




From Principles to Practice: Governing AI in the Corporation

GOVERNANCE & SUSTAINABILITY CENTER

From Principles to Practice: Governing AI in the Corporation

» This report examines how large companies are putting AI into governance practice, drawing on a survey of 130 corporate governance, sustainability, and citizenship leaders, as well as insights from Chatham House Rule convenings. It shows that AI governance is becoming more formalized and embedded in oversight, risk, and control structures, but remains uneven as companies begin to confront AI's broader workforce and environmental implications at scale.

Trusted Insights for What's Ahead®


- As AI is recognized as a material business issue—with more than 80% of S&P 500 companies disclosing AI-related risks in 2025 filings—responsible AI is moving from a niche ethics issue toward a central place in corporate governance, risk management, and operations.
 - 80% of surveyed executives expected AI to drive substantial productivity and efficiency gains, but 75% also anticipated major disruption to jobs and workforce structures.
 - Nearly two-thirds of surveyed executives said their company has developed enterprise-wide AI principles and more than half reported a centralized cross-functional AI council, suggesting governance is moving toward institutionalization.
 - AI is increasingly being folded into formal enterprise risk processes, with 70% of governance respondents saying it now appears in the enterprise risk inventory or heat map, with a focus on immediate risks such as cybersecurity and data breach.
 - Board engagement on AI is rising but fluency remains mixed: only 23% of governance respondents rated their board's AI literacy as high, while most boards are building oversight through existing committee structures and briefings rather than redesigning architecture.
 - The AI regulatory landscape is fragmented and only 9% of governance respondents said they are very prepared for emerging rules, with concern focused on the EU and US states.
 - Environmental and workforce dimensions remain less embedded in governance but will likely become material as AI expands across infrastructure, products, and communications.
-

Responsible AI in the Corporation

Figure 1

A strong majority of large-cap US public companies publicly acknowledge AI as a material business risk

Share of companies disclosing AI-related risks in Form 10-K, S&P 500, 2023–2025



Note: Some companies report more than one AI-related risk.
Source: The Conference Board/ESGAUGE, 2026

AI is moving quickly into core business processes across large companies, from internal productivity and software development to customer-facing functions, risk management, and decision support. As use expands, so does the need to govern AI as an enterprise issue rather than a narrow technology issue. That shift is already visible in the market: AI-related risk disclosure has risen sharply in recent S&P 500 annual filings (Figure 1), signaling that boards and management increasingly view AI as a material business issue.

In the corporate setting, AI governance refers to the policies, controls, accountability structures, and escalation mechanisms companies use to oversee how AI systems are developed, procured, deployed, monitored, and, where necessary, constrained. The discussion is also gradually widening to include adjacent issues such as AI's energy demands and workforce implications.

This report examines that shift through the perspectives of corporate governance, sustainability, and citizenship leaders. These functions do not own every AI decision but they are often closest to the points where AI becomes an enterprise issue, including board oversight, disclosure, regulatory exposure, stakeholder trust, workforce effects, and broader questions of accountability. Survey respondents see AI primarily as a source of productivity and efficiency gains but they also expect significant disruption to jobs and workforce structures (Figure 2). Companies do not view AI only as a tool for growth but are recognizing its broader implications for workforce planning, trust, governance, and environmental impact.

Figure 2

Surveyed leaders expect major productivity gains from AI, alongside workforce disruption

Percentage of survey respondents that agree with the following statement:




Note: 130 governance, sustainability, and citizenship executives responded.
Source: The Conference Board, 2026

Building AI Governance

Figures 3a and 3b

Most companies have formal AI principles, with implementation focused on internal governance, system design, and training

Q: Does your company have formal principles, guidelines, or policies for AI?



Q: How are your company's AI principles put into practice? (Select all that apply)



Note: 125 governance, sustainability, and citizenship executives responded.
Source: The Conference Board, 2026


Many companies have moved beyond broad discussion and begun to put formal governance around AI in place. Nearly two-thirds of respondents reported that their company has an enterprise-wide framework or set of principles for responsible AI (Figure 3a), suggesting that AI governance is becoming more structured and less experimental. This represents a significant shift from just a few years ago.

Those principles are being put into practice primarily through governance and oversight, system design, and employee training (Figure 3b). Fewer companies reported embedding them in supplier standards, public communications, or broader sustainability; environmental, social & governance (ESG); and citizenship efforts. For now, most firms are building internal discipline before extending expectations across the value chain.

Figure 4

About half of surveyed firms have a centralized enterprise-wide AI steering committee or equivalent

Q: Which statement best describes your company's internal coordination on AI?



Note: 125 governance, sustainability, and citizenship executives responded.
Source: The Conference Board, 2026

That internal focus is reinforced by how companies are organizing governance in practice. More than half of respondents report a centralized, enterprise-wide AI governance council or steering group (Figure 4), while others describe more decentralized but still coordinated approaches. AI governance is no longer being left to ad hoc experimentation across business units.


The composition of these structures is also telling. Where enterprise-wide AI bodies exist, they are most often populated by technology leaders (88% of survey respondents), legal and compliance teams (84%), and business units (73%). Sustainability (19%) and corporate citizenship (8%) functions are much less commonly represented. In practice, many companies are building responsible AI primarily through a technology, legal, compliance, and enterprise-risk lens. That is a sensible starting point but it also helps explain why current governance is centered on issues that fit most easily into existing control systems.

What Corporate AI Governance Is Prioritizing

Figure 5

Most surveyed executives consider cybersecurity-related challenges to be the most significant AI risk

Q: Which of the following AI-related risks are most significant to your company? (Select up to three)



Note: 125 governance, sustainability, and citizenship executives responded.
Source: The Conference Board, 2026

The risks companies prioritize provide one of the clearest signals of what AI governance means in practice. Respondents most frequently identified cybersecurity and data breach risk, privacy and data protection, legal liability, regulatory compliance, and operational or implementation challenges as the leading concerns (Figure 5).

The prominence of cybersecurity is notable. As AI expands the attack surface and accelerates phishing, fraud, and code exploitation, many companies are treating it as a core AI governance issue rather than a separate technical concern. That also argues for a prominent role for chief information security officers in AI oversight.


The ranking is also revealing for what sits lower on the agenda. Environmental impacts, workforce effects, and other broader external consequences rank well below the immediate enterprise risks that fit more easily into existing legal, compliance, risk, and control frameworks. That does not mean those issues are unimportant. It suggests they are less embedded in current governance systems and harder for companies to translate into near-term action.

Board and Enterprise Oversight of AI

Figure 6

AI is integrated into enterprise risk management at most surveyed firms

Q: How is oversight of AI integrated into your company's enterprise risk management (ERM) framework or equivalent processes? (Select all that apply)




Note: 52 governance executives responded.
Source: The Conference Board, 2026

As AI governance becomes more formalized, it is also moving into the broader control environment. Among governance respondents, 70% said AI appears in the enterprise risk inventory or heat map (Figure 6). Many also reported that AI-related risks are assessed through operational, compliance, cybersecurity, or privacy processes; reported to senior management or the board; and assigned to specific owners. This marks a shift from just a few years ago, when AI governance at many firms was more fragmented and informal.

Figure 7

More boards are situating AI oversight in a technology committee

Reported board committee responsible for oversight of AI, S&P 500, 2023–2025



Note: Data only includes S&P 500 companies that explicitly disclose oversight of AI in board committee charters.
Source: The Conference Board/ESGAUGE, 2026


Board involvement in AI is also increasing but the oversight model is still evolving. Among S&P 500 companies that disclosed where board-level AI oversight sits, the audit committee is the most common locus (Figure 7). Since 2023, however, the share assigning oversight to the audit or finance committee has fallen, while the share assigning it to a technology committee or equivalent has risen.

That shift suggests a growing recognition that AI is not just a controls or compliance issue but also a strategic, technical, and operational one. Greater risk committee involvement also points to a parallel shift toward treating AI more explicitly as an enterprise risk issue. The practical implication is not that one model fits all but that companies should be clear about where AI oversight sits, how issues are escalated, and when full-board attention is needed.


Figures 8a and 8b

Few governance executives consider their board to be highly AI literate

Q: How would you characterize your board's fluency on AI issues?



Q: How is your board building or planning to build its understanding of AI? (Select all that apply)



Note: 52 governance executives responded.
Source: The Conference Board, 2026

Notably, only 23% of governance respondents said their board has a high level of fluency on AI (Figure 8a). This suggests board attention to AI is advancing faster than board capability, even as AI becomes more material to strategy, operations, risk, and reputation.


Most respondents said their boards are building fluency through regular board or committee discussions, internal briefings, and advice from external experts (Figure 8b). Far fewer reported more structured approaches such as formal director education, scenario exercises, or recruiting board members specifically for AI expertise.

This incremental approach is understandable but has limits. Regular discussion is not the same as deep fluency and boards that rely too heavily on periodic updates may struggle to distinguish between lower-risk uses of AI and the applications that warrant more active challenge, escalation, or oversight. For boards, the practical implication is not necessarily to recruit AI specialists immediately but to build enough collective fluency to question management effectively, understand where AI creates exposure, and recognize when existing oversight routines are no longer sufficient.

Figure 9

Large-cap boards are adding more technology and cybersecurity expertise than specific AI expertise

Share of all independent board directors reporting selected skills and experience, S&P 500, 2021–2025



Source: The Conference Board/ESGAUGE, 2026

S&P 500 disclosures on director skills and experience point in a similar direction. Where boards are adapting, they appear to be doing so more through broader technology, cybersecurity, and digital capability than through a rapid push to recruit directors explicitly identified as AI experts.


That is likely to be the more durable approach. For most companies, the goal is not to build a board of technologists but to ensure directors have enough fluency, and access to the right internal and external expertise, to oversee AI with appropriate rigor.

Readiness for the AI Regulatory and Governance Landscape

Figure 10

Companies are aligning AI governance with leading US and EU policy developments

Q: Which external frameworks or standards are most influential on your AI governance approach?
(Select up to three)



Note: 52 governance executives responded.
Source: The Conference Board, 2026

The next phase of AI governance will depend less on whether companies have articulated principles and more on whether they can operationalize them under growing legal, regulatory, and operational pressure. Governance respondents pointed to a relatively small set of external reference points as most influential, particularly the National Institute of Standards and Technology (NIST) [AI Risk Management Framework](#), the [EU AI Act](#), and emerging state-level rules in the US (Figure 10).

Key AI Governance Frameworks and Regulatory Signals

NIST AI Risk Management Framework: Developed by the US government's standards agency as a voluntary framework to help organizations identify, assess, govern, and manage AI risk across the lifecycle. Its four core elements—govern, map, measure, and manage—help companies assign accountability, define use cases, assess risk, and apply controls in a structured way.


The EU AI Act: The EU's comprehensive law governing AI systems built around a risk-based approach. It turns many responsible-AI concepts into legal obligations. For high-risk systems, it pushes companies toward demonstrable compliance through requirements such as risk management, documentation, transparency, human oversight, robustness, cybersecurity, and ongoing monitoring. Even for many US-based companies not directly in scope, it helps define what mature AI governance looks like in practice.

US state-level developments: The US state landscape is fragmented but consequential. Colorado provides the clearest example of a more prescriptive approach to high-risk AI, with requirements tied to discrimination risk, documentation, impact assessment, and consumer protections. California has relied more on applying existing consumer protection, civil rights, privacy, and competition laws to AI. The practical implication is a growing patchwork of legal pressure that raises the premium on governance systems that can hold up across jurisdictions.

Figure 11

Very few surveyed governance leaders are fully prepared to comply with emerging AI regulations

Q: How prepared is your company to comply with emerging AI-specific regulations (e.g., EU AI Act, US state laws)?



Note: 52 governance executives responded.
Source: The Conference Board, 2026

Even so, preparedness remains limited. Only a minority of governance respondents said their company is very prepared for AI-specific regulation, while most described their organization as only somewhat prepared (Figure 11). The external environment is evolving quickly and unevenly, with a mix of voluntary frameworks, binding rules, and emerging state-level requirements. For companies, the challenge is not simply tracking legal developments but building a governance model that can hold up across jurisdictions, use cases, and legal theories; and adapt as AI becomes more embedded across the business.


In practice, companies should focus on:

- Maintaining a current enterprise-wide inventory of AI uses, including internally developed tools, vendor systems, and employee use of generative AI;
- Strengthening data governance and controls, including data quality, access, provenance, retention, model inputs, and downstream use;
- Establishing a clear method for distinguishing lower-risk from higher-risk use cases;
- Assigning explicit ownership and improving cross-functional accountability across technology, legal, compliance, risk, privacy, security, procurement, human resources, and relevant business teams;
- Putting in place consistent review, approval, and escalation processes, including clear triggers for legal, compliance, risk, privacy, security, and audit that flow up to the board;
- Defining what constitutes an AI incident and using scenario planning or tabletop exercises to test whether governance, escalation, and oversight processes work in practice;
- Strengthening oversight of third-party models, tools, and vendors, including diligence on limitations, safeguards, documentation, and accountability; and
- Building documentation, monitoring, and audit trails that can withstand scrutiny from management, boards, auditors, regulators, and counterparties.

Figure 12

Looking ahead, governance executives are prioritizing data governance and regulatory readiness

Q: Over the next two years, where will your company primarily focus on strengthening AI governance?
(Select up to three)



Note: 51 governance executives responded.
Source: The Conference Board, 2026

What comes next for AI governance is likely to be more operational and less rhetorical. The companies that will be best positioned are unlikely to be those with the broadest principles but those that can show disciplined governance in practice: clearer ownership, stronger controls, better documentation, more reliable escalation, and a governance model that can adapt as AI use deepens and scrutiny intensifies.

The Next Frontier of Responsible AI Governance


AI governance is becoming more formalized, but remains focused on immediate enterprise risks such as cybersecurity, privacy, compliance, legal exposure, and operational control. As AI becomes more integrated into core operations, infrastructure, workforce systems, and external-facing products and services, that narrower approach is likely to become less sufficient. The next phase of AI governance will require companies to look beyond whether AI is governed safely and lawfully and pay greater attention to how its environmental and workforce effects are identified, understood, and managed in practice.

Environmental implications: AI's emerging sustainability challenge

Figure 13

Surveyed sustainability heads are most concerned about data center expansion, emissions, and water use

Q: Which aspects of AI's environmental footprint are of most concern to your company? (Select up to three)



Note: 49 corporate sustainability executives responded.
Source: The Conference Board, 2026

For most companies, the environmental dimension of AI remains less developed than more familiar governance concerns such as privacy, security, and compliance. Survey responses from corporate sustainability leaders suggest that AI's environmental effects are increasingly recognized, particularly around data-center electricity demand, grid carbon intensity, water use, and the broader resource footprint of AI infrastructure (Figure 13.)

However, action remains at an early stage:

- No surveyed sustainability executives reported quantified AI-specific environmental targets or metrics, though 28% expected to establish them within the next year;
- Only 17% said their company currently reports on AI's environmental impact, and even then only indirectly within broader disclosures, while 22% expected such reporting to emerge;


- Internal capability is similarly uneven, with more than half of sustainability respondents describing their function's AI fluency as basic or low;
- At the same time, over 70% expected stakeholder attention to AI's environmental impact to increase over the next two years, driven primarily by activists, employees, and consumers.

This matters because AI's environmental footprint is likely to become harder to ignore as use scales and the broader infrastructure demands associated with AI use become more visible. However, environmental teams are often only **lightly represented in enterprise AI governance structures** and few companies appear to treat AI-related energy, water, or emissions impacts as a distinct area of management focus.

Figure 14

US data centers remain heavily concentrated in a small number of states

Map of US data center intensity, by state, as of April 2026



Source: The Conference Board, 2026; based on [USA Data Centers](#), Data Center Map, accessed April 2026

The physical footprint behind AI helps explain why this issue is moving closer to the mainstream governance agenda. As data-center capacity expands and remains concentrated in specific regions (Figure 14), AI deployment is becoming harder to separate from broader questions of electricity demand, grid pressure, carbon intensity, water use, siting, and local scrutiny.

For many companies, these issues will surface indirectly at first—through cloud usage, vendor choices, procurement decisions, infrastructure strategy, or stakeholder pressure—rather than


through a standalone AI sustainability program. The more practical near-term task is to integrate environmental considerations more deliberately into existing AI and sustainability governance where they may become material.

Workforce and people implications: Trust, disruption, and community effects

Figure 15

Corporate philanthropy leaders are most concerned about AI's effects on workforce development, trust, and digital inclusion

Q: Which social or community impacts related to AI are of most relevance or concern to your company? (Select up to three)



Note: 29 corporate citizenship executives responded.
Source: The Conference Board, 2026

The workforce and people dimension of AI is becoming more clear but remains less embedded in most governance models. Survey responses from corporate citizenship and philanthropy leaders pointed to workforce displacement, reskilling, misinformation, trust, and unequal access to AI's benefits as the leading concerns (Figure 15).

Workforce displacement stands out in particular and aligns with the broader finding that 75% of surveyed executives expected AI to contribute to large-scale workforce disruption. While debate over AI's effect on jobs remains polarized, the credible view is that AI will automate some tasks, reshape many roles, create some new ones, and shorten the transition time for workers whose skills lose value. At the same time, reskilling alone may not be enough, especially if workers are retrained into roles that may themselves become newly exposed to automation. The core risk is not only job loss but disorderly transition: roles changing faster than workers and institutions can adapt, with consequences for trust, inequality, and community stability.

These issues are harder to govern because they cut across human resources, legal, communications, compliance, public affairs, citizenship, and business leadership rather than fitting neatly into traditional technology or control functions (Figure 16). Companies cannot control the broader labor market but they can govern their own contribution to disruption more deliberately by identifying exposed roles and tasks, linking AI deployment to workforce planning, escalating material impacts, and clarifying ownership. From a community engagement perspective, it likely means investing in reskilling and internal mobility, partnering with colleges and workforce nonprofits, supporting digital inclusion, and helping workers and communities navigate the transition.

Figure 16

Responsibilities for AI's workforce and community impacts are dispersed within organizations

Q: Where does primary responsibility for addressing the social and community impacts of AI reside within your organization?



Note: 29 corporate citizenship executives responded.
Source: The Conference Board, 2026

From a governance perspective, the immediate task is not to elevate every workforce or community issue into a board-level AI risk. It is to connect AI governance more directly to the functions responsible for workforce strategy, communications, stakeholder engagement, and community impact. Stronger practice will depend on clearer ownership, better cross-functional coordination, and earlier escalation of issues that otherwise fall outside the boundaries of technology or legal oversight.

Conclusion

As AI becomes more deeply embedded across operations, infrastructure, workforce systems, and external-facing products and services, governance will need to broaden. The next phase will favor companies that move beyond principles and periodic discussion to build clear ownership, stronger controls, more reliable escalation, and documentation that can withstand scrutiny from boards, regulators, auditors, and counterparties.

For governance, sustainability, and citizenship leaders, the task now is to help ensure that AI governance is not only controlled and compliant but also connected to the workforce, environmental, and broader enterprise effects likely to become more material over time.

About This Report

This report draws on two main inputs:

- A confidential survey conducted by The Conference Board® [Governance & Sustainability Center](#) from January to March 2026 of 130 senior executives across corporate governance (52 respondents), sustainability (49), and citizenship (29). Respondents skewed toward senior roles and larger companies, with 87% from US-headquartered firms.
- A 2025 series of four Chatham House Rule roundtables on responsible AI convened by the Governance & Sustainability Center. The series brought together more than 250 professionals across governance, ESG, sustainability, citizenship, and AI-related roles.

ABOUT THE AUTHOR



Andrew Jones, PhD is a Principal Researcher at The Conference Board Governance & Sustainability Center in New York, where he leads research and thought leadership on corporate governance, sustainability, and corporate citizenship. A prolific author and speaker, he frequently engages with and advises corporate leaders and their teams on a wide range of issues, including AI governance and AI's environmental and social dimensions.

THE CONFERENCE BOARD is the Member-driven think tank that delivers **Trusted Insights for What's Ahead®**. Founded in 1916, we are a nonpartisan, not-for-profit entity holding 501(c)(3) tax-exempt status in the United States.

© 2026 The Conference Board, Inc.