OASIS OPEN

Inside AI Governance:

What Business Leaders Need to Know Before Someone Else Decides for Them

AN OASIS OPEN WHITE PAPER

oasis-open.org

BLUF

Boards and investors are demanding AI adoption. CEOs are responding quickly, and often without a clear answer to the most basic business question: what, specifically, is AI supposed to do for this organization, against our business objectives, with specific and measurable outcomes? Without those answers, AI is not a strategy but an untethered tactic: an expense with unclear ROI that, beyond known integration and third-party risks, carries unknown perils. The organizations getting this right start with the business case: define the objective, commit to measurement, and build governance around outcomes. The frameworks they need are being built now by practitioners, led by the Coalition for Secure AI (CoSAI). We invite leaders to engage and shape them, or wait and comply with whatever others decide.

Three Theses

1. There is intense pressure to adopt AI, but most business leaders are responding to competitive fear, not strategy, and have not done the hard work of defining what AI will do for their business or how they will know if it worked.

2. Without strategy, organizations cannot measure what AI is actually doing, including its effects on productivity, revenues, and risk posture, and infosec, legal, and risk functions are left setting policy in a vacuum.

3. The regulatory landscape is fragmenting faster than it is converging, but the governance frameworks regulators will eventually codify are being written now by practitioners. Engage today or comply later with whatever others decided.

1. The Pressure Is Real. The Strategy Is Not.

Adoption is broad, results are thin, and the driver is fear not strategy.

Every board wants AI. Investors ask about it on earnings calls. Competitors announce it in press releases. CEOs are responding quickly, and often without time to answer the question that should come first. What, specifically, is AI supposed to do for this organization? And what results would demonstrate it has succeeded?

The data suggests most CEOs have not answered that question before deploying. A February 2026 study by NBER economists (Working Paper 34836, Yotzov et al.) surveyed approximately 6,000 executives across the United States, United Kingdom, Germany, and Australia and found that nearly seven in ten businesses now use some form of AI. Yet nearly nine-in-ten of those organizations reported no measurable productivity impact over three years. The aggregate

productivity gain attributable to AI across the full sample was 0.29 percent. Executives project that to rise to about 1.4 percent over the next three years, but the gap between deployment activity and realized benefit is already large and widening. Additional research finds that only 39 percent of organizations attribute any EBIT impact to AI, and fewer than 1 in 20 see impact above 5 percent.

The driver is not a strategy gap. It is a pressure dynamic. Competitive fear creates urgency. That urgency produces deployment decisions made by whoever moves fastest, not whoever thinks clearest. Organizations acquire tools before they define what those tools are for, and build AI capabilities before they build the measurement frameworks that would tell them whether those capabilities are working.

Structured conversations with OASIS and CoSAI's members found broad support for this picture. These are senior practitioners at some of the industry's largest technology and enterprise firms, with deep experience in technology and AI creation, implementation, and governance, not merely a general business sample. The signal was consistent: organizational optimism about AI, limited confidence in the ability to measure its impact, and governance structures still being designed rather than operated.

OASIS Open and CoSAI internal research identifies CEO ownership as a critical factor to AI initiative success. A meta-analysis of available studies finds that fewer than a third of organizations have it. CoSAI's AI Incident Response Framework identifies the pattern directly: "Organizations throw controls at AI without strategy. No threat progression model, no capability prioritization, just reactive patching and cleanups." As Melissa Pint, Chief Digital Officer of Frontier, has put it, "There is no such thing as a technology strategy. There's only a business strategy that technology supports."

Before approving any AI deployment, CEOs should require written answers to three questions. What business problem does this solve? How will we measure whether it solved it? What does failure look like? If those answers do not exist, the organization is not ready to deploy, it is ready to spend.

2. Without Measurement, Governance Is Guesswork

No strategy means no measurement, means no policy, means no governance… and the consequences are concrete.

> "You can have all the AI in the world, but if it's on a shaky data foundation, then it's not going to bring you any value." Carol Clements, Chief Digital and Technology Officer, JetBlue Airways

That observation applies equally to AI governance. Without a clear strategy, there is nothing to measure. Without measurement, there is no policy. Without policy, there is no governance. The chain breaks at the first link.

Our meta-analysis of available research and our structured conversations with OASIS and CoSAI members revealed a striking pattern: organizations cluster at the extremes. Either a company has substantial governance structures in place, built deliberately and operating as intended, or it has none. There is no such thing as a company that is mostly governed.

The data infrastructure problem runs underneath all of it. Only 22 percent of organizations are confident their IT architecture can support new AI applications (Databricks/Economist Impact, 2024), with barriers concentrated in access control (49 percent), sensitive data protection (42 percent), data silos (38 percent), and provenance tracking (31 percent). Organizations are deploying AI on a foundation they already know is inadequate.

Third-party risk compounds the problem. CoSAI's "Establish Risks and Controls for the AI Supply Chain" identifies visibility gaps across critical infrastructure as a primary concern and calls for comprehensive data governance frameworks with specific AI considerations, robust encryption and access controls, and ethical data collection methodologies.

The risks are concrete: reduced AI-driven efficiency benefits, increased costs of correcting mistakes and inaccuracies, regulatory non-compliance and remediation expenses, intellectual property loss, and privacy failures. What is less often acknowledged is that these risks are being managed, where they are managed at all, without the strategic clarity that would allow organizations to prioritize, measure, or learn from their mitigation efforts.

The Drift/Salesloft breach of August 2025 illustrates what this costs in practice. Approximately 700 customer organizations were affected. The cause was not an AI failure, but a governance failure: Attackers obtained OAuth tokens through prior phishing and used them to impersonate a trusted third-party integration, bypassing multi-factor authentication. The attackers moved laterally across connected Salesforce, Google Workspace, and Slack environments, exfiltrating business contacts, account records, and API keys and credentials that had been stored in support cases.

Three practices address this. Threat modeling before any deployment: map data access, failure modes, and blast radius in plain business terms, and repeat after every major vendor upgrade. "What is the cost of wrong?" asks Clint Bruce, a former US Navy Special Warfare Officer. Blast-radius reduction: limit access to only what is essential, document who made those decisions, and treat vendor defaults as a starting position to be narrowed, not accepted. Instrumentation: logging, monitoring, data loss prevention, and automated response must be in place before go-live, capable of distinguishing human actions from AI agent actions.

Every deferred governance decision increases the cost of eventual correction.

3. The Governance Infrastructure Is Being Built Now

The regulatory response to AI is real, accelerating, and not converging fast enough to be useful. Multistate.ai's tracking of state legislation documents 1,208 AI-related bills introduced across all 50 states in 2025, of which 145 were enacted. As of March 2026, lawmakers in 45 states have already introduced 1,561 more. Most enacted laws are narrow: deepfakes, fraud, public sector use cases. Comprehensive cross-sectoral AI governance legislation has stalled in nearly every state that attempted it. At the federal level, no national framework exists. The Brookings Institution's January 2026 analysis of 385 state AI bills found that legislative activity tracks ideology and fiscal capacity, not risk exposure: "AI legislation depends on both ability and appetite, and the absence of either reliably constrains policymaking." The pattern reproduces familiar divides in U.S. policy innovation. There is no reason to expect it to resolve quickly.

Waiting for regulatory clarity is itself a risk decision. It is a decision to let others write the rules, to accept compounding exposure during the interim, and to arrive at compliance rather than governance. The question is not whether requirements are coming. It is whether your organization will have had any hand in shaping what they look like.

There is a better model, and it has worked before.

In 1994, Netscape introduced SSL as a voluntary technical standard. It took 20 years to travel from that voluntary standard to binding regulatory requirement: PCI DSS mandated TLS for cardholder data in 2005; NIST SP 800-52 formalized TLS for federal systems in 2014; NIST SP 800-52 Rev. 2 mandated TLS 1.2 as the minimum in 2017. TLS 1.3, by contrast, went from draft to RFC in four years.

The lesson OASIS Board member Pablo Breuer draws from this arc applies directly: "Solve your challenges and present a solution rather than waiting for regulators to cause new problems."

When an entire industry converges on a standard that works, regulators recognize and ratify it rather than re-litigate it. Technical bodies move faster than legislatures. Standards emerge from practitioners who understand the technical realities, edge cases, and implementation costs. A single open standard applies globally without treaty negotiation or jurisdictional friction.

CoSAI and OASIS Open are building that infrastructure now. CoSAI's active workstreams are producing a Secure AI Framework, AI risk governance standards, software supply chain security frameworks for AI systems, and data lineage tracking for usage, privacy, and compliance.

CoSAI's Preparing Defenders of AI Systems framework identifies gaps in the governance landscape that "demand immediate investment, research, and innovation," and commits to "continually updating its content to reflect emerging technologies, new threat vectors, and evolving best practices, ensuring that our guidance remains current and actionable for all stakeholders."

This is living, practitioner-led governance built from real-world experience, not a static compliance checklist.

The organizations that shaped SSL are not the ones that waited until PCI DSS was ready before they began taking encryption seriously.

Engage now. CoSAI and OASIS are ready to talk today.

4. Four Actions Companies Should Take Now

- Define the business objective in plain language before selecting any technology; if you can't state what problem AI will solve and how you'll measure it, you're not ready to deploy
- Require documented cross-functional risk analysis (threat model, blast radius, instrumentation plan) before approving any implementation; this is a business risk problem, not an IT problem
- Get the data foundation right before pointing AI at it; provenance and metadata governance are preconditions, not afterthoughts
- Engage with OASIS Open and CoSAI now; the frameworks regulators will eventually codify are being written today; participate and shape them.

Or wait and comply with whatever others decide.

Works Cited

Coalition for Secure AI. “AI Incident Response Framework, V1.0.” OASIS Open, 27 Oct. 2025, www.coalitionforsecureai.org/wp-content/uploads/2026/03/AI-Incident-Response-1.pdf.

Coalition for Secure AI. “Establish Risks and Controls for the AI Supply Chain, V1.0.” OASIS Open, 12 June 2025, www.coalitionforsecureai.org/wp-content/uploads/2026/03/risks-and-controls-for-the-ai-supply-chain-v1.pdf.

Coalition for Secure AI. “Model Context Protocol (MCP) Security.” OASIS Open, 8 Jan. 2026, www.coalitionforsecureai.org/wp-content/uploads/2026/03/model-context-protocol-security-1.pdf.

Coalition for Secure AI. “Preparing Defenders of AI Systems, V1.0.” OASIS Open, 14 July 2025, www.coalitionforsecureai.org/wp-content/uploads/2026/03/preparing-defenders-of-ai-systems.pdf.

Bick, Alexander, et al. “The Rapid Adoption of Generative AI.” NBER Working Paper No. 32966, National Bureau of Economic Research, Feb. 2025, www.nber.org/papers/w32966.

Bruce, Clint. Personal communication with Nick Selby. 2025.

The Conference Board / Committee for Economic Development. “AI and the C-Suite: Implications for CEO Strategy in 2026.” The Conference Board, 15 Jan. 2026, www.conference-board.org/research/ced-policy-backgrounders/ai-and-the-c-suite-implications-for-ceo-strategy-in-2026.

Congressional Research Service. “Regulating Artificial Intelligence: U.S. and International Approaches and Considerations for Congress.” Report R48555, Library of Congress, 4 June 2025, www.congress.gov/crs-product/R48555.

Daly, Mary C. “The AI Moment? Possibilities, Productivity, and Policy.” FRBSF Economic Letter, no. 2026-06, Federal Reserve Bank of San Francisco, Feb. 2026, www.frbsf.org/research-and-insights/publications/economic-letter/2026/02/ai-moment-possibilities-productivity-policy/.

Economist Impact. “Unlocking Enterprise AI: A Global Study of 1,100 Technologists.” Economist Impact, supported by Databricks, Nov. 2024, www.databricks.com/resources/analyst-research/unlocking-enterprise-ai-opportunities-and-strategies.

Denford, James S., et al. “Why AI Policy Thrives in Some States and Fades in Others.” Brookings Institution, 14 Jan. 2026, www.brookings.edu/articles/why-ai-policy-thrives-in-some-states-and-fades-in-others/.

FINRA Cyber and Analytics Unit. “Cybersecurity Alert: Salesloft Drift AI Supply Chain Attack.” FINRA, 2025, www.finra.org/rules-guidance/guidance/salesloft-drift-AI-supply-chain-attack.

Humlum, Anders, and Emilie Vestergaard. “Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI.” NBER Working Paper No. 33777, National Bureau of Economic Research, May 2025, www.nber.org/papers/w33777.

Maslej, Nestor, et al. AI Index Report 2025. Stanford Institute for Human-Centered Artificial Intelligence, 8 Apr. 2025, hai.stanford.edu/ai-index/2025-ai-index-report.

MultiState. “State AI Legislation Tracker 2026: All 50 States.” MultiState, 2026, www.multistate.ai/artificial-intelligence-ai-legislation.

Selby, N., “How FOMO Is Turning AI Into a Cybersecurity Nightmare.” EPSD, Inc, December, 2025, https://www.inc.com/nick-selby/how-fomo-is-turning-ai-into-a-cybersecurity-nightmare/91261473

Singla, A., et al. “The State of AI in 2025: Agents, Innovation, and Transformation.” McKinsey and Company, Mar. 2025, www.mckinsey.com/capabilities/quantumblack/our-insights/the-state-of-ai.

United States, Government Accountability Office. “Artificial Intelligence: Federal Efforts Guided by Requirements and Advisory Groups.” GAO-25-107933, 9 Sept. 2025, www.gao.gov/products/gao-25-107933.

Yotzov, Ivan, et al. “Firm Data on AI.” NBER Working Paper No. 34836, National Bureau of Economic Research, Feb. 2026, www.nber.org/papers/w34836.