



Center for
AI Oversight

The AI Oversight Program

Oversight governance that enables Decision Velocity

ADVANCING AI OVERSIGHT GOVERNANCE FOR BOARDS, EXECUTIVES, AND REGULATORS

Center for AI Oversight · A 501(c)(3) nonprofit · cfaio.org



Who we are

The Center for AI Oversight is an independent, nonprofit educational institution. It defines oversight governance for AI: what boards and executives must govern, and who is accountable.

Mission

Advancing AI oversight governance as a board-level discipline for boards, executives, and regulators.

Standing

Guided by former senior leaders from the Federal Reserve, the SEC, the FCC, the U.S. Treasury, and HHS. A contributing member of the NIST AI Safety Institute Consortium.

Method

It defines what must be governed and who is accountable. It does not prescribe how, leaving operational design to operators.

An educational institution in the tradition of independent research and policy bodies.



Why the Program exists

The Center defines what boards and executives must govern in AI, and who is accountable, so they can make informed decisions at the speed the technology demands.

Governance is not a brake on AI adoption. It is what makes confident adoption possible.



An independent educational institution. The Program is the standard of care it teaches.

The convictions the Center is built on

The Bedrock Principles, in brief

- I** **Responsible advancement is a shared obligation.** Scaled to position and consequence. The base duty is the same: make informed decisions about the AI you are touching.
- II** **AI requires a native oversight lens.** Distinct from the frames built for cyber, model risk, and third-party risk. It asks what must be governed, who is accountable, and how it is evidenced.
- III** **The program is what makes AI governance defensible.** A charter, decision rights, escalation paths, reporting cadence, and an evidentiary record. The committee's role is to govern the program.
- IV** **Oversight done well is the infrastructure of speed.** Risk tolerances calibrated in advance and embedded in strategy. A risk identified is information that enables a decision, not an obstacle that stops one.
- V** **AI must advance. Clarity of obligation makes that possible at speed.** What slows serious institutions is not the duty but ambiguity about what the duty requires. The Center brings consistency to the expectation.



The duty of oversight is already settled

Caremark and its line define what boards and executives owe

Under Delaware law, the board owes a duty of oversight. Caremark, sharpened by Marchand and the Boeing board settlement, requires directors to put a system in place to monitor the corporation's mission-critical risks, and to act on the red flags it surfaces.

It reaches the top

The obligation runs to the public company board and the C-suite executives accountable to it. It is not delegable downward.

AI is in scope

As AI becomes a material, mission-critical risk, it falls squarely inside the existing duty of oversight.

Regulators expect it

Across financial services, healthcare, and public disclosure, regulators increasingly expect demonstrable oversight, not adoption alone.

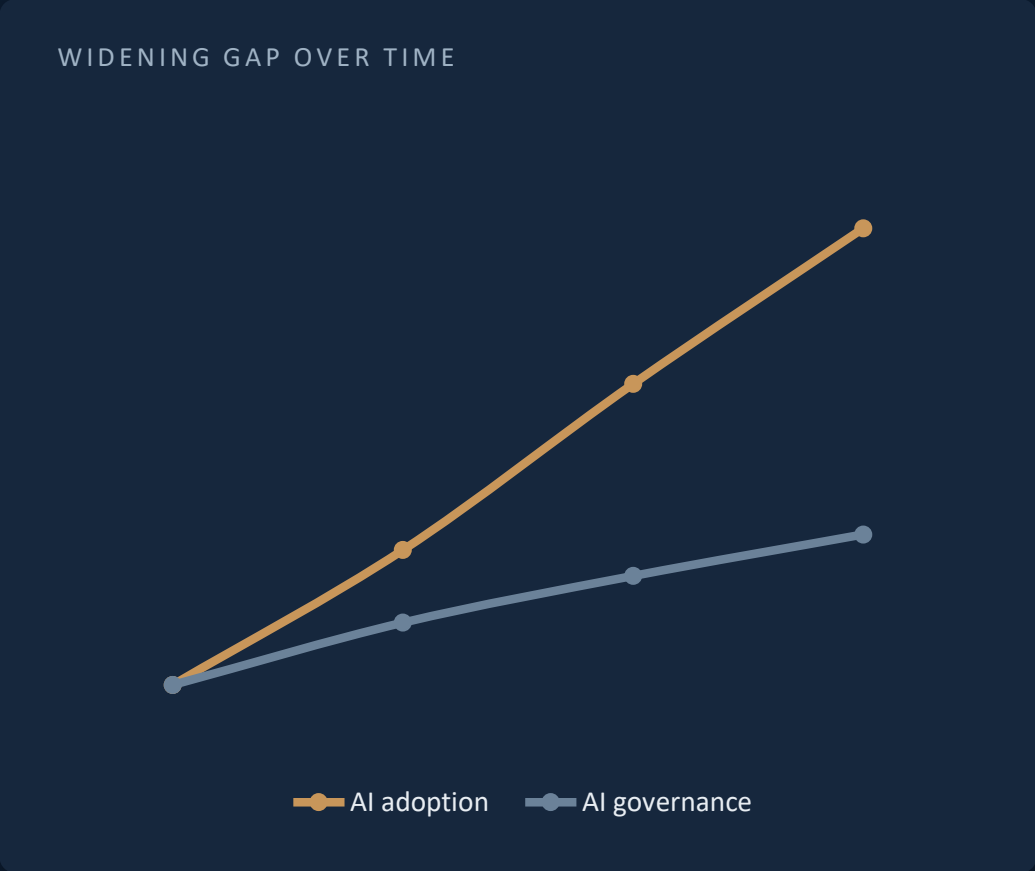
The duty is not in question. Whether the board can demonstrate it discharged the duty is.

The Velocity Gap

The problem the Program is built to close

AI is being adopted faster than most organizations can govern it. The gap between the speed of AI adoption and the speed of AI governance is where fiduciary risk, regulatory exposure, and missed strategic opportunity all live.

The Program is the institutional capability that closes it. The organizations that govern AI fastest will outperform those that govern it best on paper.



Why organizations need a Program

A different relationship with decisions, speed, and risk



Informed decisions

Boards and executives must be able to show they made an informed decision about AI, in evidence, not in assertion.



A different speed

The environment rewards faster, risk-informed decisions. Oversight has to operate at the speed of adoption, not behind it.



Risk as strategy

Risk tolerances set in advance become decision points, not refusals. Risk becomes an input to strategy rather than a brake on it.



Defensible governance

Governance that meets regulatory expectations and protects against legal liability, demonstrable when a regulator, court, or auditor asks.

Defensible oversight governance is the Program

What looks like oversight, but is not

The Committee Fallacy. A committee discussing AI is not a program.

The Assurance Trap. Auditing model controls while ignoring the executive decisions that set risk.

Framework Fatigue. A patchwork of standards with nothing internal to anchor them.

What the duty actually requires: a Program

- A charter and clear decision rights
- Risk tolerances set in advance
- Escalation paths and materiality thresholds
- A reporting cadence to the Governing Body
- A traceable evidentiary record

The program is the governance. *It is what makes oversight defensible when the duty is tested.*

What the Program defines, and what it leaves to operations

Govern the system, not the technology

WHAT we define · WHO is accountable

- The governance mandate and charter
- Risk appetite approved by the Governing Body
- Accountability across the three lines
- Standards AI systems must demonstrably meet
- Escalation protocols and materiality thresholds
- Evidence requirements for assurance

HOW is left to operational layers

- Specific control implementations
- Technical testing methodologies
- Model validation procedures
- Data quality process design
- Tool and platform selection
- Day-to-day operational execution

Clear boundaries produce faster, more confident decisions. The Program defines the architecture. It does not write the operational playbook.

The five pillars of the Program

Five pillars, fourteen domains, seventy-six controls

I Agile Governance *The Constitution*

Strategic direction, authority, and adaptive policy that guide the AI lifecycle.

II Risk-Informed System *The Guardrails*

A repeatable method to identify, assess, and prioritize AI risk against tolerance.

III AI Trust and Assurance *The Evidence*

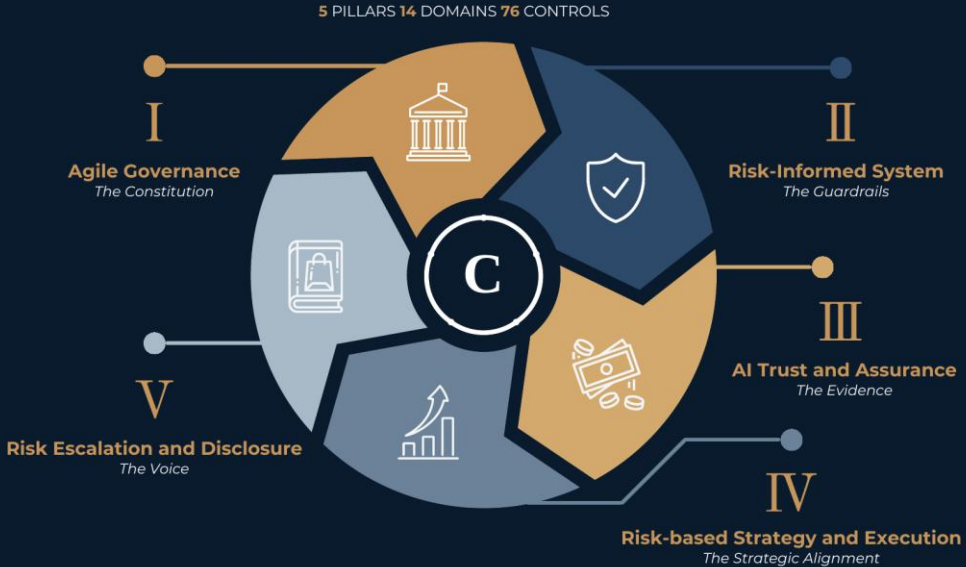
Lifecycle oversight of models, data, transparency, and security that proves trust.

IV Risk-Based Strategy and Execution *The Strategic Alignment*

The link between oversight and capital, value, and third-party risk.

V Risk Escalation and Disclosure *The Voice*

How critical risks reach leadership and how the organization discloses them.



The fourteen governance domains

Each pillar is established by its domains

I Agile Governance

The Constitution

- 01 AI Governance Program and Policy Framework
- 02 AI Governance Structure, Oversight, and Resources
- 03 Governance Program Assurance and Continuous Learning

II Risk-Informed System

The Guardrails

- 04 AI Risk Methodology, Scope, and Tolerance
- 05 Risk Intelligence and Threat Landscape

III AI Trust and Assurance

The Evidence

- 06 AI Model Risk and Agentic Lifecycle Oversight
- 07 AI Data Governance Oversight
- 08 AI Transparency, Explainability, and Human Oversight
- 09 AI Security and Resilience Assurance

IV Risk-Based Strategy and Execution

The Strategic Alignment

- 10 Risk-Informed Strategy, Resources, and Organizational Readiness
- 11 AI Value Realization and Operational Resilience Oversight
- 12 Third-Party AI and Supply Chain Governance

V Risk Escalation and Disclosure

The Voice

- 13 AI Risk Escalation and Disclosure Protocols
- 14 Validation of Escalation and Governance Effectiveness

Built on converging authority

Each control mapped to the sources that define the standard of care



Technical Standards

NIST AI RMF · ISO/IEC 42001 · ISO 38507

Repeatable risk processes and a voluntary governance structure.



Legislation and Regulation

EU AI Act · SEC · FTC · state laws

Legally binding obligations and risk-tiered classification.



Judicial Precedent

Caremark · Marchand · Boeing

Boards must maintain functioning oversight systems for material risks.



The AI Oversight Program

The interoperability layer

Fuses standards, law, and precedent into one executable system.



The outcome: Decision Velocity

Decision Velocity

The ability to make faster, risk-informed decisions backed by evidence, without sacrificing accountability.

The program is the proof point. It is also the accelerant.

Faster

Risk tolerances set in advance become decision points, not vetoes.

Evidenced

On demand, the Program shows the informed decision and the disclosure behind it.

Accountable

Speed without losing the chain from policy to practice to assurance.

How the Center engages

An educational institution. Four forms of work.



The Oversight Program

A practical model for board-level AI oversight, mapped to the laws, regulations, and case law that define what boards are accountable for.



Publications and Research

Co-developed works with institutional partners, including the Private Director's Body of Knowledge and the Internal Audit assurance supplement, plus regulatory analysis.



Education and Training

NASBA-compliant education for directors, audit executives, and risk leaders. Governance competency for senior professionals, not technology tutorials.



Convening and Speaking

Keynotes, panels, board briefings, and roundtables that bring the Center's thinking to conferences, boardrooms, and regulator forums.



Guided by leaders who built the regulatory landscape

The advisory board represents the federal agencies that built the oversight and regulatory infrastructure governing AI adoption. A contributing member of the NIST AI Safety Institute Consortium.

Federal Reserve

Former Executive Vice President, Federal Reserve Bank of Boston

Securities and Exchange Commission

Former Regional Director, Southeast Region

Federal Communications Commission

Former Bureau Chief, Public Safety · RADM (Ret.)

U.S. Department of the Treasury

Former CTO and Deputy Chief AI Officer

White House · HHS

Former Deputy Assistant to the President for Domestic Policy; Former Chief of Staff, HHS

Senior Fellow

Former Head of Technology Risk, PayPal and The Hartford

Former senior leaders from the Federal Reserve, the SEC, the FCC, the U.S. Treasury, and HHS guide the Center's work.



Build Trust. Enable Speed. Lead Responsibly.

Advancing AI oversight governance for boards, executives, and regulators.

cfaio.org

info@cfaio.org

Brian J. Allen, Esq., Executive Director, Center for AI Oversight