schellman
Quality. above all.

RESEARCH REPORT

2026 State of AI Governance

Bridging the Gap Between AI Governance Confidence and AI Governance Maturity

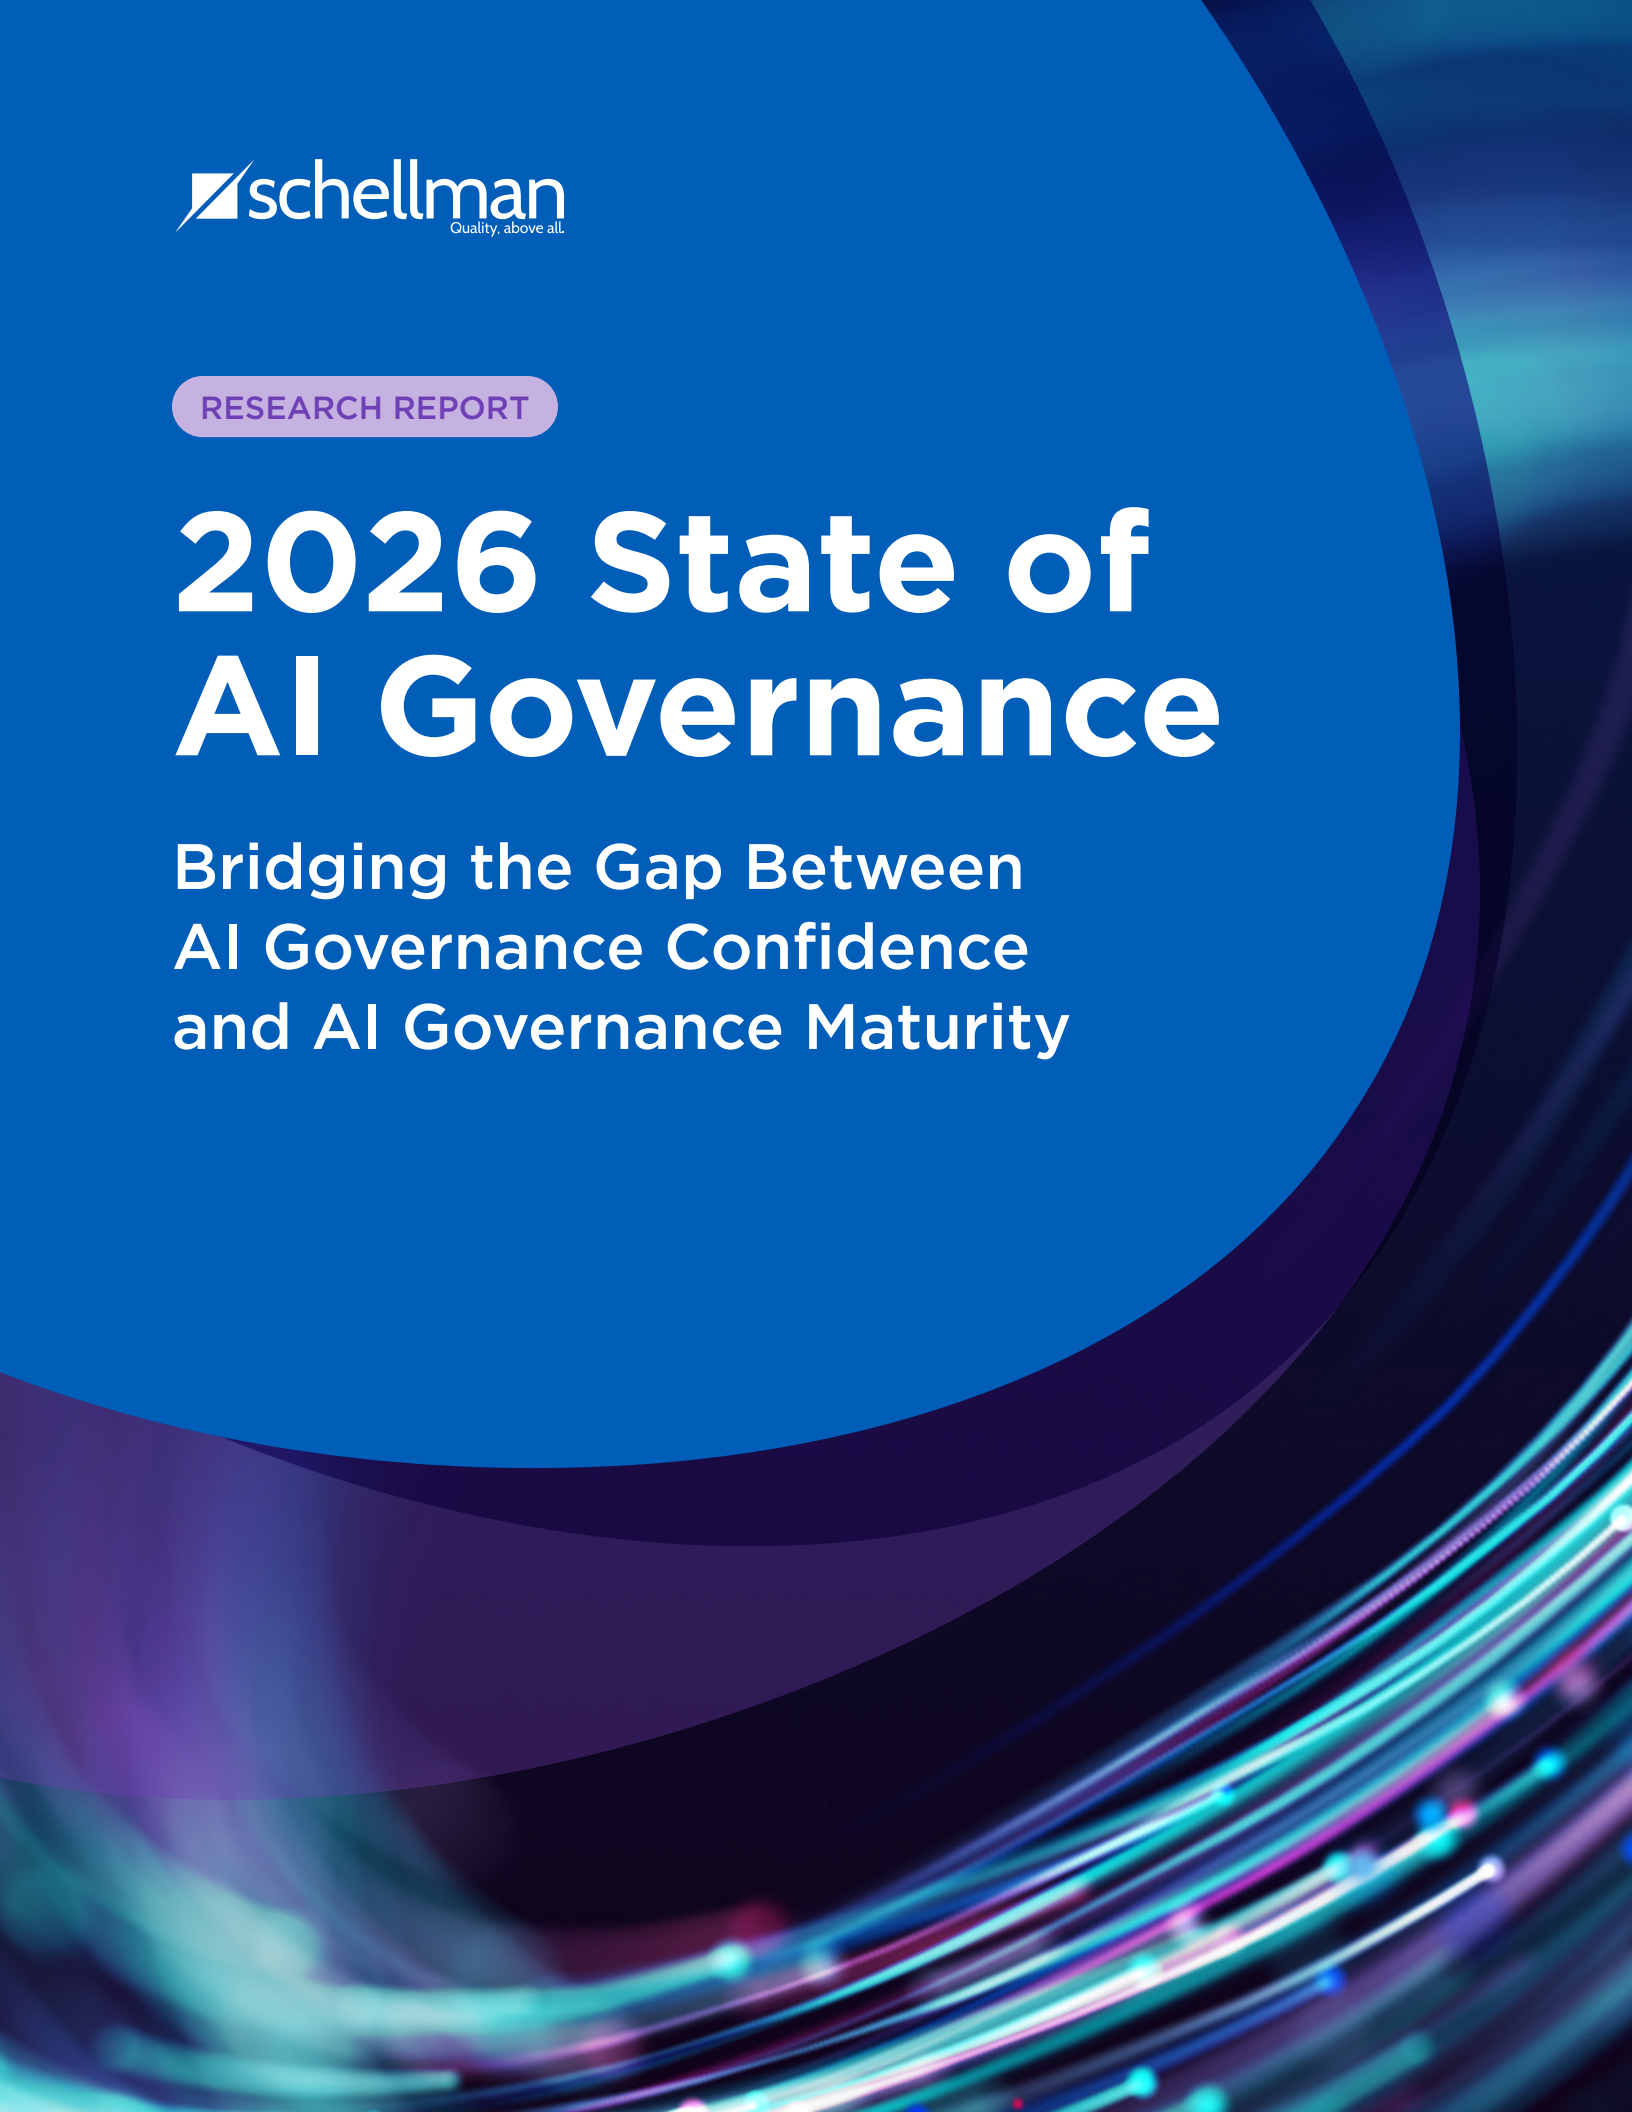


Table of contents

Executive Summary 3

THEME 1 Funding Fuels False Confidence 4

THEME 2 Bridging the Agentic AI Governance Gap 6

THEME 3 The Accountability Vacuum 9

THEME 4 An Impending Regulatory Reckoning 12

THEME 5 The Trust Dividend: AI Governance as a Performance Driver 16

Conclusion: Rethinking the Future of AI Governance 18

Executive Summary

AI adoption has moved past the experimentation phase. Enterprises are scaling AI organization-wide, and the governance programs they originally designed to protect them are struggling to keep up. Schellman's 2026 State of AI Governance survey, fielded across more than 500 U.S.-based professionals at organizations with more than 500 employees and over $100 million in annual revenue, puts data behind a dynamic that most executives sense but few have quantified.

The results expose a conundrum: Most organizations are investing in AI governance, most boards are receiving some level of AI governance reporting, and most leaders say they are confident in their governance capabilities, yet fewer than one in three have achieved operational maturity. The gap between governance activity and governance readiness lies at the heart of this paradox, and most organizations are less prepared to close that gap than they realize.

But it's not just an internal issue. Increasingly, organizations are finding themselves under pressure from two distinct directions. On one hand, the global regulatory landscape for AI is hardening, with major frameworks activating and shifting, while a widening patchwork of national and subnational requirements is creating compliance obligations that most enterprises are not ready to meet.

At the same time, AI capabilities are evolving and expanding. Autonomous AI agents are now in production at nearly half of surveyed organizations, introducing a new class of governance challenges that most established programs were not initially designed to address.

In this report, we outline what these challenges look like and explore how governance-mature organizations are attempting to solve them. We examine the current state of AI regulation, the arrival of agentic AI, and the policy frameworks aimed at keeping businesses accountable and secure.

This report identifies five core themes, each aligned to the results unearthed in the survey data:

1. **Funding fuels false confidence:** Most organizations allocate resources to AI governance, yet investment and operational readiness are fundamentally misaligned.
2. **Bridging the agentic AI governance gap:** Autonomous agents are in production at nearly half of surveyed organizations, but governance controls remain incomplete and inconsistently applied.
3. **The accountability vacuum:** Responsibility for AI governance is concentrated in a single executive, while critical governance gaps remain invisible to boards.
4. **An impending regulatory reckoning:** Organizations are aware of incoming regulations but lack the operational readiness to meet them, creating compounding compliance risk.
5. **The trust dividend:** Governance maturity is a measurable driver of business performance, competitive advantage, and customer trust.

THEME 1

Funding Fuels False Confidence

It's clear that the gaps in governance and innovation are not due to a lack of funding. In fact, the vast majority of respondents (90%) say their organizations have already allocated funding specifically for AI governance. But that investment also reveals where the divide begins to widen. Building governance capacity is an important first step, but funding initiatives is not the same as operationalizing them, and funding alone does not mean an organization is audit-ready.



of organizations describe their AI governance as fully mature — fully operationalized and continuously monitored.

The result is a reality where organizational confidence in AI governance capabilities continues to outpace actual program maturity and operational readiness.

While 9 out of 10 respondents say they have the funding in place, how it is implemented reveals significant inconsistencies.



maintain a formal AI governance policy.



have documented AI-specific incident response procedures.



have a formal, documented AI acceptable use policy that is actively communicated to employees.

Consistency matters when dealing with enterprise governance. The fact that there's so much variance between core policy approaches points to an imbalance in governance implementation. Even more striking, however, is what respondents are doing even when policies do exist. Among organizations with a formal acceptable use policy, 28% still address violations informally on a case-by-case basis.

Policy without consistent enforcement is not audit-ready governance. This gap between policy and enforcement is the primary driver of compliance risk, where independent assessment delivers the highest strategic value.

Seventy-four percent of respondents describe their organizations as capable of passing an AI compliance audit today. That confidence is not matched by structural evidence: Only 27% describe their AI governance as fully mature, and 57% maintain formal governance policies.

The next step is to ensure governance funding is directed toward the operating model for AI oversight, not just to standalone tools or policy development.

Mature programs require clear ownership, defined roles, formal risk and impact assessments, lifecycle controls, documentation, monitoring, audits, and continuous improvement.

Internal AI use cases span departments such as software development and engineering; governance, risk, and compliance (GRC); HR; accounting; and finance. On the surface, these use cases may look similar operationally, but each comes with its own unique set of risks and oversight needs. Applying a generic AI governance policy across the board can leave some areas underfunded and resources misallocated. A mature AI governance policy will allocate funding to the practical work required to manage each area safely.

Ultimately, the survey findings reveal a widening gap between self-assessed readiness and actual governance maturity. Shrinking that gap is an important first piece of the governance puzzle. And making that happen requires investment across four critical areas:

1. **Governance tooling:** Invest in governance platforms that turn AI policies into enforceable workflows. Without tooling, enforcement remains manual, inconsistent, and difficult to prove.
2. **Process and documentation:** Invest in writing and documenting policies, procedures, and standards, including creating playbooks and building decision frameworks and protocols. This is where organizations with informal enforcement often fail.
3. **Audit and assessment:** Invest in regular independent evaluation to identify and address gaps in governance. Self-assessment is important, but on its own, it can be unreliable, subjective, and often insufficient.
4. **Training and awareness:** Invest in structured, organization-wide communication and training on policies and responsibilities to ensure proper enforcement.

When we look at organizations with mature governance, these four areas consistently stand out in budget prioritization. The organizations that struggle the most tend to invest in tools alone without the proper people, process, and governance infrastructure in place.

Fortunately, there are a few signs that indicate funding isn't working. This includes existing policies that employees don't know about, inconsistent use of established governance tooling, violations addressed on an ad hoc basis rather than through documented procedures, and AI-related incidents that expose a governance control previously unknown to be inadequate.

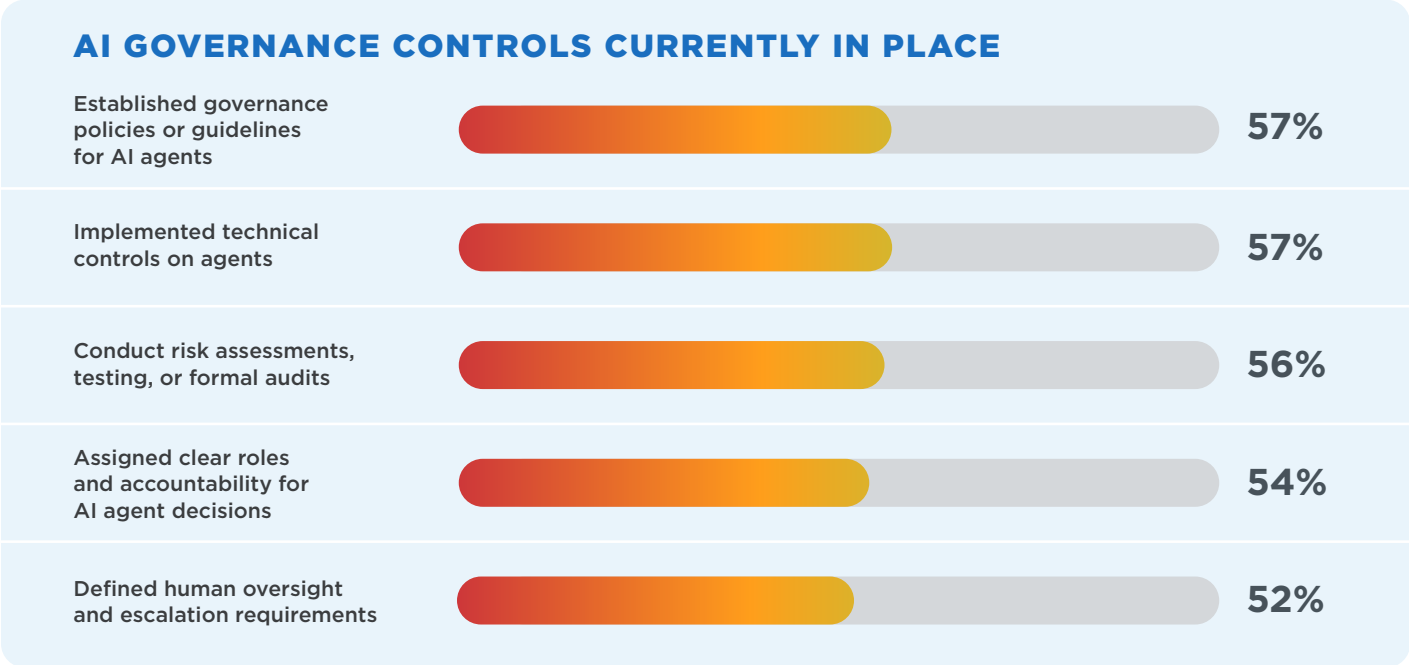
THEME 2

Bridging the Agentic AI Governance Gap

Agentic AI is quickly making the jump from experimentation to reality. Most (86%) respondents have at least put agents into testing, and almost half (46%) have agents in production within their organizations right now. Conventional wisdom might lead to the assumption that a focus on governance would slow the adoption of AI agents.

Survey respondents indicate that this assumption is not true. More than three-quarters of respondents (78%) with mature AI governance had agents in production, compared to just 22% of those who classified their governance as "developing."

But even so, in most cases, governance controls are still applied inconsistently. Among the organizations that have agents in testing or production, at least half had each of the following individual controls in place.



What makes this governance gap so impactful for AI agents is the technology's autonomous nature. Agents operate with access to enterprise systems containing sensitive data that must be strictly managed and governed. Traditional frameworks were designed to audit systems that teams build and control themselves with human input and action. And this is where agents invite a new kind of risk that can act on its own within vendor-supplied environments, with access that resembles but exceeds traditional application permissions.

Let's look at a real-world scenario. An AI agent connected to core systems, such as billing or customer support, often starts as a productivity-focused tool, but once it can execute account changes or trigger customer communications, governance must ensure those actions are reviewed and controlled.

Organizations need to govern agents with the same rigor they already apply to other high-risk technologies. The issue is not only whether an agent produces accurate outputs, but whether its decisions are explainable, reversible, accountable, and supported by technical testing and governance evidence when it interacts with sensitive data or executes real transactions.

Agents offer autonomy, but they still need oversight. Respondents revealed that the determination of how and when review is required is mixed.

38%
require human review only for high-risk or high-impact decisions.

32%
require human review for all actions, regardless of risk or impact.

22%
have established defined thresholds that would trigger a human review.

The autonomy of AI agents has left organizations grappling with a seemingly simple yet complex question: When does an agent need human review? While there is no standardized industry approach to this question, forward-thinking organizations implementing agent governance are increasingly adopting internally tiered risk frameworks to determine oversight requirements. The following exemplifies one approach organizations are piloting, though specific tier definitions and risk thresholds may vary by organization and industry:

TIER 1 LOW RISK

Low-risk activities include internal process automation, such as expense reports or scheduling. Agents can execute autonomously, and human review is not required as a documented audit trail is sufficient.

TIER 2 MEDIUM RISK

Medium-risk actions include tasks like email drafting, customer communications, and data analysis. Agents can execute with delayed review; human review is required within 24 hours, and activities can be reversed if flagged.

TIER 3 HIGH RISK*

High-risk tasks include agents having access to financial systems, customer data, external communications, or source code repositories (particularly changes affecting access management, security configurations, or infrastructure). Agents require preapproval for these actions; human review is required before execution, and a clear approval workflow is implemented.

TIER 4 CRITICAL

Critical actions involve decisions affecting customer service agreements, legal commitments, and data deletion. Agents should not operate autonomously in these instances, and human review is always required. Agents can assist with these tasks, but humans need to own all decisions.

! **Within high-risk categories, organizations may implement sub-thresholds. For example, agents with source code access may execute routine changes (dependency updates, documentation) with delayed review, while critical changes (encryption logic, access management) require preapproval.*

Organizations without a structured strategic framework typically default to either requiring review for everything, which slows innovation, or trusting the agent blindly, which creates risk. A tiered approach enables you to move fast while acting safely and cautiously where needed.

The key takeaway here is that governance doesn't have to be a hindrance to pushing AI agents from experimentation to implementation, but half-governing agents with an arbitrary or unstructured approach is not an effective solution.

A CHECKLIST FOR TECHNICAL CONTROLS

ACCESS LOGGING
Every action an agent takes should be logged with a timestamp, details on which data was accessed, and the outcome.

ROLLBACK CAPABILITY
Agents should never be able to make permanent changes without the ability to reverse any actions.

RATE LIMITING
Controls should prevent agents from spamming or overwhelming systems.

OUTPUT VALIDATION
Agent outputs must be validated before impacting systems.

INTEGRATION CONTROLS
Measures should ensure that agents can connect only to approved systems and APIs.

DATA RESIDENCY
Agent-processed data must remain in compliant regions.

AUDIT TRAILS
All agent decisions need to be traceable and explainable.

THEME 3

The Accountability Vacuum

As AI sprawl grows, organizational alignment and accountability are critical to keeping use in check, ensuring policies are implemented correctly, and maintaining consistent governance. When it comes to purchasing or adopting AI within an organization, most respondents place that responsibility firmly with the CIO or head of IT (42%), followed by the chief data officer/AI officer (16%), and CEO (10%).



of respondents say the CIO is chiefly accountable if an AI system creates risk or if something goes wrong.

What this shows is that organizations don't necessarily have an issue with creating formal ownership over decision-making. But the counter to this fact is that ownership, in its current state, is very narrow. Oftentimes, just one executive holds both the buying decision and the liability exposure, which presents a serious structural risk.

When one executive owns both adoption decisions and liability exposure, accountability shifts to the individual rather than the organization as a whole. Risk assessment suffers because a single person is evaluating risk that they are incentivized to minimize. Escalation is delayed because no one else is formally accountable. When incidents occur, the organization's response is limited to correction rather than meaningful root-cause remediation.

In most organizations, AI risk accountability has not been formally distributed across the functions most likely to encounter it daily. Distributed accountability, where risk ownership is spread across IT, security, compliance, and legal, is the industry practice for high-stakes decisions. AI governance remains largely centralized.

Top AI Governance Concerns

61%
have formally discussed AI risk exposure and liability.

59%
have discussed regulatory compliance requirements.

47%
have discussed AI incident response readiness.

60%
have discussed AI investment strategy and budget.

56%
have discussed ethical AI use and responsible AI principles.

36%
have discussed third-party or vendor AI risk.

But accountability isn't just a priority at the organizational level. Boards also have a critical role to play in this process. However, only roughly half (54%) say they regularly report AI governance to their board of directors or equivalent executive leadership body. And the boards that are having conversations about AI governance are doing so inconsistently. Most have covered the bigger trends and topics, but gaps remain in lower-priority areas.

So, what constitutes a clear accountability model in practice? Mature organizations do not eliminate executive ownership, but they avoid making a single leader the sole point of accountability. Instead, they distribute responsibility across IT, security, compliance, legal, procurement, business units, and executive leadership so each function owns the risks and decisions closest to its role.

To build out that model, organizations should follow these five steps:

1. **Designate an AI leader or governance team**
   Assign clear responsibility for maintaining visibility into how AI is being used across the organization.

2. **Formalize and enforce AI usage policies**
   Translate broad governance principles into clear expectations, approved use cases, enforcement procedures, and escalation paths.

3. **Train all AI stakeholders**
   Establish role-based training programs for developers, deployers, and users of AI systems. AI governance requires a baseline understanding across the entire AI lifecycle, from development through deployment and use.

4. **Assess AI systems regularly**
   Review AI use cases, controls, and risks on an ongoing basis to ensure governance keeps pace with adoption, regulation, and internal standards.

5. **Treat AI governance as a strategic investment**
   Allocate budget and resources toward governance capabilities that improve decision-making, accountability, and long-term AI effectiveness.

Accountability isn't limited to internal roles and teams. One of the most concerning gaps appears with third-party or vendor AI risk. Just one-third of respondents say this is a top concern for their organization, but it is simultaneously one of the biggest risk factors for AI governance. Just 69% are confident in their ability to govern AI embedded in third-party tools.



Only 36% of boards discuss third-party AI risk, yet it's the number-one blind spot in enterprise AI governance. Nearly every organization uses SaaS systems (such as customer relationship management, analytics, and collaboration tools) with embedded AI, but third-party AI is often invisible as vendors don't typically proactively disclose what data is used or how. Despite this gap, organizations are liable for the outcomes of third-party AI even though they don't control it.

Organizations that emphasize the risks associated with third-party vendors report better governance outcomes. When it comes time for the next board meeting, leaders should **center the conversation around these questions:**

- What percentage of our enterprise software has embedded AI that we're aware of?
- For our top 10 SaaS vendors, what is their AI data policy? Do we have contractual controls?
- If a third-party AI system causes harm, such as a customer data breach or discriminatory output, what is our liability?
- Do we have a contract review process that addresses AI before we sign SaaS agreements?

Concentrated responsibility and a disconnect in board focus reflect a growing blind spot for organizations. This blind spot will only grow as AI becomes deeply embedded in vendor platforms, SaaS tools, and enterprise software that organizations did not build and cannot directly inspect.

THEME 4

An Impending Regulatory Reckoning

The regulatory landscape is constantly in flux. New policies and guidelines are evolving and emerging on a global scale, leaving organizations to grapple with issues such as data sovereignty while also navigating the transformative power of AI on existing governance frameworks.

Fortunately, most organizations are aware of AI regulation. Eighty-four percent of manager-level respondents and 94% of directors are aware of state-level regulations that may apply to their organizations.



of organizations have experienced AI-related incidents or near-misses in the past 24 months, but only 44% have documented AI-specific incident response procedures.

As it stands, respondents have adopted a wide range of governance frameworks to meet regulatory requirements.



have the NIST cybersecurity framework in place.



established a proprietary or internal AI governance framework.



use the ISO/IEC 27001 framework.

Unfortunately, preparation is not as widespread or consistent as awareness. With major global frameworks activating and a rapidly expanding patchwork of national and subnational requirements creating overlapping obligations, most enterprises are better at knowing what's coming than building programs capable of meeting it. This realization is even more visible when looking across regions. Although nearly all organizations (94%) operate in regions with AI regulatory requirements, preparedness has not followed closely behind.



have taken specific actions to prepare for U.S. regulations.



have taken actions to prepare for EU AI Act compliance.



have taken actions to prepare for APAC requirements.

For businesses that operate across the world, compliance needs are quickly outpacing what is even possible with existing AI governance infrastructure. Organizations largely know the regulations in the space and many of the accompanying requirements. The problem is determining what does, or doesn't, apply to the organization itself. Some policies have gone into effect, some will be active in the near future, and others remain on the horizon for 2027 and beyond. As an example, let's look at the timeline for the EU AI Act.

EU AI ACT ENFORCEMENT TIMELINE
(Updated with Digital Omnibus changes)



Important caveat: These enforcement dates reflect the May 7, 2026, provisional agreement on the EU AI Act Digital Omnibus. The Omnibus has not yet been formally adopted or published in the Official Journal (expected before August 2, 2026). Until formal adoption, these updated dates are not legally binding. If formal adoption does not occur before August 2, 2026, the original (unamended) enforcement deadlines apply from that date. Organizations should monitor the EU Official Journal for official publication to confirm binding status.

The EU AI Act is a highly complex piece of regulation, and organizations across the globe are left unsure of how every stage will, or will not, impact them. Yet it is far from the only example. Considerations also need to account for regulations ranging from state-level policies in the U.S., such as the Texas Responsible AI Governance Act (TRAIGA) and New York's RAISE Act, to broader policies like Europe's GDPR.

Success starts with building governance programs that are resilient enough to absorb regulatory changes without having to start from scratch each time a deadline shifts or a new jurisdiction activates. Specifically, organizations can tackle this challenge by building a regulatory tracking and response system. The vast majority (94%) of respondents to Schellman's survey have awareness of the regulations that may apply to them, and most have even entered the assessment phase. But that only takes an organization partway to maturity.

Mature governance programs should be built on these four steps:

STEP 1: AWARENESS

- Subscribe to regulatory tracking services
- Join industry associations that publish regulatory updates
- Monitor government and agency announcements

STEP 2: ASSESSMENT

- For each regulation, identify which business processes are affected
- Map to your current controls and policies
- Identify compliance gaps

STEP 3: ROADMAP

- Prioritize by impact, with high-risk regulations first
- Sequence implementation (consider planned enforcement dates)
- Allocate budget, resources, and timeline

STEP 4: IMPLEMENTATION

- Update policies and controls
- Train teams and communicate expectations
- Audit compliance
- Document for regulators

THEME 5

The Trust Dividend: AI Governance as a Performance Driver

The focus on governance almost always comes from a place of compliance and risk management. But that perspective is highly limiting. AI governance, when implemented correctly, can become a measurable driver of business performance. Businesses with mature governance consistently report better operational outcomes, greater confidence in scaling AI, and tangible commercial advantages.

More than half (57%) of respondents say effective AI governance translates directly into improved efficiency, 43% indicate it makes AI scaling and innovation easier, and 49% believe it improves readiness for new regulations. There is a direct thread from investing in AI governance and outcomes that yield performance gains that, in a competitive market landscape, could help carve out an advantage for organizations.



It's easy to get lost in the idea of improved efficiency. So, considering 57% of organizations report improved efficiency as a benefit of effective AI governance, how do we characterize this?

Across the enterprise landscape, we see benefits that include:

- **Faster sales cycles:** Governance documentation, certifications, and independent assessments provide buyers with evidence they can review during due diligence.
- **Lower incident exposure:** Mature programs identify risks earlier through monitoring, escalation paths, and documented response procedures.
- **More predictable compliance costs:** Proactive governance reduces the need for rushed remediation when new requirements take effect.
- **Stronger customer retention:** Customers are more likely to trust vendors that can demonstrate responsible AI controls and accountability.
- **Competitive differentiation:** Governance becomes a proof point in RFPs, vendor reviews, and enterprise sales conversations.

The benefits also extend beyond operational enhancements. AI governance certifications are becoming a standard part of the RFP process and vendor evaluations as prospective customers increasingly want to work with organizations that are secure and operating with a mature governance infrastructure. These certifications are quickly becoming a differentiator externally rather than a compliance checkbox.

Those expectations aren't just applicable to AI-native businesses. Interest in AI governance certification is nearly universal, from professional services to SaaS companies. The reality is that a broad set of customers now expect organizations to prove that AI is being governed responsibly.

The survey data shows organizations recognize governance as a competitive advantage, confirming the shift we've been seeing take shape. In 2024-2025, governance was a "nice to have" for most organizations. In 2026, AI governance certifications like ISO 42001 have become standard in RFPs. By 2027 and beyond, AI governance certifications will become mandatory for many customer contracts. Customers want independent verification, not self-assessment, and early certification builds competitive advantages.

As customers demand verification and certification, AI governance will be increasingly treated as part of a business's go-to-market plans rather than just a compliance checkbox. From documenting controls to maintaining audit-ready evidence, customers want to know where an organization stands.

Trust is paramount. The question businesses ask is shifting from "do you have an AI governance program?" to "can you prove it?"



CONCLUSION

Rethinking the Future of AI Governance

The growing gap in AI governance isn't an awareness issue or a lack of funding. Most organizations have funding, policies, board reporting, and governance frameworks already in place. But just getting AI governance programs off the ground does not translate to an audit-ready, mature AI governance stance. It may look good on paper, but in practice, enforcement is inconsistent and incomplete. That's where the gap exists.

The organizations that close this gap first will be better positioned to scale AI with confidence, respond more quickly to regulatory changes, demonstrate accountability to customers and partners, and reduce the likelihood that innovation creates unmanaged risk.

In the future, effective AI governance will not be defined by adoption alone. It will be determined by whether organizations can prove that their governance programs are mature and ready to adapt to regulatory shifts and the ever-evolving nature of AI.

READY TO STRENGTHEN YOUR AI GOVERNANCE PROGRAM?

Contact an expert to learn how Schellman is helping organizations assess their AI governance maturity, build practical roadmaps, prepare for emerging regulatory expectations, and pursue recognized certifications and assessments.

As the first ANAB-accredited ISO 42001 certification body and the first authorized auditor for AIUC-1, Schellman brings deep experience helping organizations turn AI governance from a policy exercise into a provable, audit-ready program.



RESEARCH METHODOLOGY

The 2026 State of AI Governance survey, commissioned by Schellman and fielded by Researchscape, examines the views of 525 U.S.-based professionals involved in AI governance. The survey was fielded from April 13 to May 11, 2026. To qualify, respondents had to be in selected departments and roles at companies with more than 500 employees and an annual revenue of $100 million or more. The results are not weighted.

ABOUT SCHELLMAN

Schellman is a leading provider of attestation and compliance services. We are the only company in the world that is a CPA firm, a globally licensed PCI Qualified Security Assessor, an ISO Certification Body, HITRUST CSF Assessor, a FedRAMP Independent Assessor, and most recently, an APEC Accountability Agent. Renowned for expertise tempered by practical experience, Schellman's professionals provide superior client service balanced by steadfast independence. Our approach builds successful, long-term relationships and allows our clients to achieve multiple compliance objectives through a single third-party assessor. Learn more at schellman.com.

schellman
Quality. above all.

schellman.com

4010 W Boy Scout Blvd / Ste 600 / Tampa, FL 33607

Contact an expert or call 1.866.254.0000

Outside of the United States, please dial: +1.813.288.8833

Follow @Schellman on:

©2026 Schellman Compliance, LLC. All rights reserved.