 1,000 or more employees in France, Germany, Italy, the Netherlands, Saudi Arabia, Spain, the UAE and the UK. The data was collected between 14 and 24 September 2026.

Further Reading

The agentic AI risk and governance topics covered in this report reflect the findings of the survey. To learn more about the wider risk landscape around AI, please visit [**ibm.com/trust/responsible-technology**](https://ibm.com/trust/responsible-technology).

