From blueprint to reality

Execute effective AI governance in a volatile landscape

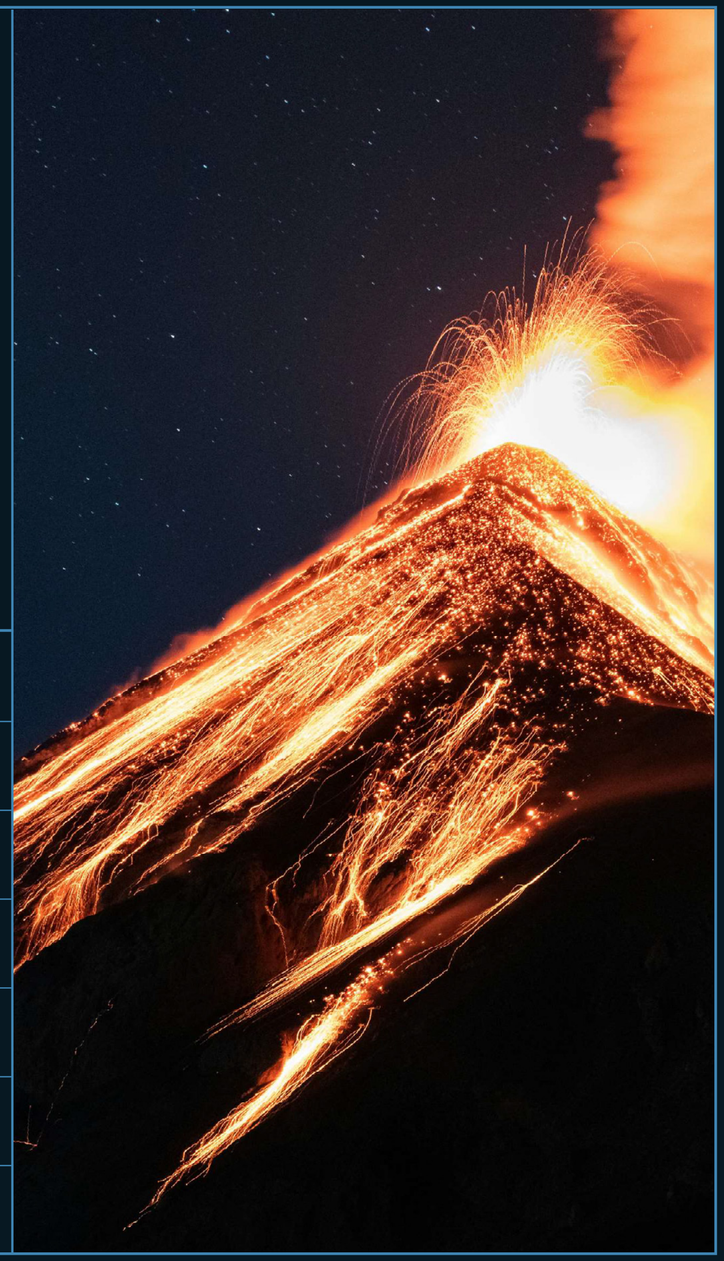

AUDITBOARD.COM

Table of contents

Executive summary: AI governance at a crossroads

With 82% of organizations reporting moderate to extensive deployment of AI tools across functions, AI adoption is no longer speculative. It's operational. Organizations across sectors are racing to integrate generative and machine learning tools into their core business processes, seeking productivity gains and competitive advantages. But this momentum has triggered a parallel challenge: managing the associated risks.

Business leaders are now tasked with building robust AI governance programs at a time when the technology, threats, and regulatory expectations are all evolving rapidly, and without a consistent benchmark for what "robust" governance really looks like. The result is a growing disconnect between policy creation and policy execution.

Nearly all organizations surveyed report awareness of regulatory developments and express deep concern about AI-related risks. Yet implementation lags behind. Many have drafted policies, but few have embedded AI governance into their organizations' operational fabric.

This "policy-practice gap" is emerging as a new risk frontier, one rooted not in ignorance but in executional uncertainty, cultural fragmentation, and misaligned ownership.

This report draws on survey data from 412 GRC and audit professionals to examine:

- Why this gap exists despite high awareness and urgency
- How cultural and structural factors are greater barriers than technology
- Where organizations are overconfident and underprepared
- What practical steps can help embed governance into day-to-day operations

Success will require more than compliance checklists. It will demand cross-functional accountability, clear governance ownership, and a commitment to a risk-aware culture at every level of decision-making.

WHAT GRC LEADERS NEED TO KNOW:

- Most orgs have policies; few have operational controls
- Risk and compliance are often sidelined in governance design
- Shadow AI and third-party risks are underestimated and under-managed
- Cross-functional governance must go beyond checklists

Section 1: AI's explosive growth, risks, and regulatory imperatives

AI is advancing fast; **more than 75% of organizations report using or planning to use multiple forms of AI, including generative, predictive, and classification systems, making governance not just urgent, but increasingly complex.** Yet this wave of adoption is outpacing the systems designed to manage it. As a result, many organizations now face a critical mismatch between the speed of innovation and the maturity of governance.

In our survey, 86% of respondents said their organization is aware of AI regulations that are coming or already in force. Many are familiar with major frameworks such as the NIST AI Risk Management Framework, the EU AI Act , and national guidelines like Canada's Directive on Automated Decision-Making. This awareness suggests that governance is on the radar, but awareness does not equal preparedness.

Despite high levels of concern, with over 80% of respondents saying their organizations are "very" or "extremely" concerned about AI risks, implementation is still lagging. AI systems are being deployed faster than oversight structures can keep up, leading to ad hoc governance, uneven accountability, and increased exposure to legal, ethical, and operational failures.

This disconnect is emerging as a global challenge, shaped in part by uneven regulatory landscapes. In the European Union, the passage of the AI Act marks a significant shift, introducing binding obligations based on risk tiers and requiring documentation, oversight, and enforcement mechanisms. In contrast, the United States has emphasized voluntary frameworks like NIST's, with sector-specific oversight evolving at a slower pace. The UK and Canada have taken a principles-based approach, prioritizing transparency and fairness through guidelines rather than laws.

Amid this regulatory patchwork, many organizations are gravitating toward the NIST AI RMF as a de facto standard. Though non-mandatory, 49% of surveyed organizations are aligning with it, not because they're required to, but because they see strategic upside. The NIST framework helps companies prepare for likely regulation, signals responsibility to customers and investors, and provides internal clarity around roles and processes. For many, it functions as both a risk shield and a reputational asset.

As some jurisdictions resist regulatory mandates and others accelerate them, one thing is clear: **Governance is no longer optional.** In a fragmented policy environment, internal governance has become a business-critical function. Organizations that treat governance as a core capability, not a compliance box-checking exercise, will be better positioned to manage risk, build trust, and respond to a rapidly evolving regulatory landscape.

Image credit: Peter Olexa

Section 2: Making AI governance work in practice

Despite high awareness and growing concern, most organizations remain early in their AI governance journeys. Many have policies in place or in development, but few have made the leap from documentation to disciplined execution. In our survey, **only 25% of organizations said they have a fully implemented AI governance program.** While that figure will likely rise in the coming year, it reveals a striking lag between intent and action.

In the meantime, most efforts are focused on policy drafting, principles development, and internal messaging around responsible AI use. These steps are important, but insufficient on their own. Without integration into business workflows, technical environments, and operational routines, even the best-written policies will remain theoretical.

The gap becomes more apparent when we look at specific governance components. While organizations are investing in complex efforts like AI usage monitoring (45%), risk assessments (44%), and third-party model evaluations (40%), far fewer have implemented foundational practices. Only 28% have usage logging, 25% maintain model documentation, and just 23% enforce access controls for AI systems. **Many are trying to solve the most difficult parts of governance first without a clear foundation to build on.**

What's driving this gap? In part, it's a function of competing pressures. Business units want to move fast with AI to capture efficiency gains and market advantage. Risk teams want to implement controls and slow things down. In the absence of coordinated structures, policy creation becomes the lowest-friction path to showing progress, even if that progress hasn't yet reached production environments.

The result is a governance maturity curve with a steep slope: early enthusiasm and documentation, followed by a harder, slower climb toward real-world accountability. Organizations understand the risks; they've written them into their policies. But turning those policies into daily practice remains a work in progress.

The practice-policy gap between intent and execution

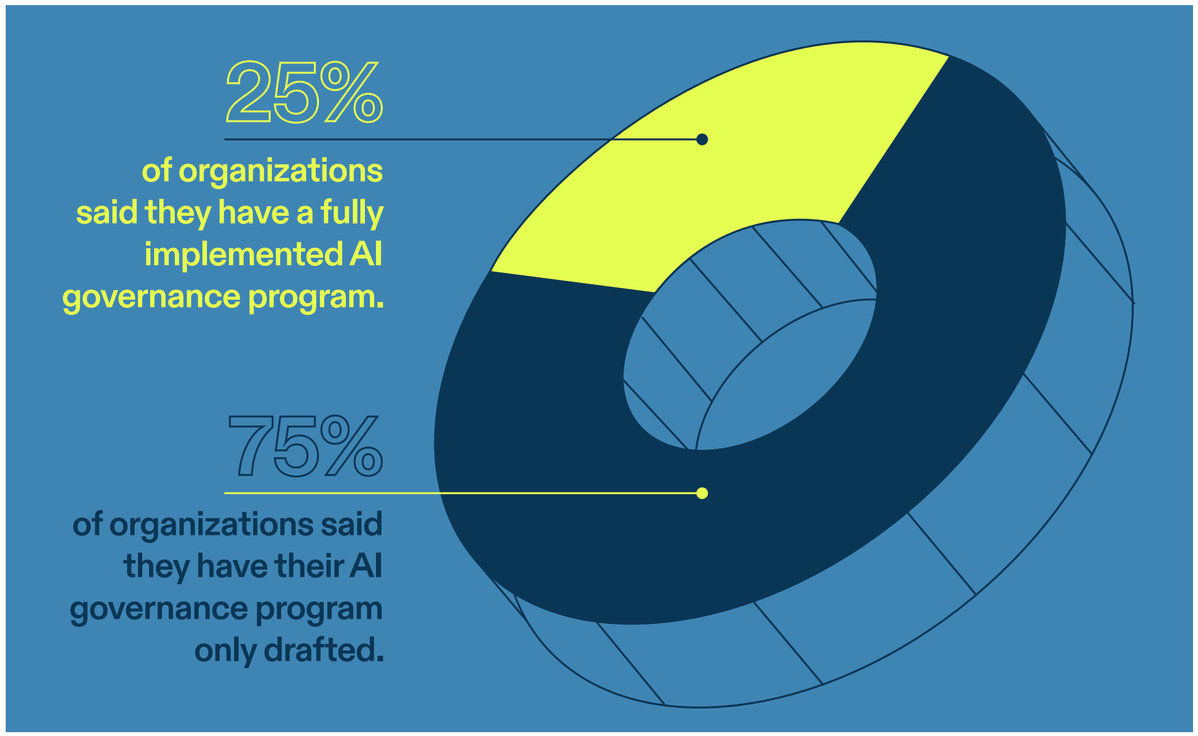

Section 3: Culture and execution: The real barriers to AI governance

If governance is lagging, it's not because organizations lack awareness or even intent. It's because they're confronting barriers that are far more cultural and structural than technical. In our survey, respondents identified the leading obstacles to AI governance as lack of clear ownership, insufficient internal expertise, and resource constraints. **Fewer than 15% said the main problem was a lack of tools.**

Top barriers to implementing AI governance

Specifically, when asked to identify their top barriers to implementing AI governance:

44%	A LACK OF CLEAR OWNERSHIP
39%	LACK OF INTERNAL EXPERTISE
34%	LIMITED RESOURCES

Source: AuditBoard, June 2025 flash poll of 412 information security, compliance, and risk professionals

This distinction matters. Most organizations are not struggling to find dashboards or compliance software; they're struggling to determine who's accountable, how teams should coordinate, and what workflows need to change. The issue is less about capability and more about clarity.

This is why many governance efforts stall even after policies are drafted. Policy tells the organization what should happen. Culture and structure determine whether it happens. And until organizations address the cultural gaps—unclear roles, lack of collaboration, uneven accountability—the policy-practice gap will persist.

Section 4: Fragmented ownership, fragmented oversight

One of the most persistent challenges in AI governance is not whether it's on the executive radar—it is—but rather how responsibility for it is distributed across the organization. While nearly all organizations in our survey (96%) report some level of board or executive engagement with AI governance, this top-down interest has not translated into clear operational accountability.

Technical leaders often focus on innovation, performance, and scalability. Compliance, ethics, and risk mitigation may be part of the conversation, but they're rarely at the center of governance design or enforcement. And without clear accountability for integrating governance across business lines, policies often remain abstract or siloed.

The result is a governance structure that appears coherent on paper, backed by policies, executive sponsorship, and formal committees, but often lacks the operational clarity to be effective in practice. Oversight becomes fragmented not just in terms of role ownership, but also in how risks are surfaced, prioritized, and addressed across the AI lifecycle.

Without clearly defined roles, formalized handoffs, and coordinated processes between technical and risk functions, organizations are left with what might be described as **"distributed responsibility without distributed accountability."** And in a field as fast-moving and high-stakes as AI, that's a serious structural vulnerability.

This structure creates a fundamental misalignment between where AI is being built and where it should be governed.

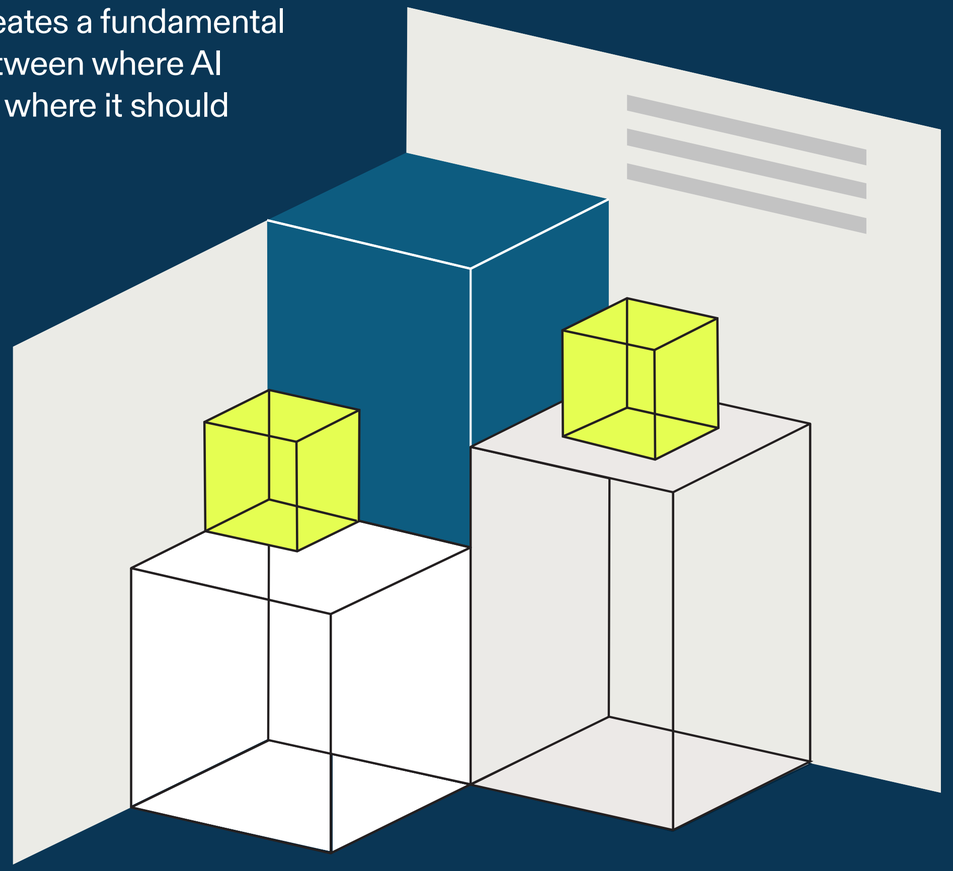

Section 5: Confidence ≠ control: The visibility gap

Across our survey, organizations expressed a high degree of confidence in their ability to oversee both third-party AI systems and unauthorized, employee-initiated tools, often referred to as “shadow AI.” Ninety-two percent of respondents said they are confident in their visibility into third-party AI use, and nearly as many (**90%**) claimed strong oversight of shadow AI within their environments. On the surface, this suggests a mature governance posture.

But dig deeper, and that confidence appears misplaced. Just **67%** of organizations report conducting formal, AI-specific risk assessments for third-party models or vendors. That leaves roughly one in three firms relying on external AI systems without a clear understanding of the risks they may pose. The visibility into shadow AI is even harder to verify. As generative tools become embedded into everyday workflows, from marketing to coding to operations, employees are increasingly adopting technologies outside the scope of official procurement, IT, or compliance processes. Many of these tools process sensitive data, make operational decisions, or generate customer-facing outputs, all without formal oversight.

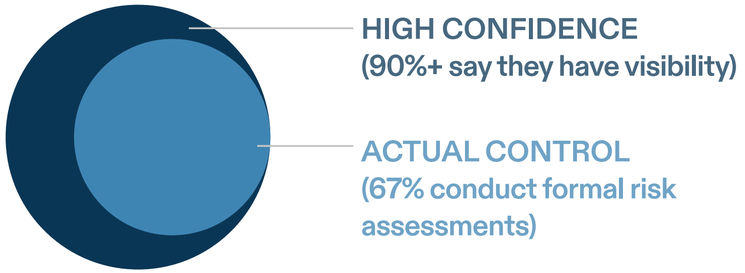

Overconfidence, in this context, becomes a risk in itself. When companies assume they have control, they’re less likely to invest in proactive auditing, centralized model inventories, or employee education. And when vulnerabilities surface, they’re often caught off guard, leading to downstream consequences.

Those consequences can be severe. Compliance violations tied to unauthorized AI use can result in legal penalties, especially as regulations like the EU AI Act and national privacy laws tighten. Privacy breaches may arise if unvetted tools access personal or regulated data. Integration with poorly secured third-party models can create new entry points for cyberattacks. Perhaps most damaging of all, public trust can erode rapidly if flawed or biased AI outputs go unmonitored until it’s too late.

In short, visibility is not the same as control, and assuming otherwise can mask critical weaknesses in an organization’s risk posture. To close this gap, organizations need more than high-level oversight. They need formal processes to identify, classify, monitor, and manage every AI system they touch, whether built internally, sourced externally, or adopted unofficially by employees.

TOP 5 GOVERNANCE RISKS

1. SHADOW AI TOOLS BYPASSING OVERSIGHT
2. THIRD-PARTY MODELS LACKING FORMAL REVIEW
3. AUTOMATION SCALING WITHOUT STRUCTURE
4. FRAGMENTED ACCOUNTABILITY ACROSS DEPARTMENTS
5. POLICIES NOT EMBEDDED IN DEV / OPS WORKFLOWS

Section 6: Automation ambition vs. foundational gaps

As the complexity and scale of AI systems continue to grow, many organizations are turning to automation in hopes of closing their governance gaps. The appeal is obvious: Automation promises efficiency, consistency, and scalability, particularly in environments where manual oversight simply can't keep up. But our findings reveal a potentially risky pattern. **Organizations are directing their automation efforts toward the most advanced and difficult governance tasks, while foundational controls remain immature or incomplete.**

ORGANIZATIONS NEED TO BUILD AUTOMATION ON SOLID FOUNDATIONS

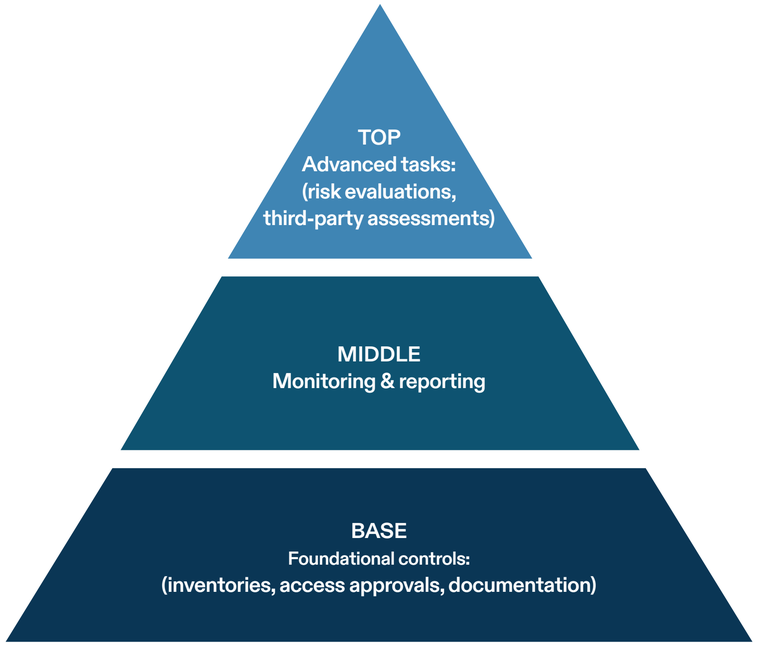

In our survey, a majority (51%) of respondents identified AI usage monitoring, 47% selected third-party AI assessments, and 45% selected risk evaluations as their top priorities for automation. These are sophisticated processes that require accurate data inputs, strong accountability frameworks, and well-defined governance policies. And yet, many of the building blocks that support such efforts—things like model inventories, usage logging, and approval workflows—are either missing or inconsistently applied. And when it comes to fairness and transparency, **59%** of organizations still rely on human review rather than technical solutions, further emphasizing the gap between ambition and operational maturity.

This mismatch between ambition and operational readiness creates a significant governance vulnerability. Automation is being used to address high-risk activities before routine controls are in place to support them. Rather than building upward from strong foundational practices, many organizations are trying to scale governance from the top down. It's an approach that risks embedding inconsistency, rather than eliminating it.

What emerges is a kind of governance illusion: Automation gives the appearance of control, but without foundational processes and clear ownership, it may simply replicate gaps at scale.

Governance, after all, is not just about surveillance; it's about structure. Without inventories to know what AI is being used, or workflows to determine how it's approved, even the best automation cannot enforce rules that don't yet exist.

Strategic automation depends on maturity. Organizations that try to leapfrog foundational governance steps risk building brittle systems that crack under regulatory scrutiny or operational pressure. A more effective approach starts with codifying the basics: documenting AI assets, formalizing review procedures, assigning responsibility, and then layering automation on top to scale what already works.

Section 7: From frameworks to execution: What success looks like

For many organizations, **success in AI governance has so far been measured by the presence of policies:** codes of conduct, ethical AI guidelines, or risk principles. These artifacts are important, but they are not outcomes. As the use of AI expands and governance expectations rise, the definition of success is beginning to shift: from drafting high-level principles to embedding enforceable practices into the core of how AI systems are developed, deployed, and managed.

This evolution is happening unevenly. In the next 12 months, most organizations (52%) plan to continue prioritizing policy development. Ethics frameworks, risk principles, and compliance guidance still top the list of AI governance initiatives. Enforcement, on the other hand, remains a future objective. Governance continues to trail behind adoption, with many companies still in the process of defining internal rules even as advanced AI systems go live across their business units.

This "build-as-you-go" approach reflects the pressure many teams are under. They are moving fast to meet innovation goals, while trying to retrofit guardrails in parallel. But as the policy layer stabilizes, the next challenge will be turning those principles into tangible operational practices.

In our research, leading organizations are beginning to define governance success using more concrete metrics. These include measures like the completeness of their AI system inventory, the percentage of AI use cases undergoing formal risk assessment, the frequency of policy violations or exceptions, and the average time to detect and respond to unauthorized AI use. Others are tracking how much of their third-party AI ecosystem is covered by risk evaluations, or how closely their practices align with external frameworks like the NIST AI RMF.

These are not just compliance indicators; they're signals of operational maturity. An organization that can maintain a current inventory of its AI systems, enforce approval workflows, and track exceptions in real time is not just compliant. It's in control. And in a field where reputational risk and regulatory expectations are evolving rapidly, control is everything.

At the heart of this shift is a mindset change. AI governance cannot remain an annual policy exercise or a reactive audit function. It must become part of daily operations, baked into how AI is evaluated, approved, and monitored at every stage of its lifecycle.

The next frontier of AI governance is operational accountability, not just policy authorship.

Section 8: Recommendations - Building a living governance program

AI governance cannot be treated as a one-time compliance project. Based on the patterns uncovered in our research, five strategic actions stand out for organizations seeking to move from policy drafting to durable, scalable execution.

1. Translate policy into practice.

Move beyond ethical principles and into execution. Define how policies apply to real-world scenarios: which teams review AI use cases, how model performance is monitored, and what happens when issues arise. Embed governance into daily decisions, not just documents.

2. Build and empower cross-functional teams

AI governance isn't owned by one function. Risk, compliance, product, legal, security, and engineering all need a seat at the table. Establish cross-functional councils with clearly defined responsibilities and decision-making authority to ensure consistent execution.

3. Automate strategically, not prematurely

Automation is powerful, but only when built on solid foundations. Focus first on core controls like AI inventories, access approvals, and documentation standards. Scaling without structure risks embedding gaps rather than closing them.

4. Train and communicate continuously

Even the best frameworks fail without user understanding. Roll out training tailored by function and seniority. Communicate policies clearly and often, with internal reporting and visible expectations to build a culture of responsible AI use.

5. Stay agile and adaptive

The regulatory and technology landscape is evolving fast. Governance must too. Shift from annual reviews to continuous updates, with teams structured to respond to new tools, risks, and regulatory changes as they emerge.

Together, these recommendations form the foundation for a governance approach that is not only compliant but sustainable. By shifting from reactive policies to proactive systems, organizations can embed governance into the DNA of their AI strategies, reducing risk, building trust, and enabling innovation with confidence.

AI governance is a team sport

AI governance is no longer a theoretical concern; it's a strategic imperative. As organizations integrate AI into critical operations, they must also grapple with the risks that come with speed, complexity, and opacity. Our research shows that the challenge today isn't awareness. Most organizations recognize that regulation is coming, that reputational risks are real, and that governance matters. **The problem lies in execution.**

From fragmented ownership to overconfidence in visibility, the obstacles to effective governance are deeply cultural and structural. Policies exist, but processes are missing. Automation is ambitious, but uneven. Teams are engaged, but not aligned. These disconnects are creating governance programs that look complete on the surface but are still fragile underneath.

To move forward, organizations must treat AI governance as a living, shared responsibility, not a siloed compliance task. It requires collaboration across functions, clear lines of accountability, and controls that are designed to evolve alongside the technology. Success will not come from one team or one tool, but from cross-functional alignment, embedded practices, and a culture of responsible innovation.

In the years ahead, the organizations that outperform will not simply be those who move the fastest with AI. They will be the ones who govern it best, with transparency, resilience, and trust built into every layer of how AI is developed and deployed.

Section 9: The business case for AI governance

Reduced audit burden
Faster evidence generation for controls and oversight, aligned with frameworks like NIST AI RMF and the EU AI act.

Early risk detection
Catch potential issues before deployment, minimizing downstream risk and brand damage.

Improved model transparency
Strengthens trust with customers, partners, and regulators by explaining how AI decisions are made.

Standardized oversight process
Makes governance scalable across business units, reducing ad hoc or siloed approaches.

Ethical alignment
Signals responsibility and values, building goodwill with employees, customers, and the public.

Clear model accountability
Ensures each system has a responsible owner, limiting operation, legal, and reputational exposure.

Automated compliance workflows
Reduces manual tracking and accelerates readiness for internal reviews and external audits.

Centralized model inventory
Enables visibilty into where AI is used, supports prioritization and portfolio-level risk management.

Bias monitoring & mitigation
Helps surface and correct fairness issues before they impact people or trigger scrutiny.

Controlled change management
Maintains stability and traceability across model updates, versions, and decommissioning.

Section 10: Appendix

- **RESEARCH METHODOLOGY**
  The survey included 412 respondents sourced from a leading global online panel provider. They were selected from the panel based on geographic and role-based quotas, as well as screening questions based on role in audit and compliance, decision-making role, company size, and how long they have been in their audit role. All participants were Audit, GRC, or IT decision-makers and purchase influencers working at companies with annual revenue of at least $100 million USD. Selected respondents were further screened based on self-reported audit and compliance knowledge and attentiveness to survey questions.

- **ROLE QUOTAS**
  The survey divided respondents into four broad roles: C-suite 20%, Lead 70%, Manager 10%. Respondents were asked to select which role – from a list of 23 options – most closely described their primary responsibility, even if none were quite right or even if they performed more than one of these roles. Answers were consolidated into those four broad roles.

- **GEOGRAPHIC QUOTAS**
  The survey included respondents from the U.S., Canada, Germany, and the UK.

- **INDUSTRY**
  Although no industry-level quotas were deployed, we monitored the data to ensure that no single industry was overrepresented in the data. The final breakdown of respondents by industry is as follows: Financial Services 12%, Retail / Ecommerce 12%, Industrial and Manufacturing 12%, Energy & Resources 12%, Transportation and Logistics (including supply chain) 12%, Life Sciences (including healthcare and pharmaceuticals) 11%, Insurance 8%, Technology 8%, Business / Professional Services 4%, Education 4%, Government / Public Sector 2%, Telecommunications 2%, and Marketing and Advertising 2%.

- **RESPONDENT SCREENS**
  - Role: All respondents were required to indicate that they were responsible for or had influence in evaluating and/or selecting audit compliance solutions or software for their organization.

  - Company size: All respondents must self-report that their companies have a minimum of 250 employees. All potential respondents from smaller companies were excluded. In total, the survey includes 3% of respondents from companies with 250-499 employees, 12% from companies with 500-999 employees, 50% from companies with 1,000 to 4,999 employees, 23% from companies with 5,000 to 9,999 employees, 8% from companies with 10,000 to 24,999 employees, 3% from companies with 25,000 to 49,999 employees, and 1% from companies with 50,000 or more employees.

  - Time in IT: Respondents must have spent a minimum of 3 years managing, planning, or purchasing compliance and/or cyber risk management software services or infrastructure in order to qualify for the survey. In total, 16% of respondents have spent 3 to 5 years in this role, 56% have spent 6 to 10 years in this role, 26% have spent 11 to 15 years in this role, and 2% have spent 16 years or more in this role.

  - Information level: In our experience, it is possible to have "qualifying respondents" who nevertheless prove to have too little information or knowledge about the space to provide useful data from which to draw insights. We therefore apply an "information" screen to respondents as well. Specifically, we ask whether or not respondents could explain certain terms to their colleagues if asked to do so. In order to qualify for this survey, a respondent must say "yes" to this question for the term "GRC (Governance, Risk, and Compliance)"

  - "Attention" level: It is easy for respondents to speed through surveys or not pay enough attention to provide useful data. We make an effort to exclude these respondents as well, as they provide generally less useful data. In this survey, respondents were screened out for "attention" reasons if they said they could explain the made-up term "CRISM Framework" to a colleague in the same question used for the Information Screen noted above.

- **RESPONDENT SCREENS**
  It is technically impossible and improper to list a margin of error for a survey of this type. The respondents for this sample were drawn from an online panel with an unknown relationship to the total universe, about which we also do not know the true demographics. As such, the exact representativeness of this, or any similarly produced sample, is unknown.

About AuditBoard

AuditBoard's mission is to be the category-defining global platform for connected risk, elevating our customers through innovation. More than 50% of the Fortune 500 trust AuditBoard to transform their audit, risk, and compliance management.

AuditBoard is top-rated by customers on G2, Capterra, and Gartner Peer Insights, and was recently ranked for the sixth year in a row as one of the fastest-growing technology companies in North America by Deloitte.

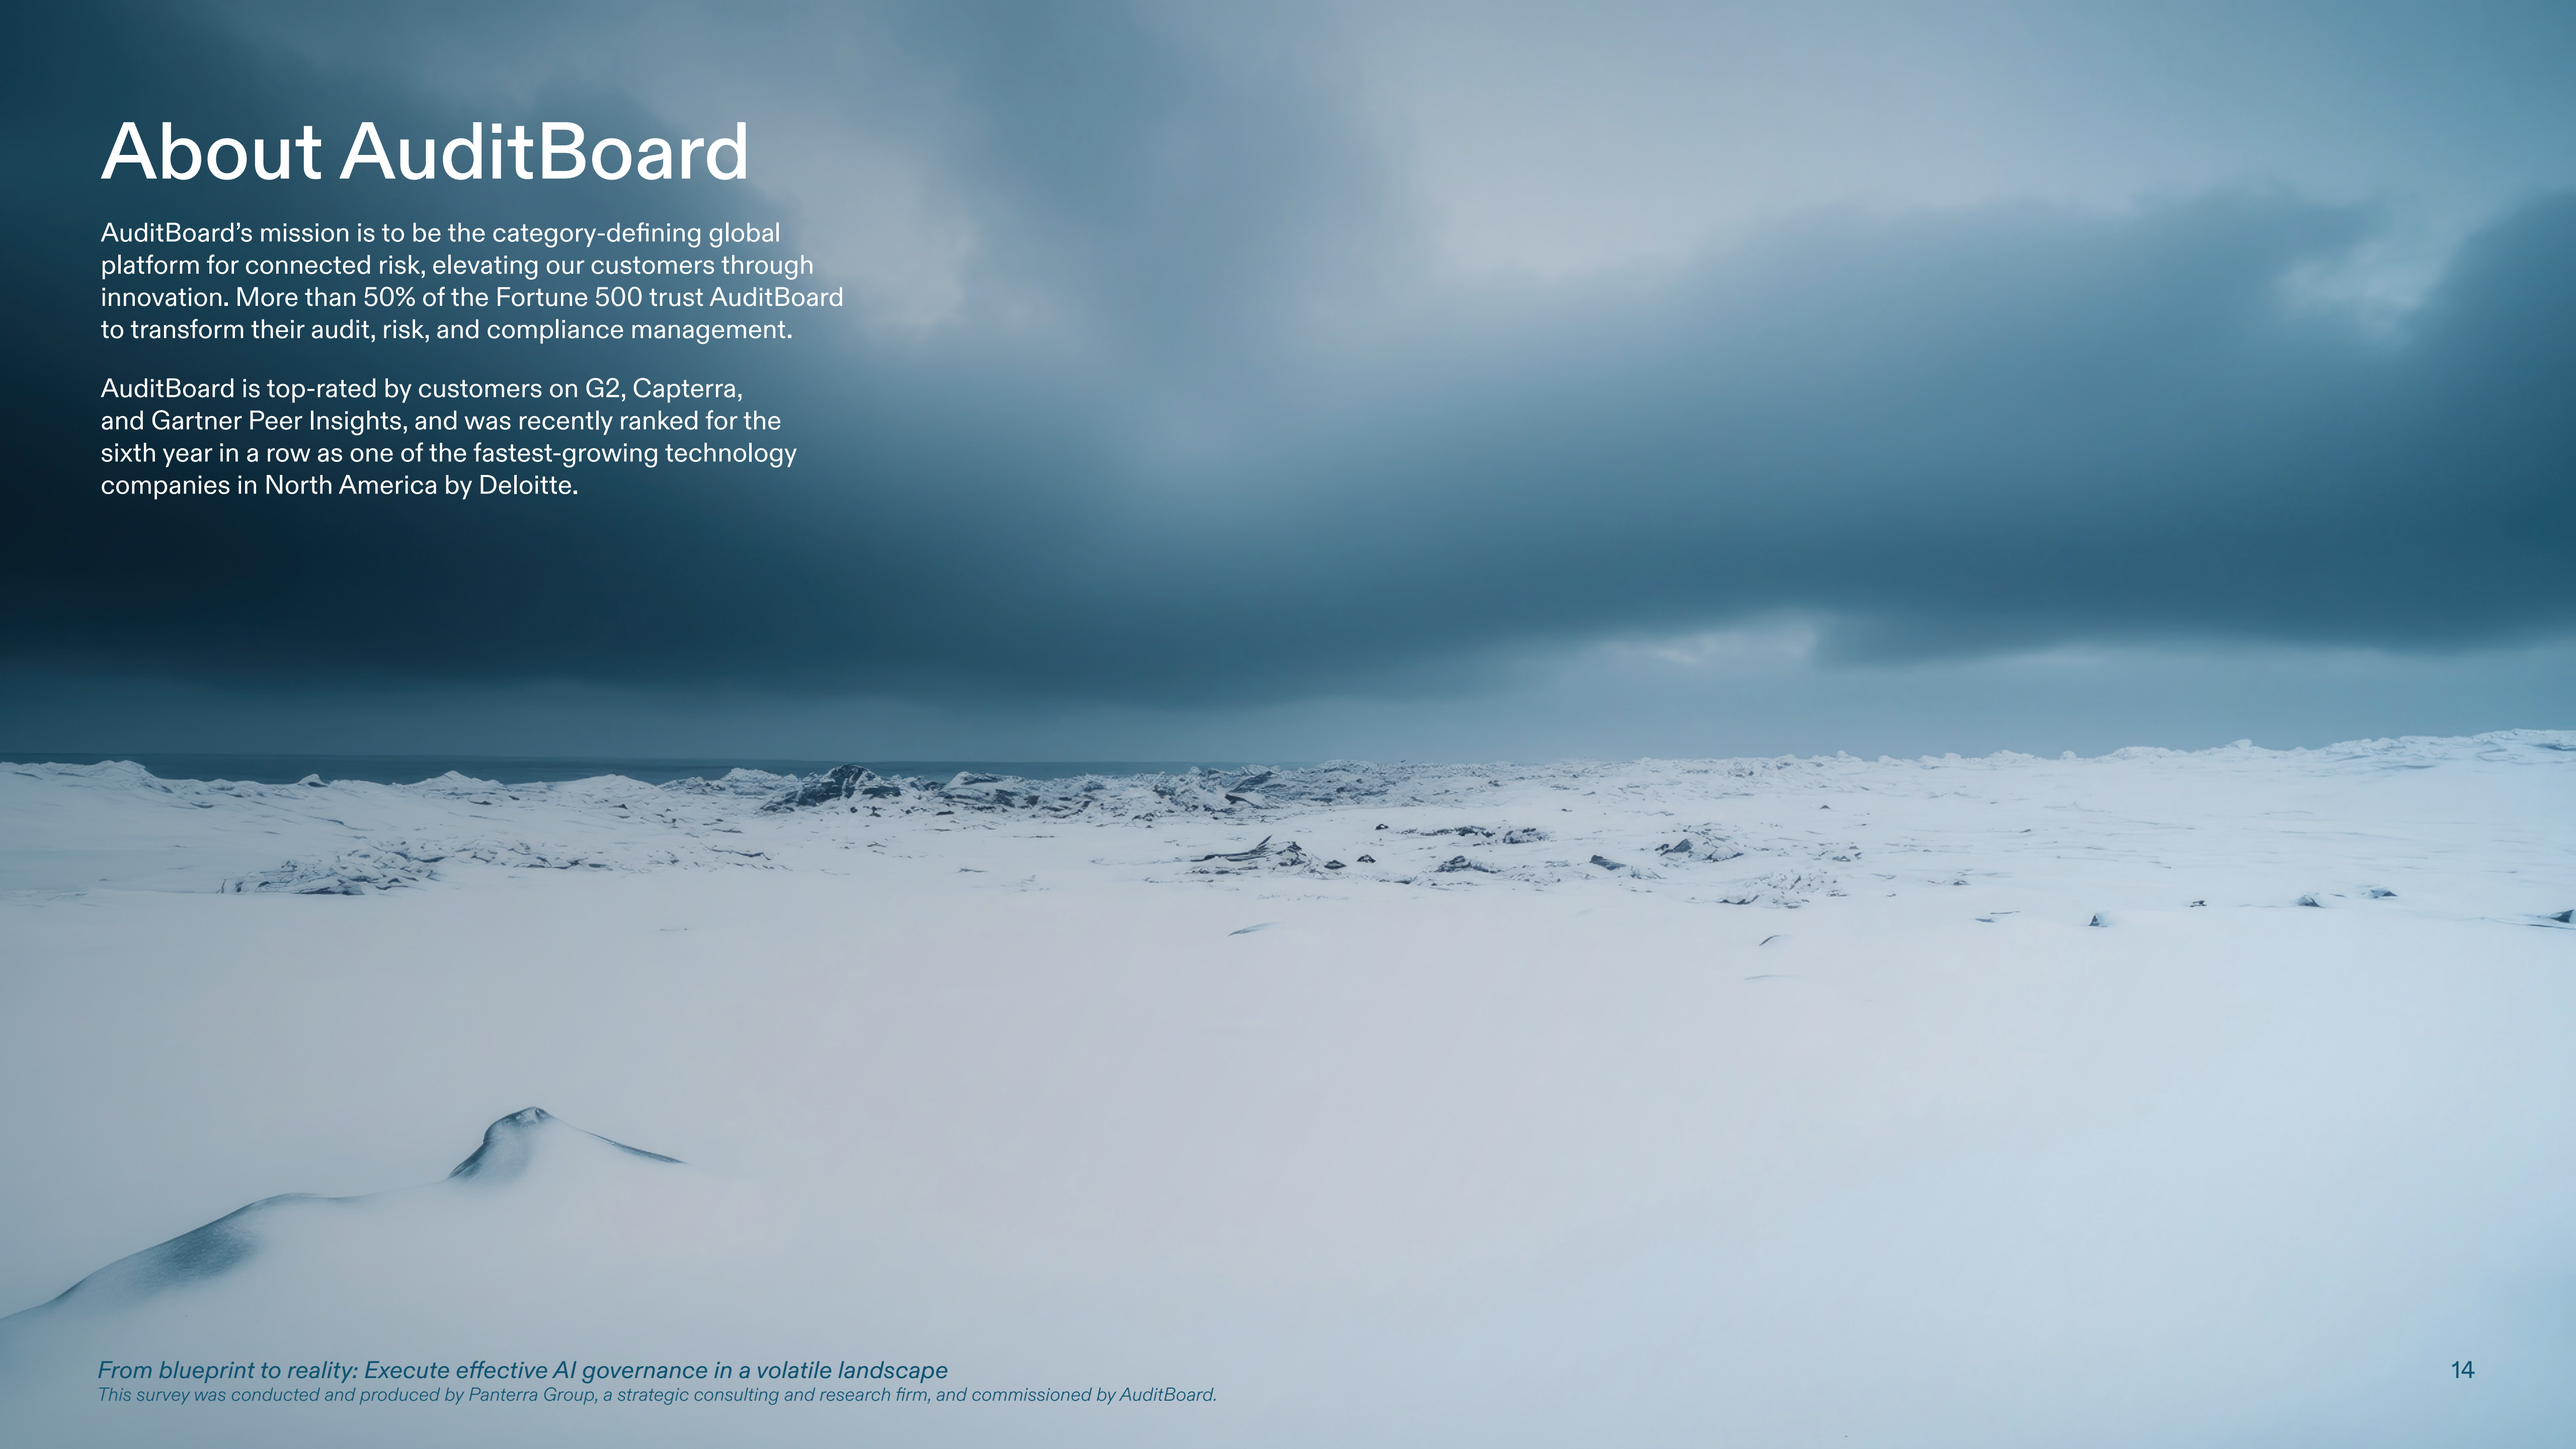


AuditBoard